31

Luc Devroye
László Györfi
Gábor Lugosi
A Probabilistic Theory of Pattern Recognition

$$
\begin{aligned}
\text { Stochastic Mechanics } & \text { Applications of } \\
\text { Random Media } & \text { Mathematics } \\
\text { Signal Processing } & \text { Stochastic Modelling } \\
\text { and Image Synthesis } & \text { and Applied Probability }
\end{aligned}
$$

Mathematical Economics
Stochastic Optimization
Stochastic Control 31

Edited by I. Karatzas
M. Yor

Advisory Board P. Brémaud
E. Carlen
R. Dobrushin
W. Fleming
D. Geman
G. Grimmett
G. Papanicolaou
J. Scheinkman

## Springer

New York
Berlin
Heidelberg
Barcelona
Budapest
Hong Kong
London
Milan
Paris
Santa Clara
Singapore
Tokyo

1 Fleming/Rishel, Deterministic and Stochastic Optimal Control (1975)
2 Marchuk, Methods of Numerical Mathematics, Second Ed. (1982)
3 Balakrishnan, Applied Functional Analysis, Second Ed. (1981)
4 Borovkov, Stochastic Processes in Queueing Theory (1976)
5 Liptser/Shiryayev, Statistics of Random Processes I: General Theory (1977)
6 Liptser/Shiryayev, Statistics of Random Processes II: Applications (1978)
7 Vorob'ev, Game Theory: Lectures for Economists and Systems Scientists (1977)

8 Shiryayev, Optimal Stopping Rules (1978)
9 Ibragimov/Rozanov, Gaussian Random Processes (1978)
10 Wonham, Linear Multivariable Control: A Geometric Approach, Third Ed. (1985)

11 Hida, Brownian Motion (1980)
12 Hestenes, Conjugate Direction Methods in Optimization (1980)
13 Kallianpur, Stochastic Filtering Theory (1980)
14 Krylov, Controlled Diffusion Processes (1980)
15 Prabhu, Stochastic Storage Processes: Queues, Insurance Risk, and Dams (1980)

16 Ibragimov/Has'minskii, Statistical Estimation: Asymptotic Theory (1981)
17 Cesari, Optimization: Theory and Applications (1982)
18 Elliott, Stochastic Calculus and Applications (1982)
19 Marchuk/Shaidourov, Difference Methods and Their Extrapolations (1983)
20 Hijab, Stabilization of Control Systems (1986)
21 Protter, Stochastic Integration and Differential Equations (1990)
22 Benveniste/Métivier/Priouret, Adaptive Algorithms and Stochastic Approximations (1990)
23 Kloeden/Platen, Numerical Solution of Stochastic Differential Equations (1992)

24 Kushner/Dupuis, Numerical Methods for Stochastic Control Problems in Continuous Time (1992)
25 Fleming/Soner, Controlled Markov Processes and Viscosity Solutions (1993)

26 Baccelli/Brémaud, Elements of Queueing Theory (1994)
27 Winkler, Image Analysis, Random Fields, and Dynamic Monte Carlo Methods: An Introduction to Mathematical Aspects (1994)
28 Kalpazidou, Cycle Representations of Markov Processes (1995)
29 Elliott/Aggoun/Moore, Hidden Markov Models: Estimation and Control (1995)

30 Hernández-Lerma/Lasserre, Discrete-Time Markov Control Processes: Basic Optimality Criteria (1996)
31 Devroye/Györfi/Lugosi, A Probabilistic Theory of Pattern Recognition (1996)
32 Maitra/Sudderth, Discrete Gambling and Stochastic Games (1996)

## Luc Devroye László Györfi Gábor Lugosi

# A Probabilistic Theory of Pattern Recognition 

With 99 Figures

Luc Devroye
School of Computer Science
McGill University
Montreal, Quebec, H3A 2A7
Canada

László Györfi
Gábor Lugosi
Department of Mathematics and Computer Science
Technical University of Budapest
Budapest
Hungary

## Managing Editors

I. Karatzas

Department of Statistics
Columbia University
New York, NY 10027, USA
M. Yor

CNRS, Laboratoire de Probabilités
Université Pierre et Marie Curie
4, Place Jussieu, Tour 56
F-75252 Paris Cedex 05, France

Mathematics Subject Classification (1991): 68T10, $68 \mathrm{~T} 05,62 \mathrm{G} 07,62 \mathrm{H} 30$

Library of Congress Cataloging-in-Publication Data
Devroye, Luc.
A probabilistic theory of pattern recognition/Luc Devroye, László Györfi, Gábor Lugosi. p. cm.

Includes bibliographical references and index.
ISBN 0-387-94618-7 (hardcover)

1. Pattern perception. 2. Probabilities. I. Györfi, László.
II. Lugosi, Gábor. III. Title.

Q327.D5 1996
$003^{\prime} .52^{\prime} 015192-\mathrm{dc} 20$ 95-44633
Printed on acid-free paper.
© 1996 Springer-Veriag New York, Inc.
All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer-Verlag New York, Inc., 175 Fifth Avenue, New York, NY 10010, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.
The use of general descriptive names, trade names, trademarks, etc., in this publication, even if the former are not especially identified, is not to be taken as a sign that such names, as understood by the Trade Marks and Merchandise Marks Act, may accordingly be used freely by anyone.

Production managed by Francine McNeill; manufacturing supervised by Jeffrey Taub.
Photocomposed copy prepared using Springer's svsing.sty macro.
Printed and bound by Braun-Brumfield, Inc., Ann Arbor, MI.
Printed in the United States of America.
98765432 (Corrected second printing, 1997)
ISBN 0-387-94618-7 Springer-Verlag New York Berlin Heidelberg SPIN 10565785

## Preface

Life is just a long random walk. Things are created because the circumstances happen to be right. More often than not, creations, such as this book, are accidental. Nonparametric estimation came to life in the fifties and sixties and started developing at a frenzied pace in the late sixties, engulfing pattern recognition in its growth. In the mid-sixties, two young men, Tom Cover and Peter Hart, showed the world that the nearest neighbor rule in all its simplicity was guaranteed to err at most twice as often as the best possible discrimination method. Tom's results had a profound influence on Terry Wagner, who became a Professor at the University of Texas at Austin and brought probabilistic rigor to the young field of nonparametric estimation. Around 1971, Vapnik and Chervonenkis started publishing a revolutionary series of papers with deep implications in pattern recognition, but their work was not well known at the time. However, Tom and Terry had noticed the potential of the work, and Terry asked Luc Devroye to read that work in preparation for his Ph.D. dissertation at the University of Texas. The year was 1974. Luc ended up in Texas quite by accident thanks to a tip by his friend and fellow Belgian Willy Wouters, who matched him up with Terry. By the time Luc's dissertation was published in 1976, pattern recognition had taken off in earnest. On the theoretical side, important properties were still being discovered. In 1977, Stone stunned the nonparametric community by showing that there are nonparametric rules that are convergent for all distributions of the data. This is called distribution-free or universal consistency, and it is what makes nonparametric methods so attractive. Yet, very few researchers were concerned with universal consistency-one notable exception was Laci Györfi, who at that time worked in Budapest amid an energetic group of nonparametric specialists that included Sándor Csibi, József Fritz, and Pál Révész.

So, linked by a common vision, Luc and Laci decided to join forces in the early eighties. In 1982, they wrote six chapters of a book on nonparametric regression function estimation, but these were never published. In fact, the notes are still in drawers in their offices today. They felt that the subject had not matured yet. A book on nonparametric density estimation saw the light in 1985. Unfortunately, as true baby-boomers, neither Luc nor Laci had the time after 1985 to write a text on nonparametric pattern recognition. Enter Gábor Lugosi, who obtained his doctoral degree under Laci's supervision in 1991. Gábor had prepared a set of rough course notes on the subject around 1992 and proposed to coordinate the project-this book-in 1993. With renewed energy, we set out to write the book that we should have written at least ten years ago. Discussions and work sessions were held in Budapest, Montreal, Leuven, and Louvain-La-Neuve. In Leuven, our gracious hosts were Ed van der Meulen and Jan Beirlant, and in Louvain-La-Neuve, we were gastronomically and spiritually supported by Léopold Simar and Irène Gijbels. We thank all of them. New results accumulated, and we had to resist the temptation to publish these in journals. Finally, in May 1995, the manuscript had bloated to such extent that it had to be sent to the publisher, for otherwise it would have become an encyclopedia. Some important unanswered questions were quickly turned into masochistic exercises or wild conjectures. We will explain subject selection, classroom use, chapter dependence, and personal viewpoints in the Introduction. We do apologize, of course, for all remaining errors.

We were touched, influenced, guided, and taught by many people. Terry Wagner's rigor and taste for beautiful nonparametric problems have infected us for life. We thank our past and present coauthors on nonparametric papers, Alain Berlinet, Michel Broniatowski, Ricardo Cao, Paul Deheuvels, András Faragó, Adam Krzyżak, Tamás Linder, Andrew Nobel, Mirek Pawlak, Igor Vajda, Harro Walk, and Ken Zeger. Tamás Linder read most of the book and provided invaluable feedback. His help is especially appreciated. Several chapters were critically read by students in Budapest. We thank all of them, especially András Antos, Miklós Csűrös, Balázs Kégl, István Páli, and Márti Pintér. Finally, here is an alphabetically ordered list of friends who directly or indirectly contributed to our knowledge and love of nonparametrics: Andrew and Roger Barron, Denis Bosq, Prabhir Burman, Tom Cover, Antonio Cuevas, Pierre Devijver, Ricardo Fraiman, Ned Glick, Wenceslao Gonzalez-Manteiga, Peter Hall, Eiichi Isogai, Ed Mack, Arthur Nádas, Georg Pflug, George Roussas, Winfried Stute, Tamás Szabados, Godfried Toussaint, Sid Yakowitz, and Yannis Yatracos.

Gábor diligently typed the entire manuscript and coordinated all contributions. He became quite a $T_{E} X p e r t$ in the process. Several figures were made by idraw and xfig by Gábor and Luc. Most of the drawings were directly programmed in PostScript by Luc and an undergraduate student at McGill University, Hisham Petry, to whom we are grateful. For Gábor, this book comes at the beginning of his career. Unfortunately, the other two authors are not so lucky. As both Luc and Laci felt that they would probably not write another book on nonparametric pattern recognition-the random walk must go on-they decided to put their general
view of the subject area on paper while trying to separate the important from the irrelevant. Surely, this has contributed to the length of the text.

So far, our random excursions have been happy ones. Coincidentally, Luc is married to Bea, the most understanding woman in the world, and happens to have two great daughters, Natasha and Birgit, who do not stray off their random courses. Similarly, Laci has an equally wonderful wife, Kati, and two children with steady compasses, Kati and János. During the preparations of this book, Gábor met a wonderful girl, Arrate. They have recently decided to tie their lives together.

On the less amorous and glamorous side, we gratefully acknowledge the research support of NSERC CANADA, FCAR QUEBEC, OTKA HUNGARY, and the exchange program between the Hungarian Academy of Sciences and the Royal Belgian Academy of Sciences. Early versions of this text were tried out in some classes at the Technical University of Budapest, Katholieke Universiteit Leuven, Universität Stuttgart, and Université Montpellier II. We would like to thank those students for their help in making this a better book.

Montreal, Quebec, Canada
Luc Devroye
Budapest, Hungary
Laci Györfi
Budapest, Hungary
Gábor Lugosi

## Contents

Preface ..... v
1 Introduction ..... 1
2 The Bayes Error ..... 9
2.1 The Bayes Problem ..... 9
2.2 A Simple Example ..... 11
2.3 Another Simple Example ..... 12
2.4 Other Formulas for the Bayes Risk ..... 14
2.5 Plug-In Decisions ..... 15
2.6 Bayes Error Versus Dimension ..... 17
Problems and Exercises ..... 18
3 Inequalities and Alternate Distance Measures ..... 21
3.1 Measuring Discriminatory Information ..... 21
3.2 The Kolmogorov Variational Distance ..... 22
3.3 The Nearest Neighbor Error ..... 22
3.4 The Bhattacharyya Affinity ..... 23
3.5 Entropy ..... 25
3.6 Jeffreys' Divergence ..... 27
$3.7 \quad F$-Errors ..... 28
3.8 The Mahalanobis Distance ..... 30
$3.9 f$-Divergences ..... 31
Problems and Exercises ..... 35
4 Linear Discrimination ..... 39
4.1 Univariate Discrimination and Stoller Splits ..... 40
4.2 Linear Discriminants ..... 44
4.3 The Fisher Linear Discriminant ..... 46
4.4 The Normal Distribution ..... 47
4.5 Empirical Risk Minimization ..... 49
4.6 Minimizing Other Criteria ..... 54
Problems and Exercises ..... 56
5 Nearest Neighbor Rules ..... 61
5.1 Introduction ..... 61
5.2 Notation and Simple Asymptotics ..... 63
5.3 Proof of Stone's Lemma ..... 66
5.4 The Asymptotic Probability of Error ..... 69
5.5 The Asymptotic Error Probability of Weighted Nearest Neighbor Rules ..... 71
$5.6 \quad k$-Nearest Neighbor Rules: Even $k$ ..... 74
5.7 Inequalities for the Probability of Error ..... 75
5.8 Behavior When $L^{*}$ Is Small ..... 78
5.9 Nearest Neighbor Rules When $L^{*}=0$ ..... 80
5.10 Admissibility of the Nearest Neighbor Rule ..... 81
5.11 The ( $k, l$ )-Nearest Neighbor Rule ..... 81
Problems and Exercises ..... 83
6 Consistency ..... 91
6.1 Universal Consistency ..... 91
6.2 Classification and Regression Estimation ..... 92
6.3 Partitioning Rules ..... 94
6.4 The Histogram Rule ..... 95
6.5 Stone's Theorem ..... 97
6.6 The $k$-Nearest Neighbor Rule ..... 100
6.7 Classification Is Easier Than Regression Function Estimation ..... 101
6.8 Smart Rules ..... 106
Problems and Exercises ..... 107
7 Slow Rates of Convergence ..... 111
7.1 Finite Training Sequence ..... 111
7.2 Slow Rates ..... 113
Problems and Exercises ..... 118
8 Error Estimation ..... 121
8.1 Error Counting ..... 121
8.2 Hoeffding's Inequality ..... 122
8.3 Error Estimation Without Testing Data ..... 124
8.4 Selecting Classifiers ..... 125
8.5 Estimating the Bayes Error ..... 128
Problems and Exercises ..... 129
9 The Regular Histogram Rule ..... 133
9.1 The Method of Bounded Differences ..... 133
9.2 Strong Universal Consistency ..... 138
Problems and Exercises ..... 142
10 Kernel Rules ..... 147
10.1 Consistency ..... 149
10.2 Proof of the Consistency Theorem ..... 153
10.3 Potential Function Rules ..... 159
Problems and Exercises ..... 161
11 Consistency of the $k$-Nearest Neighbor Rule ..... 169
11.1 Strong Consistency ..... 170
11.2 Breaking Distance Ties ..... 174
11.3 Recursive Methods ..... 176
11.4 Scale-Invariant Rules ..... 177
11.5 Weighted Nearest Neighbor Rules ..... 178
11.6 Rotation-Invariant Rules ..... 179
11.7 Relabeling Rules ..... 180
Problems and Exercises ..... 182
12 Vapnik-Chervonenkis Theory ..... 187
12.1 Empirical Error Minimization ..... 187
12.2 Fingering ..... 191
12.3 The Glivenko-Cantelli Theorem ..... 192
12.4 Uniform Deviations of Relative Frequencies from Probabilities ..... 196
12.5 Classifier Selection ..... 199
12.6 Sample Complexity ..... 201
12.7 The Zero-Error Case ..... 202
12.8 Extensions ..... 206
Problems and Exercises ..... 208
13 Combinatorial Aspects of Vapnik-Chervonenkis Theory ..... 215
13.1 Shatter Coefficients and VC Dimension ..... 215
13.2 Shatter Coefficients of Some Classes ..... 219
13.3 Linear and Generalized Linear Discrimination Rules ..... 224
13.4 Convex Sets and Monotone Layers ..... 226
Problems and Exercises ..... 229
14 Lower Bounds for Empirical Classifier Selection ..... 233
14.1 Minimax Lower Bounds ..... 234
14.2 The Case $L_{\mathcal{C}}=0$ ..... 234
14.3 Classes with Infinite VC Dimension ..... 238
14.4 The Case $L_{\mathcal{C}}>0$ ..... 239
14.5 Sample Complexity ..... 245
Problems and Exercises ..... 247
15 The Maximum Likelihood Principle ..... 249
15.1 Maximum Likelihood: The Formats ..... 249
15.2 The Maximum Likelihood Method: Regression Format ..... 250
15.3 Consistency ..... 253
15.4 Examples ..... 256
15.5 Classical Maximum Likelihood: Distribution Format ..... 260
Problems and Exercises ..... 261
16 Parametric Classification ..... 263
16.1 Example: Exponential Families ..... 266
16.2 Standard Plug-In Rules ..... 267
16.3 Minimum Distance Estimates ..... 270
16.4 Empirical Error Minimization ..... 275
Problems and Exercises ..... 276
17 Generalized Linear Discrimination ..... 279
17.1 Fourier Series Classification ..... 280
17.2 Generalized Linear Classification ..... 285
Problems and Exercises ..... 287
18 Complexity Regularization ..... 289
18.1 Structural Risk Minimization ..... 290
18.2 Poor Approximation Properties of VC Classes ..... 297
18.3 Simple Empirical Covering ..... 297
Problems and Exercises ..... 300
19 Condensed and Edited Nearest Neighbor Rules ..... 303
19.1 Condensed Nearest Neighbor Rules ..... 303
19.2 Edited Nearest Neighbor Rules ..... 309
19.3 Sieves and Prototypes ..... 309
Problems and Exercises ..... 312
20 Tree Classifiers ..... 315
20.1 Invariance ..... 318
20.2 Trees with the $X$-Property ..... 319
20.3 Balanced Search Trees ..... 322
20.4 Binary Search Trees ..... 326
20.5 The Chronological k-d Tree ..... 328
20.6 The Deep k-d Tree ..... 332
20.7 Quadtrees ..... 333
20.8 Best Possible Perpendicular Splits ..... 334
20.9 Splitting Criteria Based on Impurity Functions ..... 336
20.10 A Consistent Splitting Criterion ..... 340
20.11 BSP Trees ..... 341
20.12 Primitive Selection ..... 343
20.13 Constructing Consistent Tree Classifiers ..... 346
20.14 A Greedy Classifier ..... 348
Problems and Exercises ..... 357
21 Data-Dependent Partitioning ..... 363
21.1 Introduction ..... 363
21.2 A Vapnik-Chervonenkis Inequality for Partitions ..... 364
21.3 Consistency ..... 368
21.4 Statistically Equivalent Blocks ..... 372
21.5 Partitioning Rules Based on Clustering ..... 377
21.6 Data-Based Scaling ..... 381
21.7 Classification Trees ..... 383
Problems and Exercises ..... 383
22 Splitting the Data ..... 387
22.1 The Holdout Estimate ..... 387
22.2 Consistency and Asymptotic Optimality ..... 389
22.3 Nearest Neighbor Rules with Automatic Scaling ..... 391
22.4 Classification Based on Clustering ..... 392
22.5 Statistically Equivalent Blocks ..... 393
22.6 Binary Tree Classifiers ..... 394
Problems and Exercises ..... 395
23 The Resubstitution Estimate ..... 397
23.1 The Resubstitution Estimate ..... 397
23.2 Histogram Rules ..... 399
23.3 Data-Based Histograms and Rule Selection ..... 403
Problems and Exercises ..... 405
24 Deleted Estimates of the Error Probability ..... 407
24.1 A General Lower Bound ..... 408
24.2 A General Upper Bound for Deleted Estimates ..... 411
24.3 Nearest Neighbor Rules ..... 413
24.4 Kernel Rules ..... 415
24.5 Histogram Rules ..... 417
Problems and Exercises ..... 419
25 Automatic Kernel Rules ..... 423
25.1 Consistency ..... 424
25.2 Data Splitting ..... 428
25.3 Kernel Complexity ..... 431
25.4 Multiparameter Kernel Rules ..... 435
25.5 Kernels of Infinite Complexity ..... 436
25.6 On Minimizing the Apparent Error Rate ..... 439
25.7 Minimizing the Deleted Estimate ..... 441
25.8 Sieve Methods ..... 444
25.9 Squared Error Minimization ..... 445
Problems and Exercises ..... 446
26 Automatic Nearest Neighbor Rules ..... 451
26.1 Consistency ..... 451
26.2 Data Splitting ..... 452
26.3 Data Splitting for Weighted NN Rules ..... 453
26.4 Reference Data and Data Splitting ..... 454
26.5 Variable Metric NN Rules ..... 455
26.6 Selection of $k$ Based on the Deleted Estimate ..... 457
Problems and Exercises ..... 458
27 Hypercubes and Discrete Spaces ..... 461
27.1 Multinomial Discrimination ..... 461
27.2 Quantization ..... 464
27.3 Independent Components ..... 466
27.4 Boolean Classifiers ..... 468
27.5 Series Methods for the Hypercube ..... 470
27.6 Maximum Likelihood ..... 472
27.7 Kernel Methods ..... 474
Problems and Exercises ..... 474
28 Epsilon Entropy and Totally Bounded Sets ..... 479
28.1 Definitions ..... 479
28.2 Examples: Totally Bounded Classes ..... 480
28.3 Skeleton Estimates ..... 482
28.4 Rate of Convergence ..... 485
Problems and Exercises ..... 486
29 Uniform Laws of Large Numbers ..... 489
29.1 Minimizing the Empirical Squared Error ..... 489
29.2 Uniform Deviations of Averages from Expectations ..... 490
29.3 Empirical Squared Error Minimization ..... 493
29.4 Proof of Theorem 29.1 ..... 494
29.5 Covering Numbers and Shatter Coefficients ..... 496
29.6 Generalized Linear Classification ..... 501
Problems and Exercises ..... 505
30 Neural Networks ..... 507
30.1 Multilayer Perceptrons ..... 507
30.2 Arrangements ..... 511
30.3 Approximation by Neural Networks ..... 517
30.4 VC Dimension ..... 521
$30.5 \quad L_{1}$ Error Minimization ..... 526
30.6 The Adaline and Padaline ..... 531
30.7 Polynomial Networks ..... 532
30.8 Kolmogorov-Lorentz Networks and Additive Models ..... 534
30.9 Projection Pursuit ..... 538
30.10 Radial Basis Function Networks ..... 540
Problems and Exercises ..... 542
31 Other Error Estimates ..... 549
31.1 Smoothing the Error Count ..... 549
31.2 Posterior Probability Estimates ..... 554
31.3 Rotation Estimate ..... 556
31.4 Bootstrap ..... 556
Problems and Exercises ..... 559
32 Feature Extraction ..... 561
32.1 Dimensionality Reduction ..... 561
32.2 Transformations with Small Distortion ..... 567
32.3 Admissible and Sufficient Transformations ..... 569
Problems and Exercises ..... 572
Appendix ..... 575
A. 1 Basics of Measure Theory ..... 575
A. 2 The Lebesgue Integral ..... 576
A. 3 Denseness Results ..... 579
A. 4 Probability ..... 581
A. 5 Inequalities ..... 582
A. 6 Convergence of Random Variables ..... 584
A. 7 Conditional Expectation ..... 585
A. 8 The Binomial Distribution ..... 586
A. 9 The Hypergeometric Distribution ..... 589
A. 10 The Multinomial Distribution ..... 589
A. 11 The Exponential and Gamma Distributions ..... 590
A. 12 The Multivariate Normal Distribution ..... 590
Notation ..... 591
References ..... 593
Author Index ..... 619
Subject Index ..... 627

## 1

## Introduction

Pattern recognition or discrimination is about guessing or predicting the unknown nature of an observation, a discrete quantity such as black or white, one or zero, sick or healthy, real or fake. An observation is a collection of numerical measurements such as an image (which is a sequence of bits, one per pixel), a vector of weather data, an electrocardiogram, or a signature on a check suitably digitized. More formally, an observation is a $d$-dimensional vector $x$. The unknown nature of the observation is called a class. It is denoted by $y$ and takes values in a finite set $\{1,2, \ldots, M\}$. In pattern recognition, one creates a function $g(x): \mathcal{R}^{d} \rightarrow$ $\{1, \ldots, M\}$ which represents one's guess of $y$ given $x$. The mapping $g$ is called a classifier. Our classifier errs on $x$ if $g(x) \neq y$.

How one creates a rule $g$ depends upon the problem at hand. Experts can be called for medical diagnoses or earthquake predictions--they try to mold $g$ to their own knowledge and experience, often by trial and error. Theoretically, each expert operates with a built-in classifier $g$, but describing this $g$ explicitly in mathematical form is not a sinecure. The sheer magnitude and richness of the space of $x$ may defeat even the best expert-it is simply impossible to specify $g$ for all possible $x$ 's one is likely to see in the future. We have to be prepared to live with imperfect classifiers. In fact, how should we measure the quality of a classifier? We can't just dismiss a classifier just because it misclassifies a particular $x$. For one thing, if the observation does not fully describe the underlying process (that is, if $y$ is not a deterministic function of $x$ ), it is possible that the same $x$ may give rise to two different $y$ 's on different occasions. For example, if we just measure water content of a person's body, and we find that the person is dehydrated, then the cause (the class) may range from a low water intake in hot weather to severe diarrhea. Thus, we introduce a probabilistic setting, and let $(X, Y)$ be an $\mathcal{R}^{d} \times\{1, \ldots, M\}$-valued
random pair. The distribution of ( $X, Y$ ) describes the frequency of encountering particular pairs in practice. An error occurs if $g(X) \neq Y$, and the probability of error for a classifier $g$ is

$$
L(g)=\mathbf{P}\{g(X) \neq Y\} .
$$

There is a best possible classifier, $g^{*}$, which is defined by

$$
g^{*}=\underset{g: \mathcal{R}^{d} \rightarrow\{1, \ldots, M\}}{\arg \min } \mathbf{P}\{g(X) \neq Y\} .
$$

Note that $g^{*}$ depends upon the distribution of $(X, Y)$. If this distribution is known, $g^{*}$ may be computed. The problem of finding $g^{*}$ is Bayes' problem, and the classifier $g^{*}$ is called the Bayes classifier (or the Bayes rule). The minimal probability of error is called the Bayes error and is denoted by $L^{*}=L\left(g^{*}\right)$. Mostly, the distribution of $(X, Y)$ is unknown, so that $g^{*}$ is unknown too.

We do not consult an expert to try to reconstruct $g^{*}$, but have access to a good database of pairs ( $X_{i}, Y_{i}$ ), $1 \leq i \leq n$, observed in the past. This database may be the result of experimental observation (as for meteorological data, fingerprint data, ECG data, or handwritten characters). It could also be obtained through an expert or a teacher who filled in the $Y_{i}$ 's after having seen the $X_{i}$ 's. To find a classifier $g$ with a small probability of error is hopeless unless there is some assurance that the ( $X_{i}, Y_{i}$ )'s jointly are somehow representative of the unknown distribution. We shall assume in this book that $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, the data, is a sequence of independent identically distributed (i.i.d.) random pairs with the same distribution as that of $(X, Y)$. This is a very strong assumption indeed. However, some theoretical results are emerging that show that classifiers based on slightly dependent data pairs and on i.i.d. data pairs behave roughly the same. Also, simple models are easier to understand are more amenable to interpretation.

A classifier is constructed on the basis of $X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}$ and is denoted by $g_{n} ; Y$ is guessed by $g_{n}\left(X ; X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right)$. The process of constructing $g_{n}$ is called learning, supervised learning, or learning with a teacher. The performance of $g_{n}$ is measured by the conditional probability of error

$$
\left.L_{n}=L\left(g_{n}\right)=\mathbf{P}\left\{g_{n}\left(X ; X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right) \neq Y \mid X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right)\right\} .
$$

This is a random variable because it depends upon the data. So, $L_{n}$ averages over the distribution of $(X, Y)$, but the data is held fixed. Averaging over the data as well would be unnatural, because in a given application, one has to live with the data at hand. It would be marginally useful to know the number $\mathbf{E} L_{n}$ as this number would indicate the quality of an average data sequence, not your data sequence. This text is thus about $L_{n}$, the conditional probability of error.

An individual mapping $g_{n}: \mathcal{R}^{d} \times\left\{\mathcal{R}^{d} \times\{1, \ldots, M\}\right\}^{n} \longrightarrow\{1, \ldots, M\}$ is still called a classifier. A sequence $\left\{g_{n}, n \geq 1\right\}$ is called a (discrimination) rule. Thus, classifiers are functions, and rules are sequences of functions.

A novice might ask simple questions like this: How does one construct a good classifier? How good can a classifier be? Is classifier A better than classifier B? Can we estimate how good a classifier is? What is the best classifier? This book partially
answers such simple questions. A good deal of energy is spent on the mathematical formulations of the novice's questions. For us, a rule-not a classifier-is good if it is consistent, that is, if

$$
\lim _{n \rightarrow \infty} \mathbf{E} L_{n}=L^{*}
$$

or equivalently, if $L_{n} \rightarrow L^{*}$ in probability as $n \rightarrow \infty$. We assume that the reader has a good grasp of the basic elements of probability, including notions such as convergence in probability, strong laws of large numbers for averages, and conditional probability. A selection of results and definitions that may be useful for this text is given in the Appendix. A consistent rule guarantees us that taking more samples essentially suffices to roughly reconstruct the unknown distribution of ( $X, Y$ ) because $L_{n}$ can be pushed as close as desired to $L^{*}$. In other words, infinite amounts of information can be gleaned from finite samples. Without this guarantee, we would not be motivated to take more samples. We should be careful and not impose conditions on $(X, Y)$ for the consistency of a rule, because such conditions may not be verifiable. If a rule is consistent for all distributions of ( $X, Y$ ), it is said to be universally consistent.

Interestingly, until 1977, it was not known if a universally consistent rule existed. All pre-1977 consistency results came with restrictions on ( $X, Y$ ). In 1977, Stone showed that one could just take any $k$-nearest neighbor rule with $k=k(n) \rightarrow \infty$ and $k / n \rightarrow 0$. The $k$-nearest neighbor classifier $g_{n}(x)$ takes a majority vote over the $Y_{i}$ 's in the subset of $k$ pairs $\left(X_{i}, Y_{i}\right)$ from $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ that have the smallest values for $\left\|X_{i}-x\right\|$ (i.e., for which $X_{i}$ is closest to $x$ ). Since Stone's proof of the universal consistency of the $k$-nearest neighbor rule, several other rules have been shown to be universally consistent as well. This book stresses universality and hopefully gives a reasonable account of the developments in this direction.

Probabilists may wonder why we did not use convergence with probability one in our definition of consistency. Indeed, strong consistency-convergence of $L_{n}$ to $L^{*}$ with probability one-implies convergence for almost every sample as it grows. Fortunately, for most well-behaved rules, consistency and strong consistency are equivalent. For example, for the $k$-nearest neighbor rule, $k \rightarrow \infty$ and $k / n \rightarrow 0$ together imply $L_{n} \rightarrow L^{*}$ with probability one. The equivalence will be dealt with, but it will not be a major focus of attention. Most, if not all, equivalence results are based upon some powerful concentration inequalities such as McDiarmid's. For example, we will be able to show that for the $k$-nearest neighbor rule, there exists a number $c>0$, such that for all $\epsilon>0$, there exists $N(\epsilon)>0$ depending upon the distribution of $(X, Y)$, such that

$$
\mathbf{P}\left\{L_{n}-L^{*}>\epsilon\right\} \leq e^{-c n \epsilon^{2}}, n \geq N(\epsilon)
$$

This illustrates yet another focus of the book-inequalities. Whenever possible, we make a case or conclude a proof via explicit inequalities. Various parameters can be substituted in these inequalities to allow the user to draw conclusions regarding sample size or to permit identification of the most important parameters.

The material in the book is often technical and dry. So, to stay focused on the main issues, we keep the problem simple:
A. We only deal with binary classification $(M=2)$. The class $Y$ takes values in $\{0,1\}$, and a classifier $g_{n}$ is a mapping: $\mathcal{R}^{d} \times\left\{\mathcal{R}^{d} \times\{0,1\}\right\}^{n} \rightarrow\{0,1\}$.
B. We only consider i.i.d. data sequences. We also disallow active learning, a set-up in which the user can select the $X_{i}$ 's deterministically.
C. We do not consider infinite spaces. For example, $X$ cannot be a random function such as a cardiogram. $X$ must be a $\mathcal{R}^{d}$-valued random vector. The reader should be aware that many results given here may be painlessly extended to certain metric spaces of infinite dimension.
Let us return to our novice's questions. We know that there are good rules, but just how good can a classifier be? Obviously, $L_{n} \geq L^{*}$ in all cases. It is thus important to know $L^{*}$ or to estimate it, for if $L^{*}$ is large, any classifier, including yours, will perform poorly. But even if $L^{*}$ were zero, $L_{n}$ could still be large. Thus, it would be nice to have explicit inequalities for probabilities such as

$$
\mathbf{P}\left\{L_{n} \geq L^{*}+\epsilon\right\}
$$

However, such inequalities must necessarily depend upon the distribution of $(X, Y)$. That is, for any rule,
$\liminf _{n \rightarrow \infty} \sup _{\text {all distributions of }(X, Y) \text { with } L^{*}+\epsilon<1 / 2} \mathbf{P}\left\{L_{n} \geq L^{*}+\epsilon\right\}>0$.
Universal rate of convergence guarantees do not exist. Rate of convergence studies must involve certain subclasses of distributions of $(X, Y)$. For this reason, with few exceptions, we will steer clear of the rate of convergence quicksand.

Even if there are no universal performance guarantees, we might still be able to satisfy our novice's curiosity if we could satisfactorily estimate $L_{n}$ for the rule at hand by a function $\widehat{L}_{n}$ of the data. Such functions are called error estimates. For example, for the $k$-nearest neighbor classifier, we could use the deleted estimate

$$
\widehat{L}_{n}=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{g_{n i}\left(X_{i}\right) \neq Y_{i}\right\}}
$$

where $g_{n i}\left(X_{i}\right)$ classifies $X_{i}$ by the $k$-nearest neighbor method based upon the data $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ with $\left(X_{i}, Y_{i}\right)$ deleted. If this is done, we have a distributionfree inequality

$$
\mathbf{P}\left\{\left|\widehat{L}_{n}-L_{n}\right|>\epsilon\right\} \leq \frac{6 k+1}{n \epsilon^{2}}
$$

(the Rogers-Wagner inequality), provided that distance ties are broken in an appropriate manner. In other words, without knowing the distribution of $(X, Y)$, we can state with a certain confidence that $L_{n}$ is contained in $\left[\widehat{L}_{n}-\epsilon, \widehat{L_{n}}+\epsilon\right]$. Thus, for many classifiers, it is indeed possible to estimate $L_{n}$ from the data at hand. However, it is impossible to estimate $L^{*}$ universally well: for any $n$, and any estimate of $L^{*}$ based upon the data sequence, there always exists a distribution of $(X, Y)$ for which the estimate is arbitrarily poor.

Can we compare rules $\left\{g_{n}\right\}$ and $\left\{g_{n}^{\prime}\right\}$ ? Again, the answer is negative: there exists no "best" classifier (or superclassifier), as for any rule $\left\{g_{n}\right\}$, there exists a distribution of $(X, Y)$ and another rule $\left\{g_{n}^{\prime}\right\}$ such that for all $n, \mathbf{E}\left\{L\left(g_{n}^{\prime}\right)\right\}<\mathbf{E}\left\{L\left(g_{n}\right)\right\}$. If there had been a universally best classifier, this book would have been unnecessary: we would all have to use it all the time. This nonexistence implies that the debate between practicing pattern recognizers will never end and that simulations on particular examples should never be used to compare classifiers. As an example, consider the 1-nearest neighbor rule, a simple but not universally consistent rule. Yet, among all $k$-nearest neighbor classifiers, the 1-nearest neighbor classifier is admissible--there are distributions for which its expected probability of error is better than for any $k$-nearest neighbor classifier with $k>1$. So, it can never be totally dismissed. Thus, we must study all simple rules, and we will reserve many pages for the nearest neighbor rule and its derivatives. We will for example prove the Cover-Hart inequality (Cover and Hart, 1967) which states that for all distributions of ( $X, Y$ ),

$$
\limsup _{n \rightarrow \infty} \mathbf{E} L_{n} \leq 2 L^{*}
$$

where $L_{n}$ is the probability of error with the 1-nearest neighbor rule. As $L^{*}$ is usually small (for otherwise, you would not want to do discrimination), $2 L^{*}$ is small too, and the 1-nearest neighbor rule will do just fine.

The nonexistence of a best classifier may disappoint our novice. However, we may change the setting somewhat and limit the classifiers to a certain class $\mathcal{C}$, such as all $k$-nearest neighbor classifiers with all possible values for $k$. Is it possible to select the best classifier from this class? Phrased in this manner, we cannot possibly do better than

$$
L \stackrel{\text { def }}{=} \inf _{g_{n} \in \mathcal{C}} \mathbf{P}\left\{g_{n}(X) \neq Y\right\}
$$

Typically, $L>L^{*}$. Interestingly, there is a general paradigm for picking classifiers from $\mathcal{C}$ and to obtain universal performance guarantees. It uses empirical risk minimization, a method studied in great detail in the work of Vapnik and Chervonenkis (1971). For example, if we select $g_{n}$ from $\mathcal{C}$ by minimizing

$$
\frac{1}{n} \sum_{i=1}^{n} I_{\left\{g_{n}\left(X_{i}\right) \neq Y_{i}\right\}}
$$

then the corresponding probability of error $L_{n}$ satisfies the following inequality for all $\epsilon>0$ :

$$
\mathbf{P}\left\{L_{n}>L+\epsilon\right\} \leq 8\left(n^{V}+1\right) e^{-n \epsilon^{2} / 128}
$$

Here $V>0$ is an integer depending upon the massiveness of $\mathcal{C}$ only. $V$ is called the vC dimension of $\mathcal{C}$ and may be infinite for large classes $\mathcal{C}$. For sufficiently restricted classes $\mathcal{C}, V$ is finite and the explicit universal bound given above can be used to obtain performance guarantees for the selected $g_{n}$ (relative to $L$, not $L^{*}$ ). The bound above is only valid if $\mathcal{C}$ is independent of the data pairs $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$. Fixed classes such as all classifiers that decide 1 on a halfspace and 0 on its complement are fine. We may also sample $m$ more pairs (in addition to the $n$ pairs
already present), and use the $n$ pairs as above to select the best $k$ for use in the $k$-nearest neighbor classifier based on the $m$ pairs. As we will see, the selected rule is universally consistent if both $m$ and $n$ diverge and $n / \log m \rightarrow \infty$. And we have automatically solved the problem of picking $k$. Recall that Stone's universal consistency theorem only told us to pick $k=o(m)$ and to let $k \rightarrow \infty$, but it does not tell us whether $k \approx m^{0.01}$ is preferable over $k \approx m^{0.99}$. Empirical risk minimization produces a random data-dependent $k$ that is not even guaranteed to tend to infinity or to be $o(m)$, yet the selected rule is universally consistent.

We offer virtually no help with algorithms as in standard texts, with two notable exceptions. Ease of computation, storage, and interpretation has spurred the development of certain rules. For example, tree classifiers construct a tree for storing the data, and partition $\mathcal{R}^{d}$ by certain cuts that are typically perpendicular to a coordinate axis. We say that a coordinate axis is "cut." Such classifiers have obvious computational advantages, and are amenable to interpretation-the components of the vector $X$ that are cut at the early stages of the tree are most crucial in reaching a decision. Expert systems, automated medical diagnosis, and a host of other recognition rules use tree classification. For example, in automated medical diagnosis, one may first check a patient's pulse (component \#1). If this is zero, the patient is dead. If it is below 40 , the patient is weak. The first component is cut twice. In each case, we may then consider another component, and continue the breakdown into more and more specific cases. Several interesting new universally consistent tree classifiers are described in Chapter 20.

The second group of classifiers whose development was partially based upon easy implementations is the class of neural network classifiers, descendants of Rosenblatt's perceptron (Rosenblatt, 1956). These classifiers have unknown parameters that must be trained or selected by the data, in the way we let the data pick $k$ in the $k$-nearest neighbor classifier. Most research papers on neural networks deal with the training aspect, but we will not. When we say "pick the parameters by empirical risk minimization," we will leave the important algorithmic complexity questions unanswered. Perceptrons divide the space by one hyperplane and attach decisions 1 and 0 to the two halfspaces. Such simple classifiers are not consistent except for a few distributions. This is the case, for example, when $X$ takes values on $\{0,1\}^{d}$ (the hypercube) and the components of $X$ are independent. Neural networks with one hidden layer are universally consistent if the parameters are well-chosen. We will see that there is also some gain in considering two hidden layers, but that it is not really necessary to go beyond two.

Complexity of the training algorithm-the phase in which a classifier $g_{n}$ is selected from $\mathcal{C}$-is of course important. Sometimes, one would like to obtain classifiers that are invariant under certain transformations. For example, the $k$ nearest neighbor classifier is not invariant under nonlinear transformations of the coordinate axes. This is a drawback as components are often measurements in an arbitrary scale. Switching to a logarithmic scale or stretching a scale out by using Fahrenheit instead of Celsius should not affect good discrimination rules. There exist variants of the $k$-nearest neighbor rule that have the given invariance. In character recognition, sometimes all components of a vector $X$ that represents
a character are true measurements involving only vector differences between selected points such as the leftmost and rightmost points, the geometric center, the weighted center of all black pixels, the topmost and bottommost points. In this case, the scale has essential information, and invariance with respect to changes of scale would be detrimental. Here however, some invariance with respect to orthonormal rotations is healthy.

We follow the standard notation from textbooks on probability. Thus, random variables are uppercase characters such as $X, Y$, and $Z$. Probability measures are denoted by greek letters such as $\mu$ and $\nu$. Numbers and vectors are denoted by lowercase letters such as $a, b, c, x$, and $y$. Sets are also denoted by roman capitals, but there are obvious mnemonics: $S$ denotes a sphere, $B$ denotes a Borel set, and so forth. If we need many kinds of sets, we will typically use the beginning of the alphabet $(A, B, C)$. Most functions are denoted by $f, g, \phi$, and $\psi$. Calligraphic letters such as $\mathcal{A}, \mathcal{C}$, and $\mathcal{F}$ are used to denote classes of functions or sets. A short list of frequently used symbols is found at the end of the book.

At the end of this chapter, you will find a directed acyclic graph that describes the dependence between chapters. Clearly, prospective teachers will have to select small subsets of chapters. All chapters, without exception, are unashamedly theoretical. We did not scar the pages with backbreaking simulations or quick-and-dirty engineering solutions. The methods gleaned from this text must be supplemented with a healthy dose of engineering savvy. Ideally, students should have a companion text filled with beautiful applications such as automated virus recognition, telephone eavesdropping language recognition, voice recognition in security systems, fingerprint recognition, or handwritten character recognition. To run a real pattern recognition project from scratch, several classical texts on statistical pattern recognition could and should be consulted, as our work is limited to general probability-theoretical aspects of pattern recognition. We have over 430 exercises to help the scholars. These include skill honing exercises, brainteasers, cute puzzles, open problems, and serious mathematical challenges. There is no solution manual. This book is only a start. Use it as a toy-read some proofs, enjoy some inequalities, learn new tricks, and study the art of camouflaging one problem to look like another. Learn for the sake of learning.

1 Introduction
2 The Bayes error
3 Inequalities and alternate distance measures
4 Linear discrimination
5 Nearest neighbor rules
6 Consistency
7 Slow rates of convergence
8 Error estimation
9 The regular histogram rule
10 Kernel rules
11 Consistency of the k-nearest neighbor rule
12 Vapnik-Chervonenkis theory
13 Combinatorial aspects of Vapnik-Chervonenkis theory
14 Lower bounds for empirical classifier selection
15 The maximum likelihood principle
16 Parametric classification
17 Generalized linear discrimination
18 Complexity regularization
19 Condensed and edited nearest neighbor rules
20 Tree classifiers
21 Data-dependent partitioning
22 Splitting the data
23 The resubstitution estimate
24 Deleted estimates of the error probability
25 Automatic kernel rules
26 Automatic nearest neighbor rules
27 Hypercubes and discrete spaces
28 Epsilon entropy and totally bounded sets
29 Uniform laws of large numbers
30 Neural networks
31 Other error estimates
32 Feature extraction


Figure 1.1.

## 2

## The Bayes Error

### 2.1 The Bayes Problem

In this section, we define the mathematical model and introduce the notation we will use for the entire book. Let $(X, Y)$ be a pair of random variables taking their respective values from $\mathcal{R}^{d}$ and $\{0,1\}$. The random pair $(X, Y)$ may be described in a variety of ways: for example, it is defined by the pair $(\mu, \eta)$, where $\mu$ is the probability measure for $X$ and $\eta$ is the regression of $Y$ on $X$. More precisely, for a Borel-measurable set $A \subseteq \mathcal{R}^{d}$,

$$
\mu(A)=\mathbf{P}\{X \in A\}
$$

and for any $x \in \mathcal{R}^{d}$,

$$
\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}=\mathbf{E}\{Y \mid X=x\}
$$

Thus, $\eta(x)$ is the conditional probability that $Y$ is 1 given $X=x$. To see that this suffices to describe the distribution of $(X, Y)$, observe that for any $C \subseteq \mathcal{R}^{d} \times\{0,1\}$, we have

$$
C=\left(C \cap\left(\mathcal{R}^{d} \times\{0\}\right)\right) \cup\left(C \cap\left(\mathcal{R}^{d} \times\{1\}\right)\right) \stackrel{\text { def }}{=} C_{0} \times\{0\} \cup C_{1} \times\{1\}
$$

and

$$
\begin{aligned}
\mathbf{P}\{(X, Y) \in C\} & =\mathbf{P}\left\{X \in C_{0}, Y=0\right\}+\mathbf{P}\left\{X \in C_{1}, Y=1\right\} \\
& =\int_{C_{0}}(1-\eta(x)) \mu(d x)+\int_{C_{1}} \eta(x) \mu(d x)
\end{aligned}
$$

As this is valid for any Borel-measurable set $C$, the distribution of $(X, Y)$ is determined by $(\mu, \eta)$. The function $\eta$ is sometimes called the a posteriori probability.

Any function $g: \mathcal{R}^{d} \rightarrow\{0,1\}$ defines a classifier or a decision function. The error probability of $g$ is $L(g)=\mathbf{P}\{g(X) \neq Y\}$. Of particular interest is the Bayes decision function

$$
g^{*}(x)= \begin{cases}1 & \text { if } \eta(x)>1 / 2 \\ 0 & \text { otherwise }\end{cases}
$$

This decision function minimizes the error probability.
Theorem 2.1. For any decision function $g: \mathcal{R}^{d} \rightarrow\{0,1\}$,

$$
\mathbf{P}\left\{g^{*}(X) \neq Y\right\} \leq \mathbf{P}\{g(X) \neq Y\}
$$

that is, $g^{*}$ is the optimal decision.

Proof. Given $X=x$, the conditional error probability of any decision $g$ may be expressed as

$$
\begin{aligned}
& \mathbf{P}\{g(X) \neq Y \mid X=x\} \\
& \quad=1-\mathbf{P}\{Y=g(X) \mid X=x\} \\
& \quad=1-(\mathbf{P}\{Y=1, g(X)=1 \mid X=x\}+\mathbf{P}\{Y=0, g(X)=0 \mid X=x\}) \\
& \quad=1-\left(I_{\{g(x)=1\}} \mathbf{P}\{Y=1 \mid X=x\}+I_{\{g(x)=0\}} \mathbf{P}\{Y=0 \mid X=x\}\right) \\
& \quad=1-\left(I_{\{g(x)=1\}} \eta(x)+I_{\{g(x)=0\}}(1-\eta(x))\right),
\end{aligned}
$$

where $I_{A}$ denotes the indicator of the set $A$. Thus, for every $x \in \mathcal{R}^{d}$,

$$
\begin{aligned}
& \mathbf{P}\{g(X) \neq Y \mid X=x\}-\mathbf{P}\left\{g^{*}(X) \neq Y \mid X=x\right\} \\
& \quad=\eta(x)\left(I_{\left\{g^{*}(x)=1\right\}}-I_{\{g(x)=1\}}\right)+(1-\eta(x))\left(I_{\left\{g^{*}(x)=0\right\}}-I_{\{g(x)=0\}}\right) \\
& \quad=(2 \eta(x)-1)\left(I_{\left\{g^{*}(x)=1\right\}}-I_{\{g(x)=1\}}\right) \\
& \geq 0
\end{aligned}
$$

by the definition of $g^{*}$. The statement now follows by integrating both sides with respect to $\mu(d x)$.


Figure 2.1. The Bayes decision in the example on the left is 1 if $x>a$, and 0 otherwise.

Remark. $g^{*}$ is called the Bayes decision and $L^{*}=\mathbf{P}\left\{g^{*}(X) \neq Y\right\}$ is referred to as the Bayes probability of error, Bayes error, or Bayes risk. The proof given above reveals that

$$
L(g)=1-\mathbf{E}\left\{I_{\{g(X)=1\}} \eta(X)+I_{\{g(X)=0\}}(1-\eta(X))\right\}
$$

and in particular,

$$
L^{*}=1-\mathbf{E}\left\{I_{\{\eta(X)>1 / 2\}} \eta(X)+I_{\{\eta(X) \leq 1 / 2\}}(1-\eta(X))\right\}
$$

We observe that the a posteriori probability

$$
\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}=\mathbf{E}\{Y \mid X=x\}
$$

minimizes the squared error when $Y$ is to be predicted by $f(X)$ for some function $f: \mathcal{R}^{d} \rightarrow \mathcal{R}:$

$$
\mathbf{E}\left\{(\eta(X)-Y)^{2}\right\} \leq \mathbf{E}\left\{(f(X)-Y)^{2}\right\}
$$

To see why the above inequality is true, observe that for each $x \in \mathcal{R}^{d}$,

$$
\begin{aligned}
& \mathbf{E}\left\{(f(X)-Y)^{2} \mid X=x\right\} \\
&= \mathbf{E}\left\{(f(x)-\eta(x)+\eta(x)-Y)^{2} \mid X=x\right\} \\
&=(f(x)-\eta(x))^{2}+2(f(x)-\eta(x)) \mathbf{E}\{\eta(x)-Y \mid X=x\} \\
&+\mathbf{E}\left\{(\eta(X)-Y)^{2} \mid X=x\right\} \\
&=(f(x)-\eta(x))^{2}+\mathbf{E}\left\{(\eta(X)-Y)^{2} \mid X=x\right\} .
\end{aligned}
$$

The conditional median, i.e., the function minimizing the absolute error $\mathbf{E}\{\mid f(X)-$ $Y \mid\}$ is even more closely related to the Bayes rule (see Problem 2.12).

### 2.2 A Simple Example

Let us consider the prediction of a student's performance in a course (pass/fail) when given a number of important factors. First, let $Y=1$ denote a pass and let $Y=0$ stand for failure. The sole observation $X$ is the number of hours of study per week. This, in itself, is not a foolproof predictor of a student's performance, because for that we would need more information about the student's quickness of mind, health, and social habits. The regression function $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$ is probably monotonically increasing in $x$. If it were known to be $\eta(x)=x /(c+x)$, $c>0$, say, our problem would be solved because the Bayes decision is

$$
g^{*}(x)= \begin{cases}1 & \text { if } \eta(x)>1 / 2 \\ 0 & \text { otherwise. }\end{cases}
$$

The corresponding Bayes error is

$$
L^{*}=L\left(g^{*}\right)=\mathbb{E}\{\min (\eta(X), 1-\eta(X))\}=\mathbb{E}\left\{\frac{\min (c, X)}{c+X}\right\}
$$

While we could deduce the Bayes decision from $\eta$ alone, the same cannot be said for the Bayes error $L^{*}$-it requires knowledge of the distribution of $X$. If $X=c$ with probability one (as in an army school, where all students are forced to study $c$ hours per week), then $L^{*}=1 / 2$. If we have a population that is nicely spread out, say, $X$ is uniform on $[0,4 c]$, then the situation improves:

$$
L^{*}=\frac{1}{4 c} \int_{0}^{4 c} \frac{\min (c, x)}{c+x} d x=\frac{1}{4} \log \frac{5 e}{4} \approx 0.305785
$$

Far away from $x=c$, discrimination is really simple. In general, discrimination is much easier than estimation because of this phenomenon.

### 2.3 Another Simple Example

Let us work out a second simple example in which $Y=0$ or $Y=1$ according to whether a student fails or passes a course. $X$ represents one or more observations regarding the student. The components of $X$ in our example will be denoted by $T, B$, and $E$ respectively, where $T$ is the average number of hours the students watches TV, $B$ is the average number of beers downed each day, and $E$ is an intangible quantity measuring extra negative factors such as laziness and learning difficulties. In our cooked-up example, we have

$$
Y= \begin{cases}1 & \text { if } T+B+E<7 \\ 0 & \text { otherwise }\end{cases}
$$

Thus, if $T, B$, and $E$ are known, $Y$ is known as well. The Bayes classifier decides 1 if $T+B+E<7$ and 0 otherwise. The corresponding Bayes probability of error is zero. Unfortunately, $E$ is intangible, and not available to the observer. We only have access to $T$ and $B$. Given $T$ and $B$, when should we guess that $Y=1$ ? To answer this question, one must know the joint distribution of $(T, B, E)$, or, equivalently, the joint distribution of $(T, B, Y)$. So, let us assume that $T, B$, and $E$ are i.i.d. exponential random variables (thus, they have density $e^{-u}$ on $[0, \infty)$ ). The Bayes rule compares $\mathbf{P}\{Y=1 \mid T, B\}$ with $\mathbf{P}\{Y=0 \mid T, B\}$ and makes a decision consistent with the maximum of these two values. A simple calculation shows that

$$
\begin{aligned}
\mathbf{P}\{Y=1 \mid T, B\} & =\mathbf{P}\{T+B+E<7 \mid T, B\} \\
& =\mathbf{P}\{E<7-T-B \mid T, B\} \\
& =\max \left(0,1-e^{-(7-T-B)}\right)
\end{aligned}
$$

The crossover between two decisions occurs when this value equals $1 / 2$. Thus, the Bayes classifier is as follows:

$$
g^{*}(T, B)= \begin{cases}1 & \text { if } T+B<7-\log 2=6.306852819 \ldots \\ 0 & \text { otherwise }\end{cases}
$$

Of course, this classifier is not perfect. The probability of error is

$$
\begin{aligned}
& \mathbf{P}\left\{g^{*}(T, B) \neq Y\right\} \\
&= \mathbf{P}\{T+B<7-\log 2, T+B+E \geq 7\} \\
&+\mathbf{P}\{T+B \geq 7-\log 2, T+B+E<7\} \\
&= \mathbf{E}\left\{e^{-(7-T-B)} I_{\{T+B<7-\log 2\}}\right\} \\
&+\mathbf{P}\left\{\left(1-e^{-(7-T-B)}\right) I_{\{7>T+B \geq 7-\log 2\}}\right\} \\
&= \int_{0}^{7-\log 2} x e^{-x} e^{-(7-x)} d x+\int_{7-\log 2}^{7} x e^{-x}\left(1-e^{-(7-x)}\right) d x
\end{aligned}
$$

$$
\text { (since the density of } T+B \text { is } u e^{-u} \text { on }[0, \infty) \text { ) }
$$

$$
=e^{-7}\left(\frac{(7-\log 2)^{2}}{2}+2(8-\log 2)-8-\frac{7^{2}}{2}+\frac{(7-\log 2)^{2}}{2}\right)
$$

$$
\left(\text { as } \int_{x}^{\infty} u e^{-u} d u=(1+x) e^{-x}\right)
$$

$$
=0.0199611 \ldots
$$

If we have only access to $T$, then the Bayes classifier is allowed to use $T$ only. First, we find

$$
\begin{aligned}
\mathbf{P}\{Y=1 \mid T\} & =\mathbf{P}\{E+B<7-T \mid T\} \\
& =\max \left(0,1-(1+7-T) e^{-(7-T)}\right)
\end{aligned}
$$

The crossover at $1 / 2$ occurs at $T=c \stackrel{\text { def }}{=} 5.321653009 \ldots$, so that the Bayes classifier is given by

$$
g^{*}(T)= \begin{cases}1 & \text { if } T<c \\ 0 & \text { otherwise } .\end{cases}
$$

The probability of error is

$$
\begin{aligned}
& \mathbb{P}\left\{g^{*}(T) \neq Y\right\} \\
&= \mathbf{P}\{T<c, T+B+E \geq 7\}+\mathbf{P}\{T \geq c, T+B+E<7\} \\
&= \mathbf{E}\left\{(1+7-T) e^{-(7-T)} I_{\{T<c\}}\right\} \\
&+\mathbf{P}\left\{\left(1-(1+7-T) e^{-(7-T)}\right) I_{\{7>T \geq c\}}\right\} \\
&= \int_{0}^{c} e^{-x}(1+7-x) e^{-(7-x)} d x+\int_{c}^{7} e^{-x}\left(1-(1+7-x) e^{-(7-x)}\right) d x \\
&= e^{-7}\left(\frac{8^{2}}{2}-\frac{(8-c)^{2}}{2}+e^{-(c-7)}-1-\frac{(8-c)^{2}}{2}+\frac{1}{2}\right) \\
&= 0.02235309002 \ldots .
\end{aligned}
$$

The Bayes error has increased slightly, but not by much. Finally, if we do not have access to any of the three variables, $T, B$, and $E$, the best we can do is see which
class is most likely. To this end, we compute

$$
\mathbf{P}\{Y=0\}=\mathbf{P}\{T+B+E \geq 7\}=\left(1+7+7^{2} / 2\right) e^{-7}=.02963616388 \ldots .
$$

If we set $g \equiv 1$ all the time, we make an error with probability $0.02963616388 \ldots$.
In practice, Bayes classifiers are unknown simply because the distribution of ( $X, Y$ ) is unknown. Consider a classifier based upon ( $T, B$ ). Rosenblatt's perceptron (see Chapter 4) looks for the best linear classifier based upon the data. That is, the decision is of the form

$$
g(T, B)= \begin{cases}1 & \text { if } a T+b B<c \\ 0 & \text { otherwise }\end{cases}
$$

for some data-based choices for $a, b$ and $c$. If we have lots of data at our disposal, then it is possible to pick out a linear classifier that is nearly optimal. As we have seen above, the Bayes classifier happens to be linear. That is a sheer coincidence, of course. If the Bayes classifier had not been linear-for example, if we had $Y=I_{\left\{T+B^{2}+E<7\right\}}$-then even the best perceptron would be suboptimal, regardless of how many data pairs one would have. If we use the 3-nearest neighbor rule (Chapter 5), the asymptotic probability of error is not more than 1.3155 times the Bayes error, which in our example is about 0.02625882705 . The example above also shows the need to look at individual components, and to evaluate how many and which components would be most useful for discrimination. This subject is covered in the chapter on feature extraction (Chapter 32).

### 2.4 Other Formulas for the Bayes Risk

The following forms of the Bayes error are often convenient:

$$
\begin{aligned}
L^{*} & =\inf _{g: \mathcal{R}^{d} \rightarrow\{0,1\}} \mathbf{P}\{g(X) \neq Y\} \\
& =\mathbf{E}\{\min \{\eta(X), 1-\eta(X)\}\} \\
& =\frac{1}{2}-\frac{1}{2} \mathbf{E}\{|2 \eta(X)-1|\} .
\end{aligned}
$$



Figure 2.2. The Bayes decision when classconditional densities exist. In the figure on the left, the decision is 0 on $[a, b]$ and 1 elsewhere.

In special cases, we may obtain other helpful forms. For example, if $X$ has a density $f$, then

$$
\begin{aligned}
L^{*} & =\int \min (\eta(x), 1-\eta(x)) f(x) d x \\
& =\int \min \left((1-p) f_{0}(x), p f_{1}(x)\right) d x
\end{aligned}
$$

where $p=\mathbf{P}\{Y=1\}$, and $f_{i}(x)$ is the density of $X$ given that $Y=i . p$ and $1-p$ are called the class probabilities, and $f_{0}, f_{1}$ are the class-conditional densities. If $f_{0}$ and $f_{1}$ are nonoverlapping, that is, $\int f_{0} f_{1}=0$, then obviously $L^{*}=0$. Assume moreover that $p=1 / 2$. Then

$$
\begin{aligned}
L^{*} & =\frac{1}{2} \int \min \left(f_{0}(x), f_{1}(x)\right) d x \\
& =\frac{1}{2} \int f_{1}(x)-\left(f_{1}(x)-f_{0}(x)\right)_{+} d x \\
& =\frac{1}{2}-\frac{1}{4} \int\left|f_{1}(x)-f_{0}(x)\right| d x
\end{aligned}
$$

Here $g_{+}$denotes the positive part of a function $g$. Thus, the Bayes error is directly related to the $L_{1}$ distance between the class densities.


Figure 2.3. The shaded area is the $L_{1}$ distance between the class-conditional densities.

### 2.5 Plug-In Decisions

The best guess of $Y$ from the observation $X$ is the Bayes decision

$$
g^{*}(x)=\left\{\begin{array}{ll}
0 & \text { if } \eta(x) \leq 1 / 2 \\
1 & \text { otherwise }
\end{array}= \begin{cases}0 & \text { if } \eta(x) \leq 1-\eta(x) \\
1 & \text { otherwise } .\end{cases}\right.
$$

The function $\eta$ is typically unknown. Assume that we have access to nonnegative functions $\tilde{\eta}(x), 1-\tilde{\eta}(x)$ that approximate $\eta(x)$ and $1-\eta(x)$ respectively. In this case it seems natural to use the plug-in decision function

$$
g(x)= \begin{cases}0 & \text { if } \tilde{\eta}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

to approximate the Bayes decision. The next well-known theorem (see, e.g., Van Ryzin (1966), Wolverton and Wagner (1969a), Glick (1973), Csibi (1971), Györfi (1975), (1978), Devroye and Wagner (1976b), Devroye (1982b), and Devroye and Györfi (1985)) states that if $\tilde{\eta}(x)$ is close to the real a posteriori probability in $L_{1}$-sense, then the error probability of decision $g$ is near the optimal decision $g^{*}$.

Theorem 2.2. For the error probability of the plug-in decision $g$ defined above, we have

$$
\mathbf{P}\{g(X) \neq Y\}-L^{*}=2 \int_{\mathcal{R}^{d}}|\eta(x)-1 / 2| I_{\left\{g(x) \neq g^{*}(x)\right\}} \mu(d x),
$$

and

$$
\mathbf{P}\{g(X) \neq Y\}-L^{*} \leq 2 \int_{\mathcal{R}^{d}}|\eta(x)-\tilde{\eta}(x)| \mu(d x)=2 \mathbf{E}|\eta(X)-\tilde{\eta}(X)| .
$$

Proof. If for some $x \in \mathcal{R}^{d}, g(x)=g^{*}(x)$, then clearly the difference between the conditional error probabilities of $g$ and $g^{*}$ is zero:

$$
\mathbf{P}\{g(X) \neq Y \mid X=x\}-\mathbf{P}\left\{g^{*}(X) \neq Y \mid X=x\right\}=0 .
$$

Otherwise, if $g(x) \neq g^{*}(x)$, then as seen in the proof of Theorem 2.1, the difference may be written as

$$
\begin{aligned}
\mathbf{P}\{g(X) \neq Y \mid X=x\} & -\mathbf{P}\left\{g^{*}(X) \neq Y \mid X=x\right\} \\
& =(2 \eta(x)-1)\left(I_{\left\{g^{*}(x)=1\right\}}-I_{\{g(x)=1\}}\right) \\
& =|2 \eta(x)-1| I_{\left\{g(x) \neq g^{*}(x)\right\}} .
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\mathbf{P}\{g(X) \neq Y\}-L^{*} & =\int_{\mathcal{R}^{d}} 2|\eta(x)-1 / 2| L_{\left\{g(x) \neq \boldsymbol{g}^{*}(x)\right\}} \mu(d x) \\
& \leq \int_{\mathcal{R}^{d}} 2|\eta(x)-\tilde{\eta}(x)| \mu(d x),
\end{aligned}
$$

since $g(x) \neq g^{*}(x)$ implies $|\eta(x)-\tilde{\eta}(x)| \geq|\eta(x)-1 / 2|$.
When the classifier $g(x)$ can be put in the form

$$
g(x)= \begin{cases}0 & \text { if } \tilde{\eta}_{1}(x) \leq \tilde{\eta}_{0}(x) \\ 1 & \text { otherwise },\end{cases}
$$

where $\tilde{\eta}_{1}(x), \tilde{\eta}_{0}(x)$ are some approximations of $\eta(x)$ and $1-\eta(x)$, respectively, the situation differs from that discussed in Theorem 2.2 if $\tilde{\eta}_{0}(x)+\tilde{\eta}_{1}(x)$ does not necessarily equal to one. However, an inequality analogous to that of Theorem 2.2 remains true:

Theorem 2.3. The error probability of the decision defined above is bounded from above by

$$
\mathbf{P}\{g(X) \neq Y\}-L^{*} \leq \int_{\mathcal{R}^{d}}\left|(1-\eta(x))-\tilde{\eta}_{0}(x)\right| \mu(d x)+\int_{\mathcal{R}^{d}}\left|\eta(x)-\tilde{\eta}_{1}(x)\right| \mu(d x)
$$

The proof is left to the reader (Problem 2.9).
Remark. Assume that the class-conditional densities $f_{0}, f_{1}$ exist and are approximated by the densities $\tilde{f}_{0}(x), \tilde{f}_{1}(x)$. Assume furthermore that the class probabilities $p=\mathbf{P}\{Y=1\}$ and $1-p=\mathbf{P}\{Y=0\}$ are approximated by $\tilde{p}_{1}$ and $\tilde{p}_{0}$, respectively. Then for the error probability of the plug-in decision function

$$
g(x)= \begin{cases}0 & \text { if } \tilde{p}_{1} \tilde{f}_{1}(x) \leq \tilde{p}_{0} \tilde{f}_{0}(x) \\ 1 & \text { otherwise }\end{cases}
$$

$$
\begin{aligned}
& \mathbf{P}\{g(X) \neq Y\}-L^{*} \\
& \quad \leq \int_{\mathcal{R}^{d}}\left|(1-p) f_{0}(x)-\tilde{p}_{0} \tilde{f}_{0}(x)\right| d x+\int_{\mathcal{R}^{d}}\left|p f_{1}(x)-\tilde{p}_{1} \tilde{f}_{1}(x)\right| d x
\end{aligned}
$$

See Problem 2.10.

### 2.6 Bayes Error Versus Dimension

The components of $X$ that matter in the Bayes classifier are those that explicitly appear in $\eta(X)$. In fact, then, all discrimination problems are one-dimensional, as we could equally well replace $X$ by $\eta(X)$ or by any strictly monotone function of $\eta(X)$, such as $\eta^{7}(X)+5 \eta^{3}(X)+\eta(X)$. Unfortunately, $\eta$ is unknown in general. In the example in Section 2.3, we had in one case

$$
\eta(T, B)=\max \left(0,1-e^{-(7-T-B)}\right)
$$

and in another case

$$
\eta(T)=\max \left(0,1-(1+7-T) e^{-(7-T)}\right)
$$

The former format suggests that we could base all decisions on $T+B$. This means that if we had no access to $T$ and $B$ individually, but to $T+B$ jointly, we would be able to achieve the same results! Since $\eta$ is unknown, all of this is really irrelevant.

In general, the Bayes risk increases if we replace $X$ by $T(X)$ for any transformation $T$ (see Problem 2.1), as this destroys information. On the other hand, there exist transformations (such as $\eta(X)$ ) that leave the Bayes error untouched. For more on the relationship between the Bayes error and the dimension, refer to Chapter 32.

## Problems and Exercises

Problem 2.1. Let $T: \mathcal{X} \rightarrow \mathcal{X}^{\prime}$ be an arbitrary measurable function. If $L_{X}^{*}$ and $L_{T(X)}^{*}$ denote the Bayes error probabilities for $(X, Y)$ and $(T(X), Y)$, respectively, then prove that

$$
L_{T(X)}^{*} \geq L_{X}^{*}
$$

(This shows that transformations of $X$ destroy information, because the Bayes risk increases.)

Problem 2.2. Let $X^{\prime}$ be independent of $(X, Y)$. Prove that

$$
L_{\left(X, X^{\prime}\right)}^{*}=L_{X}^{*} .
$$

Problem 2.3. Show that $L^{*} \leq \min (p, 1-p)$, where $p, 1-p$ are the class probabilities. Show that equality holds if $X$ and $Y$ are independent. Exhibit a distribution where $X$ is not independent of $Y$, but $L^{*}=\min (p, 1-p)$.

Problem 2.4. Neyman-Pearson lemma. Consider again the decision problem, but with a decision $g$, we now assign two error probabilities,

$$
L^{(0)}(g)=\mathbf{P}\{g(X)=1 \mid Y=0\} \text { and } L^{(1)}(g)=\mathbf{P}\{g(X)=0 \mid Y=1\} .
$$

Assume that the class-conditional densities $f_{0}, f_{1}$ exist. For $c>0$, define the decision

$$
g_{c}(x)= \begin{cases}1 & \text { if } c f_{1}(x)>f_{0}(x) \\ 0 & \text { otherwise } .\end{cases}
$$

Prove that for any decision $g$, if $L^{(0)}(g) \leq L^{(0)}\left(g_{c}\right)$, then $L^{(1)}(g) \geq L^{(1)}\left(g_{c}\right)$. In other words, if $L^{(0)}$ is required to be kept under a certain level, then the decision minimizing $L^{(1)}$ has the form of $g_{c}$ for some $c$. Note that $g^{*}$ is like that.

Problem 2.5. Decisions with rejection. Sometimes in decision problems, one is allowed to say "I don't know," if this does not happen frequently. These decisions are called decisions with a reject option (see, e.g., Forney (1968), Chow (1970)). Formally, a decision $g(x)$ can have three values: 0,1 , and "reject." There are two performance measures: the probability of rejection $\mathbf{P}\{g(X)=$ "reject" $\}$, and the error probability $\mathbf{P}\{g(X) \neq Y \mid g(X) \neq$ "reject" $\}$. For a $0<c<1 / 2$, define the decision

$$
g_{c}(x)= \begin{cases}1 & \text { if } \eta(x)>1 / 2+c \\ 0 & \text { if } \eta(x) \leq 1 / 2-c \\ \text { "reject" } & \text { otherwise }\end{cases}
$$

Show that for any decision $g$, if

$$
\mathbf{P}\{g(X)=\text { "reject" }\} \leq \mathbf{P}\left\{g_{c}(X)=\text { "reject" }\right\},
$$

then

$$
\mathbf{P}\{g(X) \neq Y \mid g(X) \neq \text { "reject" }\} \geq \mathbf{P}\left\{g_{c}(X) \neq Y \mid g_{c}(X) \neq \text { "reject" }\right\} .
$$

Thus, to keep the probability of rejection under a certain level, decisions of the form of $g_{c}$ are optimal (Györfi, Györfi, and Vajda (1978)).

Problem 2.6. Consider the prediction of a student's failure based upon variables $T$ and $B$, where $Y=I_{\{T+B+E<7\}}$ and $E$ is an inaccessible variable (see Section 2.3).
(1) Let $T, B$, and $E$ be independent. Merely by changing the distribution of $E$, show that the Bayes error for classification based upon $(T, B)$ can be made as close as desired to $1 / 2$.
(2) Let $T$ and $B$ be independent and exponentially distributed. Find a joint distribution of ( $T, B, E$ ) such that the Bayes classifier is not a linear classifier.
(3) Let $T$ and $B$ be independent and exponentially distributed. Find a joint distribution of $(T, B, E)$ such that the Bayes classifier is given by

$$
g^{*}(T, B)= \begin{cases}1 & \text { if } T^{2}+B^{2}<10 \\ 0 & \text { otherwise }\end{cases}
$$

(4) Find the Bayes classifier and Bayes error for classification based on $(T, B)$ (with $Y$ as above) if ( $T, B, E$ ) is uniformly distributed on $[0,4]^{3}$.

Problem 2.7. Assume that $T, B$, and $E$ are independent uniform [ 0,4 ] random variables with interpretations as in Section 2.3. Let $Y=1(0)$ denote whether a student passes (fails) a course. Assume that $Y=1$ if and only if $T B E \leq 8$.
(1) Find the Bayes decision if no variable is available, if only $T$ is available, and if only $T$ and $B$ are available.
(2) Determine in all three cases the Bayes error.
(3) Determine the best linear classifier based upon $T$ and $B$ only.

Problem 2.8. Let $\eta^{\prime}, \eta^{\prime \prime}: \mathcal{R}^{d} \rightarrow[0,1]$ be arbitrary measurable functions, and define the corresponding decisions by $g^{\prime}(x)=I_{\left\{\eta^{\prime}(x)>1 / 2\right\}}$ and $g^{\prime \prime}(x)=I_{\left\{\eta^{\prime \prime}(x)>1 / 2\right\}}$. Prove that

$$
\left|L\left(g^{\prime}\right)-L\left(g^{\prime \prime}\right)\right| \leq \mathbf{P}\left\{g^{\prime}(X) \neq g^{\prime \prime}(X)\right\}
$$

and

$$
\left|L\left(g^{\prime}\right)-L\left(g^{\prime \prime}\right)\right| \leq \mathbf{E}\left\{|2 \eta(X)-1| I_{\left\{g^{\prime}(X) \neq g^{\prime \prime}(X)\right\}}\right\} .
$$

Problem 2.9. Prove Theorem 2.3.
Problem 2.10. Assume that the class-conditional densities $f_{0}$ and $f_{1}$ exist and are approximated by the densities $\tilde{f}_{0}$ and $\tilde{f}_{1}$, respectively. Assume furthermore that the class probabilities $p=\mathbf{P}\{Y=1\}$ and $1-p=\mathbf{P}\{Y=0\}$ are approximated by $\tilde{p}_{1}$ and $\tilde{p}_{0}$. Prove that for the error probability of the plug-in decision function

$$
g(x)= \begin{cases}0 & \text { if } \tilde{p}_{1} \tilde{f}_{1}(x) \leq \tilde{p}_{0} \tilde{f}_{0}(x) \\ 1 & \text { otherwise }\end{cases}
$$

we have

$$
\mathbf{P}\{g(X) \neq Y\}-L^{*} \leq \int_{\mathcal{R}^{d}}\left|p f_{1}(x)-\tilde{p}_{1} \tilde{f}_{1}(x)\right| d x+\int_{\mathcal{R}^{d}}\left|(1-p) f_{0}(x)-\tilde{p}_{0} \tilde{f}_{0}(x)\right| d x
$$

Problem 2.11. Using the notation of Problem 2.10, show that if for a sequence of $\tilde{f}_{m, n}(x)$ and $\tilde{p}_{m, n}(m=0,1)$,

$$
\lim _{n \rightarrow \infty} \int_{\mathcal{R}^{d}}\left|p f_{1}(x)-\tilde{p}_{1, n} \tilde{f}_{1, n}(x)\right|^{2} d x+\int_{\mathcal{R}^{d}}\left|(1-p) f_{0}(x)-\tilde{p}_{0, n} \tilde{f}_{0, n}(x)\right|^{2} d x=0
$$

then for the corresponding sequence of plug-in decisions $\lim _{n \rightarrow \infty} \mathbf{P}\left\{g_{n}(X) \neq Y\right\}=L^{*}$ (Wolverton and Wagner (1969a)). Hint: According to Problem 2.10, it suffices to show that if we are given a deterministic sequence of density functions $f, f_{1}, f_{2}, f_{3}, \ldots$, then

$$
\lim _{n \rightarrow \infty} \int\left(f_{n}(x)-f(x)\right)^{2} d x=0
$$

implies

$$
\lim _{n \rightarrow \infty} \int\left|f_{n}(x)-f(x)\right| d x=0
$$

(A function $f$ is called a density function if it is nonnegative and $\int f(x) d x=1$.) To see this, observe that

$$
\int\left|f_{n}(x)-f(x)\right| d x=2 \int\left(f_{n}(x)-f(x)\right)_{+} d x=2 \sum_{i} \int_{A_{i}}\left(f_{n}(x)-f(x)\right)_{+} d x
$$

where $A_{1}, A_{2}, \ldots$ is a partition of $\mathcal{R}^{d}$ into unit cubes, and $f_{+}$denotes the positive part of a function $f$. The key observation is that convergence to zero of each term of the infinite sum implies convergence of the whole integral by the dominated convergence theorem, since $\int\left(f_{n}(x)-f(x)\right)_{+} d x \leq \int f_{n}(x) d x=1$. Handle the right-hand side by the Cauchy-Schwarz inequality.

Problem 2.12. Define the $L_{1}$ error of a function $f: \mathcal{R}^{d} \rightarrow \mathcal{R}$ by $J(f)=\mathbf{E}\{|f(X)-Y|\}$. Show that a function minimizing $J(f)$ is the Bayes rule $g^{*}$, that is, $J^{*}=\inf _{f} J(f)=J\left(g^{*}\right)$. Thus, $J^{*}=L^{*}$. Define a decision by

$$
g(x)= \begin{cases}0 & \text { if } f(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

Prove that its error probability $L(g)=\mathbf{P}\{g(X) \neq Y\}$ satisfies the inequality

$$
L(g)-L^{*} \leq J(f)-J^{*}
$$

## 3

## Inequalities and Alternate Distance Measures

### 3.1 Measuring Discriminatory Information

In our two-class discrimination problem, the best rule has (Bayes) probability of error

$$
L^{*}=\mathbf{E}\{\min (\eta(X), 1-\eta(X))\} .
$$

This quantity measures how difficult the discrimination problem is. It also serves as a gauge of the quality of the distribution of $(X, Y)$ for pattern recognition. Put differently, if $\psi_{1}$ and $\psi_{2}$ are certain many-to-one mappings, $L^{*}$ may be used to compare discrimination based on $\left(\psi_{1}(X), Y\right)$ with that based on $\left(\psi_{2}(X), Y\right)$. When $\psi_{1}$ projects $\mathcal{R}^{d}$ to $\mathcal{R}^{d_{1}}$ by taking the first $d_{1}$ coordinates, and $\psi_{2}$ takes the last $d_{2}$ coordinates, the corresponding Bayes errors will help us decide which projection is better. In this sense, $L^{*}$ is the fundamental quantity in feature extraction.

Other quantities have been suggested over the years that measure the discriminatory power hidden in the distribution of ( $X, Y$ ). These may be helpful in some settings. For example, in theoretical studies or in certain proofs, the relationship between $L^{*}$ and the distribution of ( $X, Y$ ) may become clearer via certain inequalities that link $L^{*}$ with other functionals of the distribution. We all understand moments and variances, but how do these simple functionals relate to $L^{*}$ ? Perhaps we may even learn a thing or two about what it is that makes $L^{*}$ small. In feature selection, some explicit inequalities involving $L^{*}$ may provide just the kind of numerical information that will allow one to make certain judgements on what kind of feature is preferable in practice. In short, we will obtain more information about $L^{*}$ with a variety of uses in pattern recognition.

In the next few sections, we avoid putting any conditions on the distribution of ( $X, Y$ ).

### 3.2 The Kolmogorov Variational Distance

Inspired by the total variation distance between distributions, the Kolmogorov variational distance

$$
\begin{aligned}
\delta_{\mathrm{Ko}} & =\frac{1}{2} \mathbf{E}\{|\mathbf{P}\{Y=1 \mid X\}-\mathbf{P}\{Y=0 \mid X\}|\} \\
& =\frac{1}{2} \mathbf{E}\{|2 \eta(X)-1|\}
\end{aligned}
$$

captures the distance between the two classes. We will not need anything special to deal with $\delta_{\mathrm{KO}}$ as

$$
\begin{aligned}
L^{*} & =\mathbf{E}\left\{\frac{1}{2}-\frac{1}{2}|2 \eta(X)-1|\right\} \\
& =\frac{1}{2}-\frac{1}{2} \mathbf{E}\{|2 \eta(X)-1|\} \\
& =\frac{1}{2}-\delta_{\text {ко }} .
\end{aligned}
$$

### 3.3 The Nearest Neighbor Error

The asymptotic error of the nearest neighbor rule is

$$
L_{\mathrm{NN}}=\mathbf{E}\{2 \eta(X)(1-\eta(X))\}
$$

(see Chapter 5). Clearly, $2 \eta(1-\eta) \geq \min (\eta, 1-\eta)$ as $2 \max (\eta, 1-\eta) \geq 1$. Also, using the notation $A=\min (\eta(X), 1-\eta(X))$, we have

$$
\begin{align*}
L^{*} \leq L_{\mathrm{NN}}= & 2 \mathbf{E}\{A(1-A)\} \\
\leq & 2 \mathbf{E}\{A\} \cdot \mathbf{E}\{1-A\} \\
& (\text { by the second association inequality of Theorem A. 19) } \\
= & 2 L^{*}\left(1-L^{*}\right) \leq 2 L^{*}, \tag{3.1}
\end{align*}
$$

which are well-known inequalities of Cover and Hart (1967). $L_{\mathrm{NN}}$ provides us with quite a bit of information about $L^{*}$.

The measure $L_{\mathrm{NN}}$ has been rediscovered under other guises: Devijver and Kittler (1982, p.263) and Vajda (1968) call it the quadratic entropy, and Mathai and Rathie (1975) refer to it as the harmonic mean coefficient.


Figure 3.1. Relationship between the Bayes error and the asymptotic nearest neighbor error. Every point in the unshaded region is possible.

### 3.4 The Bhattacharyya Affinity

The Bhattacharyya measure of affinity (Bhattacharyya, (1946)) is $-\log \rho$, where

$$
\rho=\mathbf{E}\{\sqrt{\eta(X)(1-\eta(X))}\}
$$

will be referred to as the Matushita error. It does not occur naturally as the limit of any standard discrimination rule (see, however, Problem 6.11). $\rho$ was suggested as a distance measure for pattern recognition by Matushita (1956). It also occurs under other guises in mathematical statistics-see, for example, the Hellinger distance literature (Le Cam (1970), Beran (1977)).
Clearly, $\rho=0$ if and only if $\eta(X) \in\{0,1\}$ with probability one, that is, if $L^{*}=0$. Furthermore, $\rho$ takes its maximal value $1 / 2$ if and only if $\eta(X)=1 / 2$ with probability one. The relationship between $\rho$ and $L^{*}$ is not linear though. We will show that for all distributions, $L_{\mathrm{NN}}$ is more useful than $\rho$ if it is to be used as an approximation of $L^{*}$.

Theorem 3.1. For all distributions, we have

$$
\begin{aligned}
\frac{1}{2}-\frac{1}{2} \sqrt{1-4 \rho^{2}} & \leq \frac{1}{2}-\frac{1}{2} \sqrt{1-2 L_{\mathrm{NN}}} \\
& \leq L^{*} \\
& \leq L_{\mathrm{NN}} \\
& \leq \rho .
\end{aligned}
$$

Proof. First of all,

$$
\begin{aligned}
\rho^{2} & =\mathbf{E}^{2}\{\sqrt{\eta(X)(1-\eta(X))}\} \\
& \leq \mathbf{E}\{\eta(X)(1-\eta(X))\} \text { (by Jensen's inequality) }
\end{aligned}
$$

$$
\begin{aligned}
& =\frac{L_{\mathrm{NN}}}{2} \\
& \leq L^{*}\left(1-L^{*}\right) \text { (by the Cover-Hart inequality (3.1)). }
\end{aligned}
$$

Second, as $\sqrt{\eta(1-\eta)} \geq 2 \eta(1-\eta)$ for all $\eta \in[0,1]$, we see that $\rho \geq L_{\mathrm{NN}} \geq L^{*}$. Finally, by the Cover-Hart inequality,

$$
\sqrt{1-2 L_{\mathrm{NN}}} \geq \sqrt{1-4 L^{*}\left(1-L^{*}\right)}=1-2 L^{*} .
$$

Putting all these things together establishes the chain of inequalities.


Figure 3.2. The inequalities linking $\rho$ to $L^{*}$ are illustrated. Note that the region is larger than that cut out in the ( $L^{*}, L_{\mathrm{NN}}$ ) plane in Figure 3.1.

The inequality $L_{\mathrm{NN}} \leq \rho$ is due to Ito (1972). The inequality $L_{\mathrm{NN}} \geq 2 \rho^{2}$ is due to Horibe (1970). The inequality $1 / 2-\sqrt{1-4 \rho^{2}} / 2 \leq L^{*} \leq \bar{\rho}$ can be found in Kailath (1967). The left-hand side of the last inequality was shown by Hudimoto (1957). All these inequalities are tight (see Problem 3.2). The appeal of quantities like $L_{\mathrm{NN}}$ and $\rho$ is that they involve polynomials of $\eta$, whereas $L^{*}=$ $\mathbf{E}\{\min (\eta(X), 1-\eta(X))\}$ is nonpolynomial. For certain discrimination problems in which $X$ has a distribution that is known up to certain parameters, one may be able to compute $L_{\mathrm{NN}}$ and $\rho$ explicitly as a function of these parameters. Via inequalities, this may then be used to obtain performance guarantees for parametric discrimination rules of the plug-in type (see Chapter 16).

For completeness, we mention a generalization of Bhattacharyya's measure of affinity, first suggested by Chernoff (1952):

$$
\delta_{C}=-\log \left(\mathbb{E}\left\{\eta^{\alpha}(X)(1-\eta(X))^{1-\alpha}\right\}\right)
$$

where $\alpha \in(0,1)$ is fixed. For $\alpha=1 / 2, \delta_{C}=-\log \rho$. The asymmetry introduced by taking $\alpha \neq 1 / 2$ has no practical interpretation, however.

### 3.5 Entropy

The entropy of a discrete probability distribution ( $p_{1}, p_{2}, \ldots$ ) is defined by

$$
\mathcal{H}=\mathcal{H}\left(p_{1}, p_{2}, \ldots\right)=-\sum_{i=1}^{\infty} p_{i} \log p_{i}
$$

where, by definition, $0 \log 0=0$ (Shannon (1948)). The key quantity in information theory (see Cover and Thomas (1991)), it has countless applications in many branches of computer science, mathematical statistics, and physics. The entropy's main properties may be summarized as follows.
A. $\mathcal{H} \geq 0$ with equality if and only if $p_{i}=1$ for some $i$. Proof: $\log p_{i} \leq 0$ for all $i$ with equality if and only if $p_{i}=1$ for some $i$. Thus, entropy is minimal for a degenerate distribution, i.e., a distribution with the least amount of "spread."
B. $\mathcal{H}\left(p_{1}, \ldots, p_{k}\right) \leq \log k$ with equality if and only if $p_{1}=p_{2}=\cdots=p_{k}=$ $1 / k$. In other words, the entropy is maximal when the distribution is maximally smeared out. Proof:

$$
\mathcal{H}\left(p_{1}, \ldots, p_{k}\right)-\log k=\sum_{i=1}^{k} p_{i} \log \left(\frac{1}{k p_{i}}\right) \leq 0
$$

by the inequality $\log x \leq x-1, x>0$.
C. For a Bernoulli distribution $(p, 1-p)$, the binary entropy $\mathcal{H}(p, 1-p)=$ $-p \log p-(1-p) \log (1-p)$ is concave in $p$.

Assume that $X$ is a discrete random variable that must be guessed by asking questions of the type "is $X \in A$ ?," for some sets $A$. Let $N$ be the minimum expected number of questions required to determine $X$ with certainty. It is well known that

$$
\frac{\mathcal{H}}{\log 2} \leq N<\frac{\mathcal{H}}{\log 2}+1
$$

(e.g., Cover and Thomas (1991)). Thus, $\mathcal{H}$ not only measures how spread out the mass of $X$ is, but also provides us with concrete computational bounds for certain algorithms. In the simple example above, $\mathcal{H}$ is in fact proportional to the expected computational time of the best algorithm.

We are not interested in information theory per se, but rather in its usefulness in pattern recognition. For our discussion, if we fix $X=x$, then $Y$ is Bernoulli $(\eta(x))$. Hence, the conditional entropy of $Y$ given $X=x$ is

$$
\mathcal{H}(\eta(x), 1-\eta(x))=-\eta(x) \log \eta(x)-(1-\eta(x)) \log (1-\eta(x))
$$

It measures the amount of uncertainty or chaos in $Y$ given $X=x$. As we know, it takes values between 0 (when $\eta(x) \in\{0,1\}$ ) and $\log 2$ (when $\eta(x)=1 / 2$ ), and is
concave in $\eta(x)$. We define the expected conditional entropy by

$$
\begin{aligned}
\mathcal{E} & =\mathbf{E}\{\mathcal{H}(\eta(X), 1-\eta(X))\} \\
& =-\mathbf{E}\{\eta(X) \log \eta(X)+(1-\eta(X)) \log (1-\eta(X))\}
\end{aligned}
$$

For brevity, we will refer to $\mathcal{E}$ as the entropy. As pointed out above, $\mathcal{E}=0$ if and only if $\eta(X) \in\{0,1\}$ with probability one. Thus, $\mathcal{E}$ and $L^{*}$ are related to each other.

## Theorem 3.2.

A. $\mathcal{E} \leq \mathcal{H}\left(L^{*}, 1-L^{*}\right)=-L^{*} \log L^{*}-\left(1-L^{*}\right) \log \left(1-L^{*}\right)$ (Fano's inequality; Fano (1952), see Cover and Thomas (1991, p. 39)).
B. $\mathcal{E} \geq-\log \left(1-L_{\mathrm{NN}}\right) \geq-\log \left(1-L^{*}\right)$.
C. $\mathcal{E} \leq \log 2-\frac{1}{2}\left(1-2 L_{\mathrm{NN}}\right) \leq \log 2-\frac{1}{2}\left(1-2 L^{*}\right)^{2}$.


Figure 3.3. Inequalities (A) and (B) of Theorem 3.2 are illustrated here.

Proof.
Part A. Define $A=\min (\eta(X), 1-\eta(X))$. Then

$$
\begin{aligned}
\mathcal{E}= & \mathbf{E}\{\mathcal{H}(A, 1-A)\} \\
\leq & \mathcal{H}(\mathbf{E} A, 1-\mathbf{E} A) \\
& \text { (because } \mathcal{H} \text { is concave, by Jensen's inequality) } \\
= & \mathcal{H}\left(L^{*}, 1-L^{*}\right)
\end{aligned}
$$

Part B.

$$
\begin{aligned}
\mathcal{E} & =-\mathbf{E}\{A \log A+(1-A) \log (1-A)\} \\
& \geq-\mathbf{E}\left\{\log \left(A^{2}+(1-A)^{2}\right)\right\} \quad \text { (by Jensen's inequality) } \\
& =-\mathbf{E}\{\log (1-2 A(1-A))\} \\
& \geq-\log (1-\mathbf{E}\{2 A(1-A)\}) \quad \text { (by Jensen's inequality) }
\end{aligned}
$$

$$
\begin{aligned}
& =-\log \left(1-L_{\mathrm{NN}}\right) \\
& \geq-\log \left(1-L^{*}\right)
\end{aligned}
$$

PART C. By the concavity of $\mathcal{H}(A, 1-A)$ as a function of $A$, and Taylor series expansion,

$$
\mathcal{H}(A, 1-A) \leq \log 2-\frac{1}{2}(2 A-1)^{2}
$$

Therefore, by Part A,

$$
\begin{aligned}
\mathcal{E} \leq & \log 2-\mathbf{E}\left\{\frac{1}{2}(2 A-1)^{2}\right\} \\
= & \log 2-\frac{1}{2}\left(1-2 L_{\mathrm{NN}}\right) \\
\leq & \log 2-\frac{1}{2}+2 L^{*}\left(1-L^{*}\right) \\
& (\text { by the Cover-Hart inequality }(3.1)) \\
= & \log 2-\frac{1}{2}\left(1-2 L^{*}\right)^{2}
\end{aligned}
$$

REMARK. The nearly monotone relationship between $\mathcal{E}$ and $L^{*}$ will see lots of uses. We warn the reader that near the origin, $L^{*}$ may decrease linearly in $\mathcal{E}$, but it may also decrease much faster than $\mathcal{E}^{209}$. Such wide variation was not observed in the relationship between $L^{*}$ and $L_{\mathrm{NN}}$ (where it is linear) or $L^{*}$ and $\rho$ (where it is between linear and quadratic).

### 3.6 Jeffreys' Divergence

Jeffreys' divergence (1948) is a symmetric form of the Kullback-Leibler (1951) divergence

$$
\delta_{\mathrm{KL}}=\mathbf{E}\left\{\eta(X) \log \frac{\eta(X)}{1-\eta(X)}\right\}
$$

It will be denoted by

$$
\mathcal{J}=\mathbf{E}\left\{(2 \eta(X)-1) \log \frac{\eta(X)}{1-\eta(X)}\right\}
$$

To understand $\mathcal{J}$, note that the function $(2 \eta-1) \log \frac{\eta}{1-\eta}$ is symmetric about $1 / 2$, convex, and has minimum ( 0 ) at $\eta=1 / 2$. As $\eta \downarrow 0, \eta \uparrow 1$, the function becomes unbounded. Therefore, $\mathcal{J}=\infty$ if $\mathbf{P}\{\eta(X) \in\{0,1\}\}>0$. For this reason, its
use in discrimination is necessarily limited. For generalizations of $\mathcal{J}$, see Rényi (1961), Burbea and Rao (1982), Taneja (1983; 1987), and Burbea (1984). It is thus impossible to bound $\mathcal{J}$ from above by a function of $L_{\mathrm{NN}}$ and/or $L^{*}$. However, lower bounds are easy to obtain. As $x \log ((1+x) /(1-x))$ is convex in $x$ and

$$
(2 \eta-1) \log \frac{\eta}{1-\eta}=|2 \eta-1| \log \left(\frac{1+|2 \eta-1|}{1-|2 \eta-1|}\right)
$$

we note that by Jensen's inequality,

$$
\begin{aligned}
\mathcal{J} & \geq \mathbf{E}\{|2 \eta(X)-1|\} \log \left(\frac{1+\mathbf{E}\{|2 \eta(X)-1|\}}{1-\mathbf{E}\{|2 \eta(X)-1|\}}\right) \\
& =\left(1-2 L^{*}\right) \log \left(\frac{1+\left(1-2 L^{*}\right)}{1-\left(1-2 L^{*}\right)}\right) \\
& =\left(1-2 L^{*}\right) \log \left(\frac{1-L^{*}}{L^{*}}\right) \\
& \geq 2\left(1-2 L^{*}\right)^{2} .
\end{aligned}
$$

The first bound cannot be universally bettered (it is achieved when $\eta(x)$ is constant over the space). Also, for fixed $L^{*}$, any value of $\mathcal{J}$ above the lower bound is possible for some distribution of $(X, Y)$. From the definition of $\mathcal{J}$, we see that $\mathcal{J}=0$ if and only if $\eta \equiv 1 / 2$ with probability one, or $L^{*}=1 / 2$.


Figure 3.4. This figure illustrates the above lower bound on Jeffreys' divergence in terms of the Bayes error.

Related bounds were obtained by Toussaint (1974b):

$$
\mathcal{J} \geq \sqrt{1-2 L_{\mathrm{NN}}} \log \left(\frac{1+\sqrt{1-2 L_{\mathrm{NN}}}}{1-\sqrt{1-2 L_{\mathrm{NN}}}}\right) \geq 2\left(1-2 L_{\mathrm{NN}}\right)
$$

The last bound is strictly better than our $L^{*}$ bound given above. See Problem 3.7.

## $3.7 \quad F$-Errors

The error measures discussed so far are all related to expected values of concave functions of $\eta(X)=\mathbf{P}\{Y=1 \mid X\}$. In general, if $F$ is a concave function on $[0,1]$,
we define the $F$-error corresponding to $(X, Y)$ by

$$
d_{F}(X, Y)=\mathbf{E}\{F(\eta(X))\}
$$

Examples of $F$-errors are
(a) the Bayes error $L^{*}: F(x)=\min (x, 1-x)$,
(b) the asymptotic nearest neighbor error $L_{\mathrm{NN}}: F(x)=2 x(1-x)$,
(c) the Matushita error $\rho: F(x)=\sqrt{x(1-x)}$,
(d) the expected conditional entropy $\mathcal{E}: F(x)=-x \log x-(1-x) \log (1-x)$,
(e) the negated Jeffreys' divergence $-\mathcal{J}: F(x)=-(2 x-1) \log \frac{x}{1-x}$.

Hashlamoun, Varshney, and Samarasooriya (1994) point out that if $F(x) \geq$ $\min (x, 1-x)$ for each $x \in[0,1]$, then the corresponding $F$-error is an upper bound on the Bayes error. The closer $F(x)$ is to $\min (x, 1-x)$, the tighter the upper bound is. For example, $F(x)=(1 / 2) \sin (\pi x) \leq 2 x(1-x)$ yields an upper bound tighter than $L_{\mathrm{NN}}$. All these errors share the property that the error increases if $X$ is transformed by an arbitrary function.

Theorem 3.3. Let $t: \mathcal{R}^{d} \rightarrow \mathcal{R}^{k}$ be an arbitrary measurable function. Then for any distribution of $(X, Y)$,

$$
d_{F}(X, Y) \leq d_{F}(t(X), Y)
$$

Proof. Define $\eta_{t}: \mathcal{R}^{k} \rightarrow[0,1]$ by $\eta_{t}(x)=\mathbf{P}\{Y=1 \mid t(X)=x\}$, and observe that

$$
\eta_{t}(t(X))=\mathbf{E}\{\eta(X) \mid t(X)\}
$$

Thus,

$$
\begin{aligned}
d_{F}(t(X), Y) & =\mathbf{E}\left\{F\left(\eta_{t}(t(X))\right)\right\} \\
& =\mathbf{E}\{F(\mathbf{E}\{\eta(X) \mid t(X)\})\} \\
& \geq \mathbf{E}\{\mathbf{E}\{F(\eta(X)) \mid t(X)\}\} \quad \text { (by Jensen's inequality) } \\
& =\mathbf{E}\{F(\eta(X))\}=d_{F}(X, Y)
\end{aligned}
$$

REmark. We also see from the proof that the $F$-error remains unchanged if the transformation $t$ is invertible. Theorem 3.3 states that $F$-errors are a bit like Bayes errors-when information is lost (by replacing $X$ by $t(X)$ ), $F$-errors increase.

### 3.8 The Mahalanobis Distance

Two conditional distributions with about the same covariance matrices and means that are far away from each other are probably so well separated that $L^{*}$ is small. An interesting measure of the visual distance between two random variables $X_{0}$ and $X_{1}$ is the so-called Mahalanobis distance (Mahalanobis, (1936)) given by

$$
\Delta=\sqrt{\left(m_{1}-m_{0}\right)^{T} \Sigma^{-1}\left(m_{1}-m_{0}\right)}
$$

where $m_{1}=\mathbf{E} X_{1}, m_{0}=\mathbf{E} X_{0}$, are the means, $\Sigma_{1}=\mathbf{E}\left\{\left(X_{1}-m_{1}\right)\left(X_{1}-m_{1}\right)^{T}\right\}$ and $\Sigma_{0}=\mathbf{E}\left\{\left(X_{0}-m_{0}\right)\left(X_{0}-m_{0}\right)^{T}\right\}$ are the covariance matrices, $\Sigma=p \Sigma_{1}+(1-$ p) $\Sigma_{0},(\cdot)^{T}$ is the transpose of a vector, and $p=1-p$ is a mixture parameter. If $\Sigma_{1}=\Sigma_{0}=\sigma^{2} I$, where $I$ is the identity matrix, then

$$
\Delta=\frac{\left\|m_{1}-m_{0}\right\|}{\sigma}
$$

is a scaled version of the distance between the means. If $\Sigma_{1}=\sigma_{1}^{2} I, \Sigma_{0}=\sigma_{0}^{2} I$, then

$$
\Delta=\frac{\left\|m_{1}-m_{0}\right\|}{\sqrt{p \sigma_{1}^{2}+(1-p) \sigma_{0}^{2}}}
$$

varies between $\left\|m_{1}-m_{0}\right\| / \sigma_{1}$ and $\left\|m_{1}-m_{0}\right\| / \sigma_{0}$ as $p$ changes from 1 to 0 . Assume that we have a discrimination problem in which given $Y=1, X$ is distributed as $X_{1}$, given $Y=0, X$ is distributed as $X_{0}$, and $p=\mathbf{P}\{Y=1\}, 1-p$ are the class probabilities. Then, interestingly, $\Delta$ is related to the Bayes error in a general sense. If the Mahalanobis distance between the class-conditional distributions is large, then $L^{*}$ is small.

Theorem 3.4. (Devijver and Kittler (1982, p. 166)). For all distributions of $(X, Y)$ for which $\mathbf{E}\left\{\|X\|^{2}\right\}<\infty$, we have

$$
L^{*} \leq L_{\mathrm{NN}} \leq \frac{2 p(1-p)}{1+p(1-p) \Delta^{2}}
$$

REMARK. For a distribution with mean $m$ and covariance matrix $\Sigma$, the Mahalanobis distance from a point $x \in \mathcal{R}^{d}$ to $m$ is

$$
\sqrt{(x-m)^{T} \Sigma^{-1}(x-m)}
$$

In one dimension, this is simply interpreted as distance from the mean as measured in units of standard deviation. The use of Mahalanobis distance in discrimination is based upon the intuitive notion that we should classify according to the class for which we are within the least units of standard deviations. At least, for distributions that look like nice globular clouds, such a recommendation may make sense.

Proof. First assume that $d=1$, that is, $X$ is real valued. Let $u$ and $c$ be real numbers, and consider the quantity $\mathbf{E}\left\{(u(X-c)-(2 \eta(X)-1))^{2}\right\}$. We will show that if $u$ and $c$ are chosen to minimize this number, then it satisfies

$$
\begin{equation*}
0 \leq \mathbf{E}\left\{(u(X-c)-(2 \eta(X)-1))^{2}\right\}=2\left(\frac{2 p(1-p)}{1+p(1-p) \Delta^{2}}-L_{\mathrm{NN}}\right) \tag{3.2}
\end{equation*}
$$

which proves the theorem for $d=1$. To see this, note that the expression $\mathbf{E}\left\{(u(X-c)-(2 \eta(X)-1))^{2}\right\}$ is minimized for $c=\mathbf{E} X-\mathbf{E}\{2 \eta(X)-1\} / u$. Then

$$
\begin{aligned}
& \mathbf{E}\left\{(u(X-c)-(2 \eta(X)-1))^{2}\right\} \\
& \quad=\operatorname{Var}\{2 \eta(X)-1\}+u^{2} \operatorname{Var}\{X\}-2 u \operatorname{Cov}\{X, 2 \eta(X)-1\},
\end{aligned}
$$

-where $\operatorname{Cov}\{X, Z\}=\mathbf{E}\{(X-\mathbf{E} X)(Z-\mathbf{E} Z)\}$-which is, in turn, minimized for $u=\operatorname{Cov}\{X, 2 \eta(X)-1\} / \operatorname{Var}\{X\}$. Straightforward calculation shows that (3.2) indeed holds.

To extend the inequality (3.2) to multidimensional problems, apply it to the one-dimensional decision problem $(Z, Y)$, where $Z=X^{T} \Sigma^{-1}\left(m_{1}-m_{0}\right)$. Then the theorem follows by noting that by Theorem 3.3,

$$
L_{\mathrm{NN}}(X, Y) \leq L_{\mathrm{NN}}(Z, Y)
$$

where $L_{\mathrm{NN}}(X, Y)$ denotes the nearest-neighbor error corresponding to $(X, Y)$.

In case $X_{1}$ and $X_{0}$ are both normal with the same covariance matrices, we have
Theorem 3.5. (Matushita (1973); see Problem 3.11). When $X_{1}$ and $X_{0}$ are multivariate normal random variables with $\Sigma_{1}=\Sigma_{0}=\Sigma$, then

$$
\rho=\mathbf{E}\{\sqrt{\eta(X)(1-\eta(X))}\}=\sqrt{p(1-p)} e^{-\Delta^{2} / 8}
$$

If the class-conditional densities $f_{1}$ and $f_{0}$ may be written as functions of $\left(x-m_{1}\right)^{T} \Sigma_{1}^{-1}\left(x-m_{1}\right)$ and $\left(x-m_{0}\right)^{T} \Sigma_{0}^{-1}\left(x-m_{0}\right)$ respectively, then $\Delta$ remains relatively tightly linked with $L^{*}$ (Mitchell and Krzanowski (1985)), but such distributions are the exception rather than the rule. In general, when $\Delta$ is small, it is impossible to deduce whether $L^{*}$ is small or not (see Problem 3.12).

## $3.9 f$-Divergences

We have defined error measures as the expected value of a concave function of $\eta(X)$. This makes it easier to relate these measures to the Bayes error $L^{*}$ and other error probabilities. In this section we briefly make the connection to the more classical statistical theory of distances between probability measures. A general concept of these distance measures, called $f$-divergences, was introduced by Csiszár (1967).

The corresponding theory is summarized in Vajda (1989). $F$-errors defined earlier may be calculated if one knows the class probabilities $p, 1-p$, and the conditional distributions $\mu_{0}, \mu_{1}$ of $X$, given $\{Y=0\}$ and $\{Y=1\}$, that is,

$$
\mu_{i}(A)=\mathbf{P}\{X \in A \mid Y=i\} \quad i=0,1
$$

For fixed class probabilities, an $F$-error is small if the two conditional distributions are "far away" from each other. A metric quantifying this distance may be defined as follows. Let $f:[0, \infty) \rightarrow \mathcal{R} \cup\{-\infty, \infty\}$ be a convex function with $f(1)=0$. The $f$-divergence between two probability measures $\mu$ and $v$ on $\mathcal{R}^{d}$ is defined by

$$
D_{f}(\mu, \nu)=\sup _{\mathcal{A}=\left\{A_{j}\right\}} \sum_{j} v\left(A_{j}\right) f\left(\frac{\mu\left(A_{j}\right)}{v\left(A_{j}\right)}\right),
$$

where the supremum is taken over all finite measurable partitions $\mathcal{A}$ of $\mathcal{R}^{d}$. If $\lambda$ is a measure dominating $\mu$ and $\nu$-that is, both $\mu$ and $\nu$ are absolutely continuous with respect to $\lambda$-and $p=d \mu / d \lambda$ and $q=d \nu / d \lambda$ are the corresponding densities, then the $f$-divergence may be put in the form

$$
D_{f}(\mu, \nu)=\int q(x) f\left(\frac{p(x)}{q(x)}\right) \lambda(d x)
$$

Clearly, this quantity is independent of the choice of $\lambda$. For example, we may take $\lambda=\mu+\nu$. If $\mu$ and $\nu$ are absolutely continuous with respect to the Lebesgue measure, then $\lambda$ may be chosen to be the Lebesgue measure. By Jensen's inequality, $D_{f}(\mu, \nu) \geq 0$, and $D_{f}(\mu, \mu)=0$.

An important example of $f$-divergences is the total variation, or variational distance obtained by choosing $f(x)=|x-1|$, yielding

$$
V(\mu, v)=\sup _{\mathcal{A}=\left\{A_{j}\right\}} \sum_{j}\left|\mu\left(A_{j}\right)-v\left(A_{j}\right)\right| .
$$

For this divergence, the equivalence of the two definitions is stated by Scheffés theorem (see Problem 12.13).

Theorem 3.6. (Scheffé (1947)).

$$
V(\mu, v)=2 \sup _{A}|\mu(A)-v(A)|=\int|p(x)-q(x)| \lambda(d x),
$$

where the supremum is taken over all Borel subsets of $\mathcal{R}^{d}$.
Another important example is the Hellinger distance, given by $f(x)=(1-$ $\sqrt{x})^{2}$ :

$$
\begin{aligned}
H_{2}(\mu, \nu) & =\sup _{\mathcal{A}=\left\{A_{j}\right\}} 2\left(1-\sum_{j} \sqrt{\mu\left(A_{j}\right) v\left(A_{j}\right)}\right) \\
& =2\left(1-\int \sqrt{p(x) q(x)} \lambda(d x)\right)
\end{aligned}
$$

The quantity $I_{2}(\mu, \nu)=\int \sqrt{p(x) q(x)} \lambda(d x)$ is often called the Hellinger integral. We mention two useful inequalities in this respect. For the sake of simplicity, we state their discrete form. (The integral forms are analogous, see Problem 3.21.)

Lemma 3.1. (LECAM (1973)). For positive sequences $a_{i}$ and $b_{i}$, both summing to one,

$$
\sum_{i} \min \left(a_{i}, b_{i}\right) \geq \frac{1}{2}\left(\sum_{i} \sqrt{a_{i} b_{i}}\right)^{2}
$$

Proof. By the Cauchy-Schwarz inequality,

$$
\sum_{i: a_{i}<b_{i}} \sqrt{a_{i} b_{i}} \leq \sqrt{\sum_{i: a_{i}<b_{i}} a_{i}} \sqrt{\sum_{i: a_{i}<b_{i}} b_{i}} \leq \sqrt{\sum_{i: a_{i}<b_{i}} a_{i}}
$$

This, together with the inequality $(x+y)^{2} \leq 2 x^{2}+2 y^{2}$, and symmetry, implies

$$
\begin{aligned}
\left(\sum_{i} \sqrt{a_{i} b_{i}}\right)^{2} & =\left(\sum_{i: a_{i}<b_{i}} \sqrt{a_{i} b_{i}}+\sum_{i: a_{i} \geq b_{i}} \sqrt{a_{i} b_{i}}\right)^{2} \\
& \leq 2\left(\sum_{i: a_{i}<b_{i}} \sqrt{a_{i} b_{i}}\right)^{2}+2\left(\sum_{i: a_{i} \geq b_{i}} \sqrt{a_{i} b_{i}}\right)^{2} \\
& \leq 2\left(\sum_{i: a_{i}<b_{i}} a_{i}+\sum_{i: a_{i} \geq b_{i}} b_{i}\right) \\
& =2 \sum_{i}^{\min \left(a_{i}, b_{i}\right)}
\end{aligned}
$$

Lemma 3.2. (Devroye and Györfi (1985, p. 225)). Let $a_{1}, \ldots, a_{k}, b_{1}, \ldots b_{k}$ be nonnegative numbers such that $\sum_{i=1}^{k} a_{i}=\sum_{i=1}^{k} b_{i}=1$. Then

$$
\sum_{i=1}^{k} \sqrt{a_{i} b_{i}} \leq 1-\frac{\left(\sum_{i=1}^{k}\left|a_{i}-b_{i}\right|\right)^{2}}{8}
$$

Proof.

$$
\left(\sum_{i=1}^{k}\left|a_{i}-b_{i}\right|\right)^{2} \leq\left(\sum_{i=1}^{k}\left(\sqrt{a_{i}}-\sqrt{b_{i}}\right)^{2}\right)\left(\sum_{i=1}^{k}\left(\sqrt{a_{i}}+\sqrt{b_{i}}\right)^{2}\right)
$$

(by the Cauchy-Schwarz inequality)

$$
\begin{aligned}
& \leq 4 \sum_{i=1}^{k}\left(\sqrt{a_{i}}-\sqrt{b_{i}}\right)^{2} \\
& =8\left(1-\sum_{i=1}^{k} \sqrt{a_{i} b_{i}}\right),
\end{aligned}
$$

which proves the lemma.

Information divergence is obtained by taking $f(x)=x \log x$ :

$$
I(\mu, \nu)=\sup _{\mathcal{A}=\left\{A_{j}\right\}} \sum_{j} \mu\left(A_{j}\right) \log \left(\frac{\mu\left(A_{j}\right)}{\nu\left(A_{j}\right)}\right)
$$

$I(\mu, \nu)$ is also called the Kullback-Leibler number.
Our last example is the $\chi^{2}$-divergence, defined by $f(x)=(x-1)^{2}$ :

$$
\begin{aligned}
\chi^{2}(\mu, \nu) & =\sup _{\mathcal{A}=\left\{A_{j}\right\}} \sum_{j} \frac{\left(\mu\left(A_{j}\right)-v\left(A_{j}\right)\right)^{2}}{v\left(A_{j}\right)} \\
& =\int \frac{p^{2}(x)}{q(x)} \lambda(d x)-1
\end{aligned}
$$

Next, we highlight the connection between $F$-errors and $f$-divergences. Let $\mu_{0}$ and $\mu_{1}$ denote the conditional distributions of $X$ given $\{Y=0\}$ and $\{Y=1\}$. Assume that the class probabilities are equal: $p=1 / 2$. If $F$ is a concave function, then the $F$-error $d_{F}(X, Y)$ may be written as

$$
d_{F}(X, Y)=F\left(\frac{1}{2}\right)-D_{f}\left(\mu_{0}, \mu_{1}\right)
$$

where

$$
f(x)=-\frac{1}{2} F\left(\frac{x}{1+x}\right)(1+x)+F\left(\frac{1}{2}\right)
$$

and $D_{f}$ is the corresponding $f$-divergence. It is easy to see that $f$ is convex, whenever $F$ is concave. A special case of this correspondence is

$$
L^{*}=\frac{1}{2}\left(1-\frac{1}{2} V\left(\mu_{0}, \mu_{1}\right)\right)
$$

if $p=1 / 2$. Also, it is easy to verify, that

$$
\rho=\sqrt{p(1-p)} I_{2}\left(\mu_{0}, \mu_{1}\right)
$$

where $\rho$ is the Matushita error. For further connections, we refer the reader to the exercises.

## Problems and Exercises

Problem 3.1. Show that for every $\left(l_{1}, l^{*}\right)$ with $0 \leq l^{*} \leq l_{1} \leq 2 l^{*}\left(1-l^{*}\right) \leq 1 / 2$, there exists a distribution of $(X, Y)$ with $L_{\mathrm{NN}}=l_{1}$ and $L^{*}=l^{*}$. Therefore, the Cover-Hart inequalities are not universally improvable.

Problem 3.2. Tightness of the bounds. Theorem 3.1 cannot be improved.
(1) Show that for all $\alpha \in[0,1 / 2]$, there exists a distribution of ( $X, Y$ ) such that $L_{\mathrm{NN}}=L^{*}=\alpha$.
(2) Show that for all $\alpha \in[0,1 / 2]$, there exists a distribution of ( $X, Y$ ) such that $L_{\mathrm{NN}}=\alpha, L^{*}=\frac{1}{2}-\frac{1}{2} \sqrt{1-2 \alpha}$.
(3) Show that for all $\alpha \in[0,1 / 2]$, there exists a distribution of ( $X, Y$ ) such that $L^{*}=\rho=\alpha$.
(4) Show that for all $\alpha \in[0,1 / 2]$, there exists a distribution of ( $X, Y$ ) such that $L_{\mathrm{NN}}=\alpha, L^{*}=\frac{1}{2}-\frac{1}{2} \sqrt{1-4 \rho^{2}}$.

Problem 3.3. Show that $\mathcal{E} \geq L^{*}$.
Problem 3.4. For any $\alpha \leq 1$, find a sequence of distributions of ( $X_{n}, Y_{n}$ ) having expected conditional entropies $\mathcal{E}_{n}$ and Bayes errors $L_{n}^{*}$ such that $L_{n}^{*} \rightarrow 0$ as $n \rightarrow \infty$, and $\mathcal{E}_{n}$ decreases to zero at the same rate as $\left(L_{n}^{*}\right)^{\alpha}$.

PROBLEM 3.5. CONCAVITY of ERROR MEASURES. Let $Y$ denote the mixture random variable taking the value $Y_{1}$ with probability $p$ and the value $Y_{2}$ with probability $1-p$. Let $X$ be a fixed $\mathcal{R}^{d}$-valued random variable, and define $\eta_{1}(x)=\mathbf{P}\left\{Y_{1}=1 \mid X=x\right\}, \eta_{2}(x)=\mathbf{P}\left\{Y_{2}=1 \mid X=x\right\}$, where $Y_{1}, Y_{2}$ are Bernoulli random variables. Clearly, $\eta(x)=p \eta_{1}(x)+(1-p) \eta_{2}(x)$. Which of the error measures $L^{*}, \rho, L_{\mathrm{NN}}, \mathcal{E}$ are concave in $p$ for fixed joint distribution of $X, Y_{1}, Y_{2}$ ? Can every discrimination problem $(X, Y)$ be decomposed this way for some $Y_{1}, p, Y_{2}$, where $\eta_{1}(x), \eta_{2}(x) \in\{0,1\}$ for all $x$ ? If not, will the condition $\eta_{1}(x), \eta_{2}(x) \in\{0,1 / 2,1\}$ for all $x$ do?

Problem 3.6. Show that for every $l^{*} \in[0,1 / 2]$, there exists a distribution of $(X, Y)$ with $L^{*}=l^{*}$ and $\mathcal{E}=\mathcal{H}\left(L^{*}, 1-L^{*}\right)$. Thus, Fano's inequality is tight.

Problem 3.7. Toussaint's inequalities (1974b). Mimic a proof in the text to show that

$$
\mathcal{J} \geq \sqrt{1-2 L_{\mathrm{NN}}} \log \left(\frac{1+\sqrt{1-2 L_{\mathrm{NN}}}}{1-\sqrt{1-2 L_{\mathrm{NN}}}}\right) \geq 2\left(1-2 L_{\mathrm{NN}}\right)
$$

Problem 3.8. Show that $L^{*} \leq e^{-\delta_{C}}$, where $\delta_{C}$ is Chernoff's measure of affinity with parameter $\alpha \in(0,1)$.

Problem 3.9. Prove that $L^{*}=p-\mathbf{E}\left\{(2 \eta(X)-1)_{+}\right\}$, where $p=\mathbf{P}\{Y=1\}$, and $(x)_{+}=$ $\max (x, 0)$.

Problem 3.10. Show that $\mathcal{J} \geq-2 \log \rho-2 \mathcal{H}(p, 1-p)$, where $p=\mathbf{P}\{Y=1\}$ (Toussaint (1974b)).

Problem 3.11. Let $f_{1}$ and $f_{0}$ be two multivariate normal densities with means $m_{0}, m_{1}$ and common covariance matrix $\Sigma$. If $\mathbf{P}\{Y=1\}=p$, and $f_{1}, f_{0}$ are the conditional densities of
$X$ given $Y=1$ and $Y=0$ respectively, then show that

$$
\rho=\sqrt{p(1-p)} e^{-\Delta^{2} / 8}
$$

where $\rho$ is the Matushita error and $\Delta$ is the Mahalanobis distance.
Problem 3.12. For every $\delta \in[0, \infty)$ and $l^{*} \in[0,1 / 2]$ with $l^{*} \leq 2 /\left(4+\delta^{2}\right)$, find distributions $\mu_{0}$ and $\mu_{1}$ for $X$ given $Y=1, Y=0$, such that the Mahalanobis distance $\Delta=\delta$, yet $L^{*}=l^{*}$. Therefore, the Mahalanobis distance is not universally related to the Bayes risk.

Problem 3.13. Show that the Mahalanobis distance $\Delta$ is invariant under linear invertible transformations of $X$.

Problem 3.14. Lissack and Fu (1976) have suggested the measures

$$
\delta_{\mathrm{LF}}=\mathbb{E}\left\{|2 \eta(X)-1|^{\alpha}\right\}, \quad 0<\alpha<\infty .
$$

For $\alpha=1$, this is twice the Kolmogorov distance $\delta_{\text {ко }}$. Show the following:
(1) If $0<\alpha \leq 1$, then $\frac{1}{2}\left(1-\delta_{\mathrm{LF}}\right) \leq L^{*} \leq\left(1-\delta_{\mathrm{LF}}^{1 / \alpha}\right)$.
(2) If $1 \leq \alpha<\infty$, then $\frac{1}{2}\left(1-\delta_{\text {LF }}^{1 / \alpha}\right) \leq L^{*} \leq\left(1-\delta_{\text {LF }}\right)$.

Problem 3.15. Hashlamoun, Varshney, and Samarasooriya (1994) suggest using the $F$ error with the function

$$
F(x)=\frac{1}{2} \sin (\pi x) e^{-1.8063\left(x-\frac{1}{2}\right)^{2}}
$$

to obtain tight upper bounds on $L^{*}$. Show that $F(x) \geq \min (x, 1-x)$, so that the corresponding $F$-error is indeed an upper bound on the Bayes risk.
$\operatorname{Problem}$ 3.16. Prove that $L^{*} \leq \max (p(1-p))\left(1-\frac{1}{2} V\left(\mu_{0}, \mu_{1}\right)\right)$.
Problem 3.17. Prove that $L^{*} \leq \sqrt{p(1-p)} I_{2}\left(\mu_{0} \mu_{1}\right)$. $\operatorname{Hint}: \min (a, b) \leq \sqrt{a b}$.
Problem 3.18. Assume that the components of $X=\left(X^{(1)}, \ldots, X^{(d)}\right)$ are conditionally independent (given $Y$ ), and identically distributed, that is, $\mathbf{P}\left\{X^{(i)} \in A \mid Y=j\right\}=v_{j}(A)$ for $i=1, \ldots, d$ and $j=0,1$. Use the previous exercise to show that

$$
L^{*} \leq \sqrt{p(1-p)}\left(I_{2}\left(v_{0}, v_{1}\right)\right)^{d} .
$$

Problem 3.19. Show that $\chi^{2}\left(\mu_{1}, \mu_{2}\right) \geq I\left(\mu_{1}, \mu_{2}\right)$. Hint: $x-1 \geq \log x$.
Problem 3.20. Show the following analog of Theorem 3.3. Let $t: \mathcal{R}^{d} \rightarrow \mathcal{R}^{k}$ be a measurable function, and $\mu, \nu$ probability measures on $\mathcal{R}^{d}$. Define the measures $\mu_{t}$ and $\nu_{t}$ on $\mathcal{R}^{k}$ by $\mu_{t}(A)=\mu\left(t^{-1}(A)\right)$ and $v_{t}(A)=v\left(t^{-1}(A)\right)$. Show that for any convex function $f$, $D_{f}(\mu, \nu) \geq D_{f}\left(\mu_{t}, \nu_{t}\right)$.

Problem 3.21. Prove the following connections between the Hellinger integral and the total variation:

$$
\left(1-\frac{1}{2} V(\mu, \nu)\right) \geq \frac{1}{2}\left(I_{2}(\mu, \nu)\right)^{2}
$$

and

$$
(V(\mu, \nu))^{2} \leq 8\left(1-I_{2}(\mu, \nu)\right)
$$

Hint: Proceed analogously to Lemmas 3.1 and 3.2.

Problem 3.22. Pinsker's inequality. Show that

$$
(V(\mu, \nu))^{2} \leq 2 I(\mu, \nu)
$$

(Csiszár (1967), Kullback (1967), and Kemperman (1969)). Hint: First prove the inequality if $\mu$ and $\nu$ are concentrated on the same two atoms. Then define $A=\{x: p(x) \geq q(x)\}$, and the measures $\mu^{*}, \nu^{*}$ on the set $\{0,1\}$ by $\mu^{*}(0)=1-\mu^{*}(1)=\mu(A)$ and $\nu^{*}(0)=1-\nu^{*}(1)=$ $v(A)$, and apply the previous result. Conclude by pointing out that Scheffés theorem states $V\left(\mu^{*}, \nu^{*}\right)=V(\mu, \nu)$, and that $I\left(\mu^{*}, v^{*}\right) \leq I(\mu, v)$.

## 4

## Linear Discrimination

In this chapter, we split the space by a hyperplane and assign a different class to each halfspace. Such rules offer tremendous advantages - they are easy to interpret as each decision is based upon the sign of $\sum_{i=1}^{d} a_{i} x^{(i)}+a_{0}$, where $x=\left(x^{(1)}, \ldots, x^{(d)}\right)$ and the $a_{i}$ 's are weights. The weight vector determines the relative importance of the components. The decision is also easily implemented-in a standard software solution, the time of a decision is proportional to $d$-and the prospect that a small chip can be built to make a virtually instantaneous decision is particularly exciting.

Rosenblatt (1962) realized the tremendous potential of such linear rules and called them perceptrons. Changing one or more weights as new data arrive allows us to quickly and easily adapt the weights to new situations. Training or learning patterned after the human brain thus became a reality. This chapter merely looks at some theoretical properties of perceptrons. We begin with the simple onedimensional situation, and deal with the choice of weights in $\mathcal{R}^{d}$ further on. Unless one is terribly lucky, linear discrimination rules cannot provide error probabilities close to the Bayes risk, but that should not diminish the value of this chapter. Linear discrimination is at the heart of nearly every successful pattern recognition method, including tree classifiers (Chapters 20 and 21), generalized linear classifiers (Chapter 17), and neural networks (Chapter 30). We also encounter for the first time rules in which the parameters (weights) are dependent upon the data.


Figure 4.1. Rosenblatt's perceptron. The decision is based upon a linear combination of the components of the input vector.

### 4.1 Univariate Discrimination and Stoller Splits

As an introductory example, let $X$ be univariate. The crudest and simplest possible rule is the linear discrimination rule

$$
g(x)= \begin{cases}y^{\prime} & \text { if } x \leq x^{\prime} \\ 1-y^{\prime} & \text { otherwise }\end{cases}
$$

where $x^{\prime}$ is a split point and $y^{\prime} \in\{0,1\}$ is a class. In general, $x^{\prime}$ and $y^{\prime}$ are measurable functions of the data $D_{n}$. Within this class of simple rules, there is of course a best possible rule that can be determined if we know the distribution. Assume for example that $(X, Y)$ is described in the standard manner: let $\mathbf{P}\{Y=1\}=p$. Given $Y=1, X$ has a distribution function $F_{1}(x)=\mathbf{P}\{X \leq x \mid Y=1\}$, and given $Y=0$, $X$ has a distribution function $F_{0}(x)=\mathbf{P}\{X \leq x \mid Y=0\}$, where $F_{0}$ and $F_{1}$ are the class-conditional distribution functions. Then a theoretically optimal rule is determined by the split point $x^{*}$ and class $y^{*}$ given by

$$
\left(x^{*}, y^{*}\right)=\underset{\left(x^{\prime}, y^{\prime}\right)}{\arg \min } \mathbf{P}\{g(X) \neq Y\}
$$

(the minimum is always reached if we allow the values $x^{\prime}=\infty$ and $x^{\prime}=-\infty$ ). We call the corresponding minimal probability of error $L$ and note that

$$
\begin{aligned}
L= & \inf _{\left(x^{\prime}, y^{\prime}\right)}\left\{I_{\left\{y^{\prime}=0\right\}}\left(p F_{1}\left(x^{\prime}\right)+(1-p)\left(1-F_{0}\left(x^{\prime}\right)\right)\right)\right. \\
& \left.+I_{\left\{y^{\prime}=1\right\}}\left(p\left(1-F_{1}\left(x^{\prime}\right)\right)+(1-p) F_{0}\left(x^{\prime}\right)\right)\right\} .
\end{aligned}
$$

A split defined by $\left(x^{*}, y^{*}\right)$ will be called a theoretical Stoller split (Stoller (1954)).

Lemma 4.1. $L \leq 1 / 2$ with equality if and only if $L^{*}=1 / 2$.
Proof. Take $\left(x^{\prime}, y^{\prime}\right)=(-\infty, 0)$. Then the probability of error is $1-p=\mathbf{P}\{Y=0\}$. Take $\left(x^{\prime}, y^{\prime}\right)=(-\infty, 1)$. Then the probability of error is $p$. Clearly,

$$
L^{*} \leq L \leq \min (p, 1-p)
$$

This proves the first part of the lemma. For the second part, if $L=1 / 2$, then $p=1 / 2$, and for every $x, p F_{1}(x)+(1-p)\left(1-F_{0}(x)\right) \geq 1 / 2$ and $p\left(1-F_{1}(x)\right)+$ $(1-p) F_{0}(x) \geq 1 / 2$. The first inequality implies $p F_{1}(x)-(1-p) F_{0}(x) \geq p-1 / 2$, while the second implies $p F_{1}(x)-(1-p) F_{0}(x) \leq p-1 / 2$. Therefore, $L=1 / 2$ means that for every $x, p F_{1}(x)-(1-p) F_{0}(x)=p-1 / 2$. Thus, for all $x$, $F_{1}(x)=F_{0}(x)$, and therefore $L^{*}=1 / 2$.

## Lemma 4.2.

$$
L=\frac{1}{2}-\sup _{x}\left|p F_{1}(x)-(1-p) F_{0}(x)-p+\frac{1}{2}\right| .
$$

In particular, if $p=1 / 2$, then

$$
L=\frac{1}{2}-\frac{1}{2} \sup _{x}\left|F_{1}(x)-F_{0}(x)\right|
$$

Proof. Set $\rho(x)=p F_{1}(x)-(1-p) F_{0}(x)$. Then, by definition,

$$
\begin{aligned}
L= & \inf _{x} \min \{\rho(x)+1-p, p-\rho(x)\} \\
= & \frac{1}{2}-\sup _{x}\left|\rho(x)-p+\frac{1}{2}\right| \\
& (\text { since } \min \{a, b\}=(a+b-|a-b|) / 2) .
\end{aligned}
$$

The last property relates the quality of theoretical Stoller splits to the Kol-mogorov-Smirnov distance $\sup _{x}\left|F_{1}(x)-F_{0}(x)\right|$ between the class-conditional distribution functions. As a fun exercise, consider two classes with means $m_{0}=$ $\mathbf{E}\{X \mid Y=0\}, m_{1}=\mathbf{E}\{X \mid Y=1\}$, and variances $\sigma_{0}^{2}=\operatorname{Var}\{X \mid Y=0\}$ and $\sigma_{1}^{2}=$ $\operatorname{Var}\{X \mid Y=1\}$. Then the following inequality holds.

## Theorem 4.1.

$$
L^{*} \leq L \leq \frac{1}{1+\frac{\left(m_{0}-m_{1}\right)^{2}}{\left(\sigma_{0}+\sigma_{1}\right)^{2}}}
$$

REmark. When $p=1 / 2$, Chernoff (1971) proved

$$
L \leq \frac{1}{2+2 \frac{\left(m_{0}-m_{1}\right)^{2}}{\left(\sigma_{0}+\sigma_{1}\right)^{2}}} .
$$

Moreover, Becker (1968) pointed out that this is the best possible bound (see Problem 4.2).

Proof. Assume without loss of generality that $m_{0}<m_{1}$. Clearly, $L$ is smaller than the probability of error for the rule that decides 0 when $x \leq m_{0}+\Delta_{0}$, and 1 otherwise, where $m_{1}-m_{0}=\Delta_{0}+\Delta_{1}, \Delta_{0}, \Delta_{1}>0$.


Figure 4.2. The split providing the bound of Theorem 4.1.

The probability of error of the latter rule is

$$
\begin{aligned}
& p F_{1}\left(m_{0}+\Delta_{0}\right)+(1-p)\left(1-F_{0}\left(m_{0}+\Delta_{0}\right)\right) \\
& \quad=p \mathbf{P}\left\{X \leq m_{1}-\Delta_{1} \mid Y=1\right\}+(1-p) \mathbf{P}\left\{X>m_{0}+\Delta_{0} \mid Y=0\right\} \\
& \quad \leq p \frac{\sigma_{1}^{2}}{\sigma_{1}^{2}+\Delta_{1}^{2}}+(1-p) \frac{\sigma_{0}^{2}}{\sigma_{0}^{2}+\Delta_{0}^{2}}
\end{aligned}
$$

(by the Chebyshev-Cantelli inequality; see Appendix, Theorem A.17)

$$
=\frac{p}{1+\frac{\Delta_{1}^{2}}{\sigma_{1}^{2}}}+\frac{1-p}{1+\frac{\Delta_{0}^{2}}{\sigma_{0}^{2}}}
$$

$$
\left(\operatorname{take} \Delta_{1}=\left(\sigma_{1} / \sigma_{0}\right) \Delta_{0}, \text { and } \Delta_{0}=\left|m_{1}-m_{0}\right| \sigma_{0} /\left(\sigma_{0}+\sigma_{1}\right)\right)
$$

$$
=\frac{1}{1+\frac{\left(m_{0}-m_{1}\right)^{2}}{\left(\sigma_{0}+\sigma_{1}\right)^{2}}} .
$$

We have yet another example of the principle that well-separated classes yield small values for $L$ and thus $L^{*}$. Separation is now measured in terms of the largeness of $\left|m_{1}-m_{0}\right|$ with respect to $\sigma_{0}+\sigma_{1}$. Another inequality in the same spirit is given in Problem 4.1.

The limitations of theoretical Stoller splits are best shown in a simple example. Consider a uniform [0,1] random variable $X$, and define

$$
Y= \begin{cases}1 & \text { if } 0 \leq X \leq \frac{1}{3}+\epsilon \\ 0 & \text { if } \frac{1}{3}+\epsilon<X \leq \frac{2}{3}-\epsilon \\ 1 & \text { if } \frac{2}{3}-\epsilon \leq X \leq 1\end{cases}
$$

for some small $\epsilon>0$. As $Y$ is a function of $X$, we have $L^{*}=0$. If we are forced to make a trivial $X$-independent decision, then the best we can do is to set $g(x) \equiv 1$.

The probability of error is $\mathbf{P}\{1 / 3+\epsilon<X<2 / 3-\epsilon\}=1 / 3-2 \epsilon$. Consider next a theoretical Stoller split. One sees quickly that the best split occurs at $x^{\prime}=0$ or $x^{\prime}=1$, and thus that $L=1 / 3-2 \epsilon$. In other words, even the best theoretical split is superfluous. Note also that in the above example, $m_{0}=m_{1}=1 / 2$ so that the inequality of Theorem 4.1 says $L \leq 1$-it degenerates.

We now consider what to do when a split must be data-based. Stoller (1954) suggests taking $\left(x^{\prime}, y^{\prime}\right)$ such that the empirical error is minimal. He finds $\left(x^{\prime}, y^{\prime}\right)$ such that

$$
\left(x^{\prime}, y^{\prime}\right)=\underset{(x, y) \in \mathcal{R} \times\{0,1\}}{\arg \min } \frac{1}{n} \sum_{i=1}^{n}\left(I_{\left\{X_{i} \leq x, Y_{i} \neq y\right\}}+I_{\left\{X_{i}>x, Y_{i} \neq 1-y\right\}}\right)
$$

( $x^{\prime}$ and $y^{\prime}$ are now random variables, but in spite of our convention, we keep the lowercase notation for now.) We will call this Stoller's rule. The split is referred to as an empirical Stoller split. Denote the set $\{(-\infty, x] \times\{y\}\} \cup\{(x, \infty) \times\{1-y\}\}$ by $C(x, y)$. Then

$$
\left(x^{\prime}, y^{\prime}\right)=\underset{(x, y)}{\arg \min } v_{n}(C(x, y)),
$$

where $v_{n}$ is the empirical measure for the data $D_{n}=\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, that is, for every measurable set $A \in \mathcal{R} \times\{0,1\}, v_{n}(A)=(1 / n) \sum_{i=1}^{n} I_{\left\{\left(X_{i}, Y_{i}\right) \in A\right\}}$. Denoting the measure of $(X, Y)$ in $\mathcal{R} \times\{0,1\}$ by $v$, it is clear that $\mathbf{E}\left\{v_{n}(C)\right\}=$ $v(C)=\mathbf{P}\{X \leq x, Y \neq y\}+\mathbf{P}\{X>x, Y \neq 1-y\}$. Let $L_{n}=\mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}$ be the error probability of the splitting rule $g_{n}$ with the data-dependent choice ( $x^{\prime}, y^{\prime}$ ) given above, conditioned on the data. Then

$$
\begin{aligned}
L_{n}= & v\left(C\left(x^{\prime}, y^{\prime}\right)\right) \\
= & v\left(C\left(x^{\prime}, y^{\prime}\right)\right)-v_{n}\left(C\left(x^{\prime}, y^{\prime}\right)\right)+v_{n}\left(C\left(x^{\prime}, y^{\prime}\right)\right) \\
\leq & \sup _{(x, y)}\left(v(C(x, y))-v_{n}(C(x, y))\right)+v_{n}\left(C\left(x^{*}, y^{*}\right)\right) \\
& \left(\text { where }\left(x^{*}, y^{*}\right) \text { minimizes } v(C(x, y)) \text { over all }(x, y)\right) \\
\leq & 2 \sup _{(x, y)}\left|v(C(x, y))-v_{n}(C(x, y))\right|+v\left(C\left(x^{*}, y^{*}\right)\right) \\
= & 2 \sup _{(x, y)}\left|v(C(x, y))-v_{n}(C(x, y))\right|+L .
\end{aligned}
$$

From the next theorem we see that the supremum above is small even for moderately large $n$, and therefore, Stoller's rule performs closely to the best split regardless of the distribution of $(X, Y)$.

Theorem 4.2. For Stoller's rule, and $\epsilon>0$,

$$
\mathbf{P}\left\{L_{n}-L \geq \epsilon\right\} \leq 4 e^{-n \epsilon^{2} / 2}
$$

and

$$
\mathbf{E}\left\{L_{n}-L\right\} \leq \sqrt{\frac{2 \log (4 e)}{n}}
$$

Proof. By the inequality given just above the theorem,

$$
\begin{aligned}
\mathbf{P}\left\{L_{n}-L \geq \epsilon\right\} \leq & \mathbf{P}\left\{\sup _{(x, y)}\left|\nu(C(x, y))-v_{n}(C(x, y))\right| \geq \frac{\epsilon}{2}\right\} \\
\leq & \mathbf{P}\left\{\sup _{x}\left|\nu(C(x, 0))-v_{n}(C(x, 0))\right| \geq \frac{\epsilon}{2}\right\} \\
& +\mathbf{P}\left\{\sup _{x}\left|\nu(C(x, 1))-v_{n}(C(x, 1))\right| \geq \frac{\epsilon}{2}\right\} \\
\leq & 4 e^{-2 n(\epsilon / 2)^{2}}
\end{aligned}
$$

by a double application of Massart's (1990) tightened version of the Dvoretzky-Kiefer-Wolfowitz inequality (1956) (Theorem 12.9). See Problem 4.5. We do not prove this inequality here, but we will thoroughly discuss several such inequalities in Chapter 12 in a greater generality. The second inequality follows from the first via Problem 12.1.

The probability of error of Stoller's rule is uniformly close to $L$ over all possible distributions. This is just a preview of things to come, as we may be able to obtain good performance guarantees within a limited class of rules.

### 4.2 Linear Discriminants

Rosenblatt's perceptron (Rosenblatt (1962); see Nilsson (1965) for a good discussion) is based upon a dichotomy of $\mathcal{R}^{d}$ into two parts by a hyperplane. The linear discrimination rule with weights $a_{0}, a_{1}, \ldots, a_{d}$ is given by

$$
g(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{d} a_{i} x^{(i)}+a_{0}>0 \\ 0 & \text { otherwise }\end{cases}
$$

where $x=\left(x^{(1)}, \ldots, x^{(d)}\right)$.


Figure 4.3. A linear discriminant in $\mathcal{R}^{2}$ that correctly classifies all but four data points.

Its probability of error is for now denoted by $L\left(a, a_{0}\right)$, where $a=\left(a_{1}, \ldots, a_{d}\right)$. Again, we set

$$
L=\inf _{a \in \mathcal{R}^{d}, a_{0} \in \mathcal{R}} L\left(a, a_{0}\right)
$$

for the best possible probability of error within this class. Let the class-conditional distribution functions of $a_{1} X^{(1)}+\cdots+a_{d} X^{(d)}$ be denoted by $F_{0, a}$ and $F_{1, a}$, depending upon whether $Y=0$ or $Y=1$. For $L\left(a, a_{0}\right)$, we may use the bounds of Lemma 4.2 , and apply them to $F_{0, a}$ and $F_{1, a}$. Thus,

$$
L=\frac{1}{2}-\sup _{a} \sup _{x}\left|p F_{1, a}(x)-(1-p) F_{0, a}(x)-p+\frac{1}{2}\right|,
$$

which, for $p=1 / 2$, reduces to

$$
L=\frac{1}{2}-\frac{1}{2} \sup _{a} \sup _{x}\left|F_{1, a}(x)-F_{0, a}(x)\right|
$$

Therefore, $L=1 / 2$ if and only if $p=1 / 2$ and for all $a, F_{1, a} \equiv F_{0, a}$. Then apply the following simple lemma.

Lemma 4.3. (Cramér and Wold (1936)). $X_{1}$ and $X_{2}$, random variables taking values in $\mathcal{R}^{d}$, are identically distributed if and only if $a^{T} X_{1}$ and $a^{T} X_{2}$ have the same distribution for all vectors $a \in \mathcal{R}^{d}$.

Proof. Two random variables have identical distributions if and only if they have the same characteristic function-see, for example, Lukács and Laha (1964). Now, the characteristic function of $X_{1}=\left(X_{1}^{(1)}, \ldots, X_{1}^{(d)}\right)$ is

$$
\begin{aligned}
\psi_{1}(a) & =\psi_{1}\left(a_{1}, \ldots, a_{d}\right) \\
& =\mathbb{E}\left\{e^{i\left(a_{1} X_{1}^{(1)}+\cdots+a_{d} X_{1}^{(d)}\right)}\right\} \\
& =\mathbb{E}\left\{e^{i\left(a_{1} X_{2}^{(1)}+\cdots+a_{d} X_{2}^{(d)}\right)}\right\} \quad \text { (by assumption) } \\
& =\psi_{2}\left(a_{1}, \ldots, a_{d}\right),
\end{aligned}
$$

the characteristic function of $X_{2} . \square$

Thus, we have proved the following:
Theorem 4.3. $L \leq 1 / 2$ with equality if and only if $L^{*}=1 / 2$.
Thus, as in the one-dimensional case, whenever $L^{*}<1 / 2$, a meaningful ( $L<$ $1 / 2$ ) cut by a hyperplane is possible. There are also examples in which no cut improves over a rule in which $g(x) \equiv y$ for some $y$ and all $x$, yet $L^{*}=0$ and $L>1 / 4$ (say). To generalize Theorem 4.1, we offer the following result. A related inequality is shown in Problem 4.7. The idea of using Chebyshev's inequality to obtain such bounds is due to Yau and Lin (1968) (see also Devijver and Kittler (1982, p.162)).

Theorem 4.4. Let $X_{0}$ and $X_{1}$ be random variables distributed as $X$ given $Y=0$, and $Y=1$ respectively. Set $m_{0}=\mathbf{E}\left\{X_{0}\right\}, m_{1}=\mathbf{E}\left\{X_{1}\right\}$. Define also the covariance matrices $\Sigma_{1}=\mathbf{E}\left\{\left(X_{1}-m_{1}\right)\left(X_{1}-m_{1}\right)^{T}\right\}$ and $\Sigma_{0}=\mathbf{E}\left\{\left(X_{0}-m_{0}\right)\left(X_{0}-m_{0}\right)^{T}\right\}$. Then

$$
L^{*} \leq L \leq \inf _{a \in \mathcal{R}^{d}} \frac{1}{1+\frac{\left(a^{T}\left(m_{1}-m_{0}\right)\right)^{2}}{\left(\left(a^{T} \Sigma_{0} a\right)^{1 / 2}+\left(a^{T} \Sigma_{1} a\right)^{1 / 2}\right)^{2}}}
$$

Proof. For any $a \in \mathcal{R}^{d}$ we may apply Theorem 4.1 to $a^{T} X_{0}$ and $a^{T} X_{1}$. Theorem 4.4 follows by noting that

$$
\begin{aligned}
& \mathbf{E}\left\{a^{T} X_{0}\right\}=a^{T} \mathbf{E}\left\{X_{0}\right\}=a^{T} m_{0}, \\
& \mathbf{E}\left\{a^{T} X_{1}\right\}=a^{T} m_{1},
\end{aligned}
$$

and that

$$
\begin{aligned}
& \operatorname{Var}\left\{a^{T} X_{0}\right\}=\mathbf{E}\left\{a^{T}\left(X_{0}-m_{0}\right)\left(X_{0}-m_{0}\right)^{T} a\right\}=a^{T} \Sigma_{0} a, \\
& \operatorname{Var}\left\{a^{T} X_{1}\right\}=a^{T} \Sigma_{1} a .
\end{aligned}
$$

We may obtain explicit inequalities by different choices of $a . a=m_{1}-m_{0}$ yields a convenient formula. We see from the next section that $a=\Sigma\left(m_{1}-m_{0}\right)$ with $\Sigma=p \Sigma_{1}+(1-p) \Sigma_{0}$ is also a meaningful choice (see also Problem 4.7).

### 4.3 The Fisher Linear Discriminant

Data-based values for $a$ may be found by various criteria. One of the first methods was suggested by Fisher (1936). Let $\widehat{m}_{1}$ and $\widehat{m}_{0}$ be the sample means for the two classes (e.g., $\widehat{m}_{1}=\sum_{i: Y_{i}=1} X_{i} /\left|\left\{i: Y_{i}=1\right\}\right|$.) Picture projecting $X_{1}, \ldots, X_{n}$ to a line in the direction of $a$. Note that this is perpendicular to the hyperplane given by $a^{T} x+a_{0}=0$. The projected values are $a^{T} X_{1}, \ldots, a^{T} X_{n}$. These are all equal to 0 for those $X_{i}$ on the hyperplane $a^{T} x=0$ through the origin, and grow in absolute value as we flee that hyperplane. Let $\widehat{\sigma}_{1}^{2}$ and $\widehat{\sigma}_{0}^{2}$ be the sample scatters for classes 1 and 0 , respectively, that is,

$$
\widehat{\sigma}_{1}^{2}=\sum_{i: Y_{i}=1}\left(a^{T} X_{i}-a^{T} \widehat{m}_{1}\right)^{2}=a^{T} S_{1} a
$$

and similarly for $\widehat{\sigma}_{0}^{2}$, where

$$
S_{1}=\sum_{i: Y_{i}=1}\left(X_{i}-\widehat{m}_{1}\right)\left(X_{i}-\widehat{m}_{1}\right)^{T}
$$

is the scatter matrix for class 1 .

The Fisher linear discriminant is that linear function $a^{T} x$ for which the criterion

$$
J(a)=\frac{\left(a^{T} \widehat{m}_{1}-a^{T} \widehat{m}_{0}\right)^{2}}{\widehat{\sigma}_{1}^{2}+\widehat{\sigma}_{0}^{2}}=\frac{\left(a^{T}\left(\widehat{m}_{1}-\widehat{m}_{0}\right)\right)^{2}}{a^{T}\left(S_{1}+S_{0}\right) a}
$$

is maximum. This corresponds to finding a direction $a$ that best separates $a^{T} \widehat{m}_{1}$ from $a^{T} \widehat{m}_{0}$ relative to the sample scatter. Luckily, to find that $a$, we need not resort to numerical iteration-the solution is given by

$$
a=\left(S_{1}+S_{0}\right)^{-1}\left(\widehat{m}_{1}-\widehat{m}_{0}\right) .
$$

Fisher's suggestion is to replace $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ by $\left(a^{T} X_{1}, Y_{1}\right), \ldots$, ( $a^{T} X_{n}, Y_{n}$ ) and to perform one-dimensional discrimination. Usually, the rule uses a simple split

$$
g_{a_{0}}(x)= \begin{cases}1 & \text { if } a^{T} x+a_{0}>0  \tag{4.1}\\ 0 & \text { otherwise }\end{cases}
$$

for some constant $a_{0}$. Unfortunately, Fisher discriminants can be arbitrarily bad: there are distributions such that even though the two classes are linearly separable (i.e., $L=0$ ), the Fisher linear discriminant has an error probability close to 1 (see Problem 4.9).

### 4.4 The Normal Distribution

There are a few situations in which, by sheer accident, the Bayes rule is a linear discriminant. While this is not a major issue, it is interesting to identify the most important case, i.e., that of the multivariate normal distribution. The general multivariate normal density is written as

$$
f(x)=\frac{1}{\sqrt{(2 \pi)^{d} \operatorname{det}(\Sigma)}} e^{-\frac{1}{2}(x-m)^{T} \Sigma^{-1}(x-m)}
$$

where $m$ is the mean (both $x$ and $m$ are $d$-component column vectors), $\Sigma$ is the $d \times d$ covariance matrix, $\Sigma^{-1}$ is the inverse of $\Sigma$, and $\operatorname{det}(\Sigma)$ is its determinant. We write $f \sim N(m, \Sigma)$. Clearly, if $X$ has density $f$, then $m=\mathbf{E} X$ and $\Sigma=$ $\mathbf{E}\left\{(X-m)(X-m)^{T}\right\}$.

The multivariate normal density is completely specified by $d+\binom{d}{2}$ formal parameters ( $m$ and $\Sigma$ ). A sample from the density is clustered in an elliptical cloud. The loci of points of constant density are ellipsoids described by

$$
(x-m)^{T} \Sigma^{-1}(x-m)=r^{2}
$$

for some constant $r \geq 0$. The number $r$ is the Mahalanobis distance from $x$ to $m$, and is in fact useful even when the underlying distribution is not normal. It takes into account the directional stretch of the space determined by $\Sigma$.


Figure 4.4. Points at equal Mahalanobis distance from $m$.

Given a two-class problem in which $X$ has a density $(1-p) f_{0}(x)+p f_{1}(x)$ and $f_{0}$ and $f_{1}$ are both multivariate normal with parameters $m_{i}, \Sigma_{i}, i=0,1$, the Bayes rule is easily described by

$$
g^{*}(x)= \begin{cases}1 & \text { if } p f_{1}(x)>(1-p) f_{0}(x) \\ 0 & \text { otherwise }\end{cases}
$$

Take logarithms and note that $g^{*}(x)=1$ if and only if

$$
\begin{aligned}
& \left(x-m_{1}\right)^{T} \Sigma_{1}^{-1}\left(x-m_{1}\right)-2 \log p+\log \left(\operatorname{det}\left(\Sigma_{1}\right)\right) \\
& \quad<\quad\left(x-m_{0}\right)^{T} \Sigma_{0}^{-1}\left(x-m_{0}\right)-2 \log (1-p)+\log \left(\operatorname{det}\left(\Sigma_{0}\right)\right)
\end{aligned}
$$

In practice, one might wish to estimate $m_{1}, m_{0}, \Sigma_{1}, \Sigma_{0}$ and $p$ from the data and use these estimates in the formula for $g^{*}$. Interestingly, as $\left(x-m_{i}\right)^{T} \Sigma_{i}^{-1}\left(x-m_{i}\right)$ is the squared Mahalanobis distance from $x$ to $m_{i}$ in class $i$ (called $r_{i}^{2}$ ), the Bayes rule is simply

$$
g^{*}(x)= \begin{cases}1 & \text { if } r_{1}^{2}<r_{0}^{2}-2 \log ((1-p) / p)+\log \left(\operatorname{det}\left(\Sigma_{0}\right) / \operatorname{det}\left(\Sigma_{1}\right)\right) \\ 0 & \text { otherwise }\end{cases}
$$

In particular, when $p=1 / 2, \Sigma_{0}=\Sigma_{1}=\Sigma$, we have

$$
g^{*}(x)= \begin{cases}1 & \text { if } r_{1}^{2}<r_{0}^{2} \\ 0 & \text { otherwise }\end{cases}
$$

just classify according to the class whose mean is at the nearest Mahalanobis distance from $x$. When $\Sigma_{0}=\Sigma_{1}=\Sigma$, the Bayes rule becomes linear:

$$
g^{*}(x)= \begin{cases}1 & \text { if } a^{T} x+a_{0}>0 \\ 0 & \text { otherwise }\end{cases}
$$

where $a=\left(m_{1}-m_{0}\right) \Sigma^{-1}$, and $a_{0}=2 \log (p /(1-p))+m_{0}^{T} \Sigma^{-1} m_{0}-m_{1}^{T} \Sigma^{-1} m_{1}$. Thus, linear discrimination rules occur as special cases of Bayes rules for multivariate normal distributions.

Our intuition that $a$ should be in the direction $m_{1}-m_{0}$ to best separate the classes is almost right. Note nevertheless that $a$ is not perpendicular in general to the hyperplane of loci at equal distance from $m_{0}$ and $m_{1}$. When $\Sigma$ is replaced by
the standard data-based estimate, we obtain in fact the Fisher linear discriminant. Furthermore, when $\Sigma_{1} \neq \Sigma_{0}$, the decision boundary is usually not linear, and Fisher's linear discriminant must therefore be suboptimal.

Historical remarks. In early statistical work on discrimination, the normal distribution plays a central role (Anderson (1958)). For a simple introduction, we refer to Duda and Hart (1973). McLachlan (1992) has more details, and Raudys (1972; 1976) relates the error, dimensionality, and sample size for normal and nearly normal models. See also Raudys and Pikelis (1980; 1982).

### 4.5 Empirical Risk Minimization

In this section we present an algorithm that yields a classifier whose error probability is very close to the minimal error probability $L$ achievable by linear classifiers, provided that $X$ has a density. The algorithm selects a classifier by minimizing the empirical error over finitely many- $2\binom{n}{d}$-linear classifiers. For a rule

$$
\phi(x)= \begin{cases}1 & \text { if } a^{T} x+a_{0}>0 \\ 0 & \text { otherwise }\end{cases}
$$

the probability of error is

$$
L(\phi)=\mathbf{P}\{\phi(X) \neq Y\}
$$

$L(\phi)$ may be estimated by the empirical risk

$$
\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi\left(X_{i}\right) \neq Y_{i}\right\}}
$$

that is, the number of errors made by the classifier $\phi$ is counted and normalized.
Assume that $X$ has a density, and consider $d$ arbitrary data points $X_{i_{1}}, X_{i_{2}}, \ldots$, $X_{i_{d}}$ among $\left\{X_{1}, \ldots, X_{n}\right\}$, and let $a^{T} x+a_{0}=0$ be a hyperplane containing these points. Because of the density assumption, the $d$ points are in general position with probability one, and this hyperplane is unique. This hyperplane determines two classifiers:

$$
\phi_{1}(x)= \begin{cases}1 & \text { if } a^{T} x+a_{0}>0 \\ 0 & \text { otherwise }\end{cases}
$$

and

$$
\phi_{2}(x)= \begin{cases}1 & \text { if } a^{T} x+a_{0}<0 \\ 0 & \text { otherwise }\end{cases}
$$

whose empirical errors $\widehat{L}_{n}\left(\phi_{1}\right)$ and $\widehat{L}_{n}\left(\phi_{2}\right)$ may be calculated. To each $d$-tuple $X_{i_{1}}, X_{i_{2}}, \ldots, X_{i_{d}}$ of data points, we may assign two classifiers in this manner, yielding altogether $2\binom{n}{d}$ classifiers. Denote these classifiers by $\phi_{1}, \ldots, \phi_{2\binom{n}{d} \text {. Let }}$ $\widehat{\phi}$ be a linear classifier that minimizes $\widehat{L}_{n}\left(\phi_{i}\right)$ over all $i=1, \ldots, 2\binom{n}{d}$.

We denote the best possible error probability by

$$
L=\inf _{\phi} L(\phi)
$$

over the class of all linear rules, and define $\phi^{*}=\arg \min _{\phi} L(\phi)$ as the best linear rule. If there are several classifiers with $L(\phi)=L$, then we choose $\phi^{*}$ among these in an arbitrary fixed manner. Next we show that the classifier corresponding to $\widehat{\phi}$ is really very good.


Figure 4.5. If the data points are in general position, then for each linear rule there exists a linear split defined by a hyperplane crossing $d$ points such that the difference between the empirical errors is at most $d / n$.

First note that there is no linear classifier $\phi$ whose empirical error $\widehat{L}_{n}(\phi)$ is smaller than $\widehat{L}(\widehat{\phi})-d / n$. This follows from the fact that since the data points are in general position (recall the density assumption), then for each linear classifier we may find one whose defining hyperplane contains exactly $d$ data points such that the two decisions agree on all data points except possibly for these $d$ pointssee Figure 4.5. Thus, we may view minimization of the empirical error over the finite set $\left\{\phi_{1}, \ldots, \phi_{\left.2\left({ }_{d}^{( }\right)\right\}}\right)$as approximate minimization over the infinite set of linear classifiers. In Chapters 12 and 13 we will develop the full theory for rules that are found by empirical risk minimization. Theorem 4.5 just gives you a taste of things to come. Other-more involved, but also more general-proofs go back to Vapnik and Chervonenkis (1971; 1974c).

Theorem 4.5. Assume that $X$ has a density. If $\widehat{\phi}$ is found by empirical error minimization as described above, then, for all possible distributions of $(X, Y)$, ifn $\geq d$ and $2 d / n \leq \epsilon \leq 1$, we have

$$
\mathbf{P}\{L(\widehat{\phi})>L+\epsilon\} \leq e^{2 d \epsilon}\left(2\binom{n}{d}+1\right) e^{-n \epsilon^{2} / 2} .
$$

Moreover, if $n \geq d$, then

$$
\mathbf{E}\{L(\widehat{\phi})-L\} \leq \sqrt{\frac{2}{n}((d+1) \log n+(2 d+2))} .
$$

Remark. With some care Theorem 4.5 and Theorem 4.6 below can be extended so that the density assumption may be dropped. One needs to ensure that the selected
linear rule has empirical error close to that of the best possible linear rule. With the classifier suggested above this property may fail to hold if the data points are not necessarily of general position. The ideas presented here are generalized in Chapter 12 (see Theorem 12.2).

Proof. We begin with the following simple inequality:

$$
\begin{aligned}
L(\widehat{\phi})-L= & L(\widehat{\phi})-\widehat{L}_{n}(\widehat{\phi})+\widehat{L}_{n}(\widehat{\phi})-L\left(\phi^{*}\right) \\
\leq & L(\widehat{\phi})-\widehat{L}_{n}(\widehat{\phi})+\widehat{L}_{n}\left(\phi^{*}\right)-L\left(\phi^{*}\right)+\frac{d}{n} \\
& \left(\text { since } \widehat{L}_{n}(\widehat{\phi}) \leq \widehat{L}_{n}(\phi)+d / n \text { for any } \phi\right) \\
\leq & \max _{i=1, \ldots, 2\binom{n}{d}}\left(L\left(\phi_{i}\right)-\widehat{L}_{n}\left(\phi_{i}\right)\right)+\widehat{L}_{n}\left(\phi^{*}\right)-L\left(\phi^{*}\right)+\frac{d}{n} .
\end{aligned}
$$

Therefore, by the union-of-events bound, we have

$$
\begin{aligned}
& \mathbf{P}\{L(\widehat{\phi})-L>\epsilon\} \\
& \quad \leq \sum_{i=1}^{2\binom{n}{d}} \mathbf{P}\left\{L\left(\phi_{i}\right)-\widehat{L}_{n}\left(\phi_{i}\right)>\frac{\epsilon}{2}\right\}+\mathbf{P}\left\{\widehat{L}_{n}\left(\phi^{*}\right)-L\left(\phi^{*}\right)+\frac{d}{n}>\frac{\epsilon}{2}\right\} .
\end{aligned}
$$

To bound the second term on the right-hand side, observe that $n \widehat{L}_{n}\left(\phi^{*}\right)$ is binomially distributed with parameters $n$ and $L\left(\phi^{*}\right)$. By an inequality due to Chernoff (1952) and Okamoto (1958) for the tail of the binomial distribution,

$$
\mathbf{P}\left\{\widehat{L}_{n}\left(\phi^{*}\right)-L\left(\phi^{*}\right)>\frac{\epsilon}{2}-\frac{d}{n}\right\} \leq e^{-2 n\left(\frac{\epsilon}{2}-\frac{d}{n}\right)^{2}} \leq e^{2 d \epsilon} \cdot e^{-n \epsilon^{2} / 2} .
$$

We prove this inequality later (see Theorem 8.1). Next we bound one term of the sum on the right-hand side. Note that by symmetry all $2\binom{n}{d}$ terms are equal. Assume that the classifier $\phi_{1}$ is determined by the $d$-tuple of the first $d$ data points $X_{1}, \ldots, X_{d}$. We write

$$
\mathbf{P}\left\{L\left(\phi_{1}\right)-\widehat{L}_{n}\left(\phi_{1}\right)>\frac{\epsilon}{2}\right\}=\mathbf{E}\left\{\mathbf{P}\left\{\left.L\left(\phi_{1}\right)-\widehat{L}_{n}\left(\phi_{1}\right)>\frac{\epsilon}{2} \right\rvert\, X_{1}, \ldots, X_{d}\right\}\right\},
$$

and bound the conditional probability inside. Let $\left(X_{1}^{\prime \prime}, Y_{1}^{\prime \prime}\right), \ldots,\left(X_{d}^{\prime \prime}, Y_{d}^{\prime \prime}\right)$ be independent of the data and be distributed as the data $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{d}, Y_{d}\right)$. Define

$$
\left(X_{i}^{\prime}, Y_{i}^{\prime}\right)=\left\{\begin{array}{cc}
\left(X_{i}^{\prime \prime}, Y_{i}^{\prime \prime}\right) & \text { if } i \leq d \\
\left(X_{i}, Y_{i}\right) & \text { if } i>d .
\end{array}\right.
$$

Then

$$
\begin{aligned}
\mathbf{P} & \left\{\left.L\left(\phi_{1}\right)-\widehat{L}_{n}\left(\phi_{1}\right)>\frac{\epsilon}{2} \right\rvert\, X_{1}, \ldots, X_{d}\right\} \\
& \leq \mathbf{P}\left\{\left.L\left(\phi_{1}\right)-\frac{1}{n} \sum_{i=d+1}^{n} I_{\left\{\phi_{1}\left(X_{i}\right) \neq Y_{i}\right\}}>\frac{\epsilon}{2} \right\rvert\, X_{1}, \ldots, X_{d}\right\}
\end{aligned}
$$

$$
\begin{aligned}
\leq & \mathbf{P}\left\{\left.L\left(\phi_{1}\right)-\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi_{1}\left(X_{i}^{\prime}\right) \neq Y_{i}^{\prime}\right\}}+\frac{d}{n}>\frac{\epsilon}{2} \right\rvert\, X_{1}, \ldots, X_{d}\right\} \\
= & \mathbf{P}\left\{\left.L\left(\phi_{1}\right)-\frac{1}{n} \operatorname{Binomial}\left(n, L\left(\phi_{1}\right)\right)>\frac{\epsilon}{2}-\frac{d}{n} \right\rvert\, X_{1}, \ldots, X_{d}\right\} \\
& \left(\text { as } L\left(\phi_{1}\right) \text { depends upon } X_{1}, \ldots, X_{d}\right. \text { only and } \\
& \left.\left(X_{1}^{\prime}, Y_{1}^{\prime}\right) \ldots,\left(X_{d}^{\prime}, Y_{d}^{\prime}\right) \text { are independent of } X_{1}, \ldots, X_{d}\right) \\
\leq & e^{-2 n\left(\frac{\epsilon}{2}-\frac{d}{n}\right)^{2}}
\end{aligned}
$$

$$
\text { (by Theorem } 8.1 \text {; use the fact that } \epsilon \geq 2 d / n \text { ) }
$$

$$
\leq e^{-n \epsilon^{2} / 2} e^{2 d \epsilon}
$$

The inequality for the expected value follows from the probability inequality by the following simple argument: by the Cauchy-Schwarz inequality,

$$
(\mathbb{E}\{L(\widehat{\phi})-L\})^{2} \leq \mathbf{E}\left\{(L(\widehat{\phi})-L)^{2}\right\}
$$

Denoting $Z=(L(\widehat{\phi})-L)^{2}$, for any $u>0$,

$$
\begin{aligned}
\mathbf{E}\{Z\} & =\mathbf{E}\{Z \mid Z>u\} \mathbf{P}\{Z>u\}+\mathbf{E}\{Z \mid Z \leq u\} \mathbf{P}\{Z \leq u\} \\
& \leq \mathbf{P}\{Z>u\}+u \\
& \leq e^{2 d}\left(2 n^{d}+1\right) e^{-n u / 2}+u \quad \text { if } u \geq\left(\frac{2 d}{n}\right)^{2}
\end{aligned}
$$

by the probability inequality, and since $\binom{n}{d} \leq n^{d}$. Choosing $u$ to minimize the obtained expression yields the desired inequality: first verify that the minimum occurs for

$$
u=\frac{2}{n} \log \frac{n c}{2}
$$

where $c=e^{2 d}\left(2 n^{d}+1\right)$. Check that if $n \geq d$, then $u \geq(2 d / n)^{2}$. Then note that the bound $c e^{-n u / 2}+u$ equals

$$
\frac{2}{n} \log \frac{n e c}{2} \leq \frac{2}{n} \log \left(e^{2 d+2} n^{d+1}\right)=\frac{2}{n}((d+1) \log n+(2 d+2))
$$

Observe for now that the bound on $\mathbf{P}\{L(\widehat{\phi})>L+\epsilon\}$ decreases rapidly with $n$. To have an impact, it must become less than $\delta$ for small $\delta$. This happens, roughly speaking, when

$$
n \geq c \cdot \frac{d}{\epsilon^{2}}\left(\log \frac{d}{\epsilon^{2}}+\log \frac{1}{\delta}\right)
$$

for some constant $c$. Doubling $d$, the dimension, causes this minimal sample size to roughly double as well.

An important special case is when the distribution is linearly separable, that is, $L=0$. In such cases the empirical risk minimization above performs even better as the size of the error improves to $O(d \log n / n)$ from $O(\sqrt{d \log n / n})$. Clearly, the data points are linearly separable as well, that is, $\widehat{L}_{n}\left(\phi^{*}\right)=0$ with probability one, and therefore $\widehat{L}_{n}(\widehat{\phi}) \leq d / n$ with probability one.

Theorem 4.6. Assume that $X$ has a density, and that the best linear classifier has zero probability of error $(L=0)$. Then for the empirical risk minimization algorithm of Theorem 4.5, for all $n>d$ and $\epsilon \leq 1$,

$$
\mathbf{P}\{L(\widehat{\phi})>\epsilon\} \leq 2\binom{n}{d} e^{-(n-d) \epsilon},
$$

and

$$
\mathbf{E}\{L(\widehat{\phi})\} \leq \frac{d \log n+2}{n-d}
$$

Proof. By the union bound,

$$
\begin{aligned}
\mathbf{P}\{L(\widehat{\phi})>\epsilon\} & \leq \mathbf{P}\left\{\max _{i=1,2, \ldots, 2\binom{n}{d}: \hat{L}_{n}\left(\phi_{i}\right) \leq \frac{d}{n}} L\left(\phi_{i}\right)>\epsilon\right\} \\
& \leq \sum_{i=1}^{2\binom{n}{d}} \mathbf{P}\left\{\widehat{L}_{n}\left(\phi_{i}\right) \leq \frac{d}{n}, L\left(\phi_{i}\right)>\epsilon\right\} .
\end{aligned}
$$

By symmetry, this sum equals

$$
\begin{aligned}
& 2\binom{n}{d} \mathbf{P}\left\{\widehat{L}_{n}\left(\phi_{1}\right) \leq \frac{d}{n}, L\left(\phi_{1}\right)>\epsilon\right\} \\
& \quad=2\binom{n}{d} \mathbf{E}\left\{\mathbf{P}\left\{\widehat{L}_{n}\left(\phi_{1}\right) \leq \frac{d}{n}, L\left(\phi_{1}\right)>\epsilon \mid X_{1}, \ldots, X_{d}\right\}\right\}
\end{aligned}
$$

where, as in Theorem 4.5, $\phi_{1}$ is determined by the data points $X_{1}, \ldots, X_{d}$. However,

$$
\left.\begin{array}{rl}
\mathbf{P}\{ & \left\{\widehat{L}_{n}\left(\phi_{1}\right) \leq \frac{d}{n}, L\left(\phi_{1}\right)>\epsilon \mid X_{1}, \ldots, X_{d}\right\} \\
\leq & \mathbf{P}\left\{\phi_{1}\left(X_{d+1}\right)=Y_{d+1}, \ldots, \phi_{1}\left(X_{n}\right)=Y_{n}, L\left(\phi_{1}\right)>\epsilon \mid X_{1}, \ldots, X_{d}\right\} \\
& \left(\text { since all of the }(\text { at most } d) \text { errors committed by } \phi_{1}\right.
\end{array}\right\} \begin{aligned}
& \text { occur for } \left.\left(X_{1}, Y_{1}\right), \ldots,\left(X_{d}, Y_{d}\right)\right) \\
\leq & (1-\epsilon)^{n-d},
\end{aligned}
$$

since the probability that no $\left(X_{i}, Y_{i}\right)$, pair $i=d+1, \ldots, n$ falls in the set $\{(x, y)$ : $\left.\phi_{1}(x) \neq y\right\}$ is less than $(1-\epsilon)^{n-d}$ if the probability of the set is larger than $\epsilon$. The proof of the probability inequality may be completed by noting that $1-x \leq e^{-x}$.

For the expected error probability, note that for any $u>0$,

$$
\begin{aligned}
\mathbf{E}\{L(\widehat{\phi})\} & =\int_{0}^{\infty} \mathbf{P}\{L(\widehat{\phi})>t\} d t \\
& \leq u+\int_{u}^{\infty} \mathbf{P}\{L(\widehat{\phi})>t\} d t \\
& \leq u+2 n^{d} \int_{u}^{\infty} e^{-(n-d) t} d t
\end{aligned}
$$

(by the probability inequality and $\binom{n}{d} \leq n^{d}$ )
$=u+\frac{2 n^{d}}{n} e^{-(n-d) u}$.
We choose $u$ to minimize the obtained bound, which yields the desired inequality.

### 4.6 Minimizing Other Criteria

Empirical risk minimization uses extensive computations, because $\widehat{L}_{n}(\phi)$ is not a unimodal function in general (see Problems 4.10 and 4.11). Also, gradient optimization is difficult because the gradients are zero almost everywhere. In fact, given $n$ labeled points in $\mathcal{R}^{d}$, finding the best linear dichotomy is np hard (see Johnson and Preparata (1978)). To aid in the optimization, some have suggested minimizing a modified empirical error, such as

$$
\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} \Psi\left(\left(2 Y_{i}-1\right)-a^{T} X_{i}-a_{0}\right) I_{\left\{Y_{i} \neq g_{a}\left(X_{i}\right)\right\}}
$$

or

$$
\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} \Psi\left(\left(2 Y_{i}-1\right)-a^{T} X_{i}-a_{0}\right),
$$

where $\Psi$ is a positive convex function. Of particular importance here is the mean square error criterion $\Psi(u)=u^{2}$ (see, e.g., Widrow and Hoff (1960)). One can easily verify that $\widehat{L}_{n}(\phi)$ has a gradient (with respect to $\left(a, a_{0}\right)$ ) that may aid in locating a local minimum. Let $\widehat{\phi}$ denote the linear discrimination rule minimizing

$$
\mathbf{E}\left\{\left((2 Y-1)-a^{T} X-a_{0}\right)^{2}\right\}
$$

over all $a$ and $a_{0}$. A description of the solution is given in Problem 4.14.
Even in a one-dimensional situation, the mean square error criterion muddles the issue and does not give any performance guarantees:

Theorem 4.7. If $\sup _{(X, Y)}$ denotes the supremum with respect to all distributions on $\mathcal{R} \times\{0,1\}$, then

$$
\sup _{(X, Y)}(L(\widehat{\phi})-L)=1
$$

where $\widehat{\phi}$ is a linear discriminant obtained by minimizing

$$
\mathbf{E}\left\{\left((2 Y-1)-a_{1} X-a_{0}\right)^{2}\right\}
$$

over all $a_{1}$ and $a_{0}$.

REMARK. This theorem establishes the existence of distributions of $(X, Y)$ for which $L(\widehat{\phi})>1-\epsilon$ and $L<\epsilon$ simultaneously for arbitrarily small $\epsilon>0$. Therefore, minimizing the mean square error criterion is not recommended unless one has additional information regarding the distribution of $(X, Y)$.

Proof. Let $\epsilon>0$ and $\theta>0$. Consider a triatomic distribution of $(X, Y)$ :

$$
\mathbf{P}\{(X, Y)=(-\theta, 1)\}=\mathbf{P}\{(X, Y)=(1,1)\}=\epsilon / 2
$$

$$
\mathbf{P}\{(X, Y)=(0,0)\}=1-\epsilon .
$$



For $\epsilon<1 / 2$, the best linear rule decides class 0 on $[-\theta / 2, \infty)$ and 1 elsewhere, for a probability of error $L=\epsilon / 2$. The mean square error criterion asks that we minimize

$$
L(\widehat{\phi})=\left\{(1-\epsilon)(-1-v)^{2}+\frac{\epsilon}{2}(1-u-v)^{2}+\frac{\epsilon}{2}(1+u \theta-v)^{2}\right\}
$$

with respect to $a_{0}=v$ and $a_{1}=u$. Setting the derivatives with respect to $u$ and $v$ equal to zero yields

$$
u=\frac{(v-1) \theta-v}{1+\theta^{2}}, \quad \text { and } \quad v=2 \epsilon-1+\frac{\epsilon}{2} u(\theta-1)
$$

for

$$
v=\frac{(2 \epsilon-1)\left(1+\theta^{2}\right)-\frac{\epsilon}{2} \theta(\theta-1)}{1+\theta^{2}-\frac{\epsilon}{2}(1-\theta)^{2}}
$$

If we let $\epsilon \downarrow 0$ and let $\theta \uparrow \infty$, then $v \sim 3 \epsilon / 2$. Thus, for $\epsilon$ small enough and $\theta$ large enough, considering the decision at 0 only, $L(\widehat{\phi}) \geq 1-\epsilon$, because at $x=0$, $u x+v=v>0$. Thus, $L(\widehat{\phi})-L \geq 1-3 \epsilon / 2$ for $\epsilon$ small enough and $\theta$ large enough.

Others have suggested minimizing

$$
\sum_{i=1}^{n}\left(\sigma\left(a^{T} X_{i}+a_{0}\right)-Y_{i}\right)^{2}
$$

where $\sigma(u)$ is a sigmoid, that is, an increasing function from 0 to 1 such as $1 /(1+$ $e^{-u}$ ), see, for example, Wassel and Sklansky (1972), Do Tu and Installe (1975), Fritz and Györfi (1976), and Sklansky and Wassel (1979). Clearly, $\sigma(u)=I_{\{u \geq 0\}}$ provides the empirical error probability. However, the point here is to use smooth sigmoids so that gradient algorithms may be used to find the optimum. This may be viewed as a compromise between the mean squared error criteria and empirical error minimization. Here, too, anomalies can occur, and the error space is not well behaved, displaying many local minima (Hertz, Krogh, and Palmer (1991, p.108)). See, however, Problems 4.16 and 4.17.

## Problems and Exercises

Problem 4.1. With the notation of Theorem 4.1, show that the error probability $L$ of a one-dimensional theoretical Stoller split satisfies

$$
L \leq \frac{4 p(1-p)}{1+p(1-p) \frac{\left(m_{0}-m_{1}\right)^{2}}{(1-p) \sigma_{0}^{2}+p \sigma_{1}^{2}}}
$$

(Györfi and Vajda (1980)). Is this bound better than that of Theorem 4.1? Hint: For any threshold rule $g_{c}(x)=I_{\{x \geq c\}}$ and $u>0$, write

$$
\begin{aligned}
L\left(g_{c}\right) & =\mathbf{P}\{X-c \geq 0,2 Y-1=-1\}+\mathbf{P}\{X-c<0,2 Y-1=1\} \\
& \leq \mathbf{P}\{|u(X-c)-(2 Y-1)| \geq 1\} \\
& \leq \mathbf{E}\left\{(u(X-c)-(2 Y-1))^{2}\right\}
\end{aligned}
$$

by Chebyshev's inequality. Choose $u$ and $c$ to minimize the upper bound.
Problem 4.2. Let $p=1 / 2$. If $L$ is the error probability of the one-dimensional theoretical Stoller split, show that

$$
L \leq \frac{1}{2+2 \frac{\left(m_{0}-m_{1}\right)^{2}}{\left(\sigma_{0}+\sigma_{1}\right)^{2}}} .
$$

Show that the bound is achieved for some distribution when the class-conditional distributions of $X$ (that is, given $Y=0$ and $Y=1$ ) are concentrated on two points each, one of which is shared by both classes (Chernoff (1971), Becker (1968)).

Problem 4.3. Let $X$ be a univariate random variable. The distribution functions for $X$ given $Y=1$ and $Y=0$ are $F_{1}$ and $F_{0}$ respectively. Assume that the moment generating functions for $X$ exist, that is, $\mathbf{E}\left\{e^{t X} \mid Y=1\right\}=\psi_{1}(t), \mathbf{E}\left\{e^{t X} \mid Y=0\right\}=\psi_{0}(t), t \in \mathcal{R}$, where $\psi_{1}, \psi_{0}$ are finite for all $t$. In the spirit of Theorem 4.1, derive an upper bound for $L$ in function of $\psi_{1}, \psi_{0}$. Apply your bound to the case that $F_{1}$ and $F_{0}$ are both normal with possibly different means and variances.

Problem 4.4. Signals in additive gaussian noise. Let $s_{0}, s_{1} \in \mathcal{R}^{d}$ be fixed, and let $N$ be a multivariate gaussian random variable with zero mean and covariance matrix $\Sigma$. Let $\mathbf{P}\{Y=0\}=\mathbf{P}\{Y=1\}=1 / 2$, and define

$$
X= \begin{cases}s_{0}+N & \text { if } Y=0 \\ s_{1}+N & \text { if } Y=1\end{cases}
$$

Construct the Bayes decision and calculate $L^{*}$. Prove that if $\Sigma$ is the identity matrix, and $s_{0}$ and $s_{1}$ have constant components, then $L^{*} \rightarrow 0$ exponentially rapidly as $d \rightarrow \infty$.

Problem 4.5. In the last step of the proof of Theorem 4.2, we used the Dvoretzky-Kiefer-Wolfowitz-Massart inequality (Theorem 12.9). This result states that if $Z_{1}, \ldots, Z_{n}$ are i.i.d. random variables on the real line with distribution function $F(z)=\mathbf{P}\left\{Z_{1} \leq z\right\}$ and empirical distribution function $F_{n}(z)=(1 / n) \sum_{i=1}^{n} I_{\left\{Z_{i} \leq z\right\}}$, then

$$
\mathbf{P}\left\{\sup _{z \in \mathcal{R}}\left|F(z)-F_{n}(z)\right| \geq \epsilon\right\} \leq 2 e^{-2 n \epsilon^{2}}
$$

Use this inequality to conclude that

$$
\mathbf{P}\left\{\sup _{x}\left|\nu(C(x, 1))-v_{n}(C(x, 1))\right| \geq \frac{\epsilon}{2}\right\} \leq 2 e^{-2 n(\epsilon / 2)^{2}} .
$$

Hint: Map $(X, Y)$ on the real line by a one-to-one function $\psi:(\mathcal{R} \times\{0,1\}) \rightarrow \mathcal{R}$ such that $Z=\psi((X, Y))<0$ if and only if $Y=0$. Use the Dvoretzky-Kiefer-Wolfowitz-Massart inequality for $Z$.

Problem 4.6. Let $L$ be the probability of error for the best sphere rule, that is, for the rule that associates a class with the inside of a sphere $S_{x, r}$, and the other class with the outside. Here the center $x$, and radius $r$ are both variable. Show that $L=1 / 2$ if and only if $L^{*}=1 / 2$, and that $L \leq 1 / 2$.

Problem 4.7. With the notation of Theorem 4.4, show that the probability of error $L$ of the best linear discriminant satisfies

$$
L \leq \frac{4 p(1-p)}{1+p(1-p) \Delta^{2}}
$$

where

$$
\Delta=\sqrt{\left(m_{1}-m_{0}\right)^{T} \Sigma^{-1}\left(m_{1}-m_{0}\right)}
$$

is the Mahalanobis distance (Chapter 3) with $\Sigma=p \Sigma_{1}+(1-p) \Sigma_{0}$ (Györfi and Vajda(1980)). Interestingly, the upper bound is just twice the bound of Theorem 3.4 for the asymptotic nearest neighbor error. Thus, a large Mahalanobis distance does not only imply that the Bayes error is small, but also, small error probabilities may be achieved by simple linear classifiers. Hint: Apply the inequality of Problem 4.1 for the univariate random variable $X^{\prime}=a^{T} X a=\Sigma^{-1}\left(m_{1}-m_{0}\right)$.

Problem 4.8. If $m_{i}$ and $\sigma_{i}^{2}$ are the mean and variance of $a^{T} X$, given that $Y=i, i=0,1$, where $a$ is a column vector of weights, then show that the criterion

$$
J_{1}(a)=\frac{\left(m_{1}-m_{0}\right)^{2}}{\sigma_{1}^{2}+\sigma_{0}^{2}}
$$

is minimized for $a=\left(\Sigma_{1}+\Sigma_{0}\right)^{-1}\left(M_{1}-M_{0}\right)$, where $M_{i}$ and $\Sigma_{i}$ are the mean vector and covariance matrix of $X$, given $Y=i$. Also, show that

$$
J_{2}(a)=\frac{\left(m_{1}-m_{0}\right)^{2}}{p \sigma_{1}^{2}+(1-p) \sigma_{0}^{2}}
$$

is minimized for $a=\left(p \Sigma_{1}+(1-p) \Sigma_{0}\right)^{-1}\left(M_{1}-M_{0}\right)$, where $p, 1-p$ are the class probabilities. This exercise shows that if discrimination is attempted in one dimension, we might consider projections $a^{T} X$ where $a$ maximizes the weighted distance between the projected means.

Problem 4.9. In the Fisher linear discriminant rule (4.1) with free parameter $a_{0}$, show that for any $\epsilon>0$, there exists a distribution for $(X, Y), X \in \mathcal{R}^{2}$, with $L=0$ and $\mathbf{E}\left\{\|X\|^{2}\right\}<\infty$ such that $\inf _{a_{0}} \mathbf{E}\left\{L\left(g_{a_{0}}\right)\right\}>1 / 2-\epsilon$. Moreover, if $a_{0}$ is chosen to minimize the squared error

$$
\mathbf{E}\left\{\left((2 Y-1)-a^{T} X-a_{0}\right)^{2}\right\},
$$

then $\mathbf{E}\left\{L\left(g_{a_{0}}\right)\right\}>1-\epsilon$.
Problem 4.10. Find a distribution of $(X, Y)$ with $X \in \mathcal{R}^{2}$ such that with probability at least one half, $\widehat{L}_{n}(\phi)$ is not unimodal with respect to the weight vector $\left(a, a_{0}\right)$.

Problem 4.11. The following observation may help in developing a fast algorithm to find the best linear classifier in certain cases. Assume that the Bayes rule is a linear split cutting through the origin, that is, $L^{*}=L\left(a^{*}\right)$ for some coefficient vector $a^{*} \in \mathcal{R}^{d}$, where $L(a)$ denotes the error probability of the classifier

$$
g_{a}(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{d} a_{i} x^{(i)} \geq 0 \\ 0 & \text { otherwise }\end{cases}
$$

and $a=\left(a_{1}, \ldots, a_{d}\right)$. Show that $L(a)$ is unimodal as a function of $a \in \mathcal{R}^{d}$, and $L(a)$ is monotone increasing along rays pointing from $a^{*}$, that is, for any $\lambda \in(0,1)$ and $a \in \mathcal{R}^{d}$, $L(a)-L\left(\lambda a+(1-\lambda) a^{*}\right) \geq 0$ (Fritz and Györfi (1976)). Hint: Use the expression

$$
L(a)=1 / 2-\int(\eta(x)-1 / 2) \operatorname{sign}\left(\sum_{i=1}^{d} a_{i} x^{(i)}\right) \mu(d x)
$$

to show that $L(a)-L\left(\lambda a+(1-\lambda) a^{*}\right)=\int_{A}|\eta(x)-1 / 2| \mu(d x)$ for some set $A \subset \mathcal{R}^{d}$.
Problem 4.12. Let $a=\left(a_{0}, a_{1}\right)$ and

$$
\widehat{a}=\underset{a}{\arg \min } \mathbf{E}\left\{\left((2 Y-1)-a_{1} X-a_{0}\right)^{2} I_{\left\{Y_{i} \neq g_{a}\left(X_{i}\right)\right\}}\right\},
$$

and $g_{a}(x)=I_{\left\{a_{1} x+a_{0}>0\right\}}$. Show that for every $\epsilon>0$, there exists a distribution of $(X, Y)$ on $\mathcal{R} \times\{0,1\}$ such that $L(\widehat{a})-L \geq 1-\epsilon$, where $L(\widehat{a})$ is the error probability for $g_{\widehat{a}}$. Hint: Argue as in the proof of Theorem 4.7. A distribution with four atoms suffices.

Problem 4.13. Repeat the previous exercise for

$$
\widehat{a}=\underset{a}{\arg \min } \mathbf{E}\left\{\left|(2 Y-1)-a_{1} X-a_{0}\right|\right\} .
$$

Problem 4.14. Let $\phi^{*}$ denote the linear discrimination rule that minimizes the mean square error $\mathbf{E}\left\{\left(2 Y-1-a^{T} X-a_{0}\right)^{2}\right\}$ over all $a$ and $a_{0}$. As this criterion is quadratic in $\left(a, a_{0}\right)$, it is unimodal. One usually approximates $\phi^{*}$ by $\widehat{\phi}$ by minimizing $\sum_{i}\left(2 Y_{i}-1-a^{T} X_{i}-a_{0}\right)^{2}$ over all $a$ and $a_{0}$. Show that the minimal column vector $\left(a, a_{0}\right)$ is given by

$$
\left(\sum_{i} X_{i}^{\prime} X_{i}^{\prime}\right)^{-1}\left(\sum_{i}\left(2 Y_{i}-1\right) X_{i}^{\prime}\right)
$$

where $X_{i}^{\prime}=\left(X_{i}, 1\right)$ is a $(d+1)$-dimensional column vector.
Problem 4.15. The perceptron criterion is

$$
J=\sum_{i: 2 Y_{i}-1 \neq \operatorname{sign}\left(a^{T} X_{i}+a_{0}\right)}\left|a^{T} X_{i}+a_{0}\right| .
$$

Find a distribution for which $L^{*}=0, L \leq 1 / 4$, yet $\lim _{\inf }^{n \rightarrow \infty} \boldsymbol{E}\left\{L_{n}(\phi)\right\} \geq 1 / 2$, where $\phi$ is the linear discrimination rule obtained by using the $a$ and $a_{0}$ that minimize $J$.

PROBLEM 4.16. Let $\sigma$ be a monotone nondecreasing function on $\mathcal{R}$ satisfying $\lim _{u \rightarrow-\infty} \sigma(u)$ $=0$ and $\lim _{u \rightarrow \infty} \sigma(u)=1$. For $h>0$, define $\sigma_{h}(u)=\sigma(h u)$. Consider the linear discrimination rule $\widehat{\phi}$ with $a$ and $a_{0}$ chosen to minimize

$$
\sum_{i=1}^{n}\left(\sigma_{h}\left(a^{T} X_{i}+a_{0}\right)-Y_{i}\right)^{2}
$$

For every fixed $h>0$ and $0<\epsilon<1$, exhibit a distribution with $L<\epsilon$ and

$$
\liminf _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}(\widehat{\phi})\right\}>1-\epsilon .
$$

On the other hand, show that if $h$ depends on the sample size $n$ such that $h \rightarrow \infty$ as $n \rightarrow \infty$, then for all distributions, $\mathbf{E}\left\{L_{n}(\widehat{\phi})\right\} \rightarrow L$.

Problem 4.17. Given $Y=i$, let $X$ be normal with mean $m_{i}$ and covariance matrix $\Sigma_{i}$, $i=0,1$. Consider discrimination based upon the minimization of the criterion

$$
\mathbf{E}\left\{\left(Y-\sigma\left(X^{T} A X+w^{T} X+c\right)\right)^{2}\right\}
$$

with respect to $A, w$, and $c$, a $d \times d$ matrix, $d \times 1$ vector and constant respectively, where $\sigma(u)=1 /\left(1+e^{-u}\right)$ is the standard sigmoid function. Show that this is minimized for the same $A, w$, and $c$ that minimize the probability of error

$$
\mathbf{P}\left\{2 Y-1 \neq \operatorname{sign}\left(X^{T} A X+w^{T} X+c\right)\right\}
$$

and conclude that in this particular case, the squared error criterion may be used to obtain a Bayes-optimal classifier (Horne and Hush (1990)).

## 5

## Nearest Neighbor Rules

### 5.1 Introduction

Simple rules survive. The $k$-nearest neighbor rule, since its conception in 1951 and 1952 (Fix and Hodges (1951; 1952; 1991a; 1991b)), has thus attracted many followers and continues to be studied by many researchers. Formally, we define the $k$-NN rule by

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{n} w_{n i} I_{\left\{Y_{i}=1\right\}}>\sum_{i=1}^{n} w_{n i} I_{\left\{Y_{i}=0\right\}} \\ 0 & \text { otherwise, }\end{cases}
$$

where $w_{n i}=1 / k$ if $X_{i}$ is among the $k$ nearest neighbors of $x$, and $w_{n i}=0$ elsewhere. $X_{i}$ is said to be the $k$-th nearest neighbor of $x$ if the distance $\left\|x-X_{i}\right\|$ is the $k$-th smallest among $\left\|x-X_{1}\right\|, \ldots,\left\|x-X_{n}\right\|$. In case of a distance tie, the candidate with the smaller index is said to be closer to $x$. The decision is based upon a majority vote. It is convenient to let $k$ be odd, to avoid voting ties. Several issues are worth considering:
(A) Universal consistency. Establish convergence to the Bayes rule if $k \rightarrow \infty$ and $k / n \rightarrow 0$ as $n \rightarrow \infty$. This is dealt with in Chapter 11 .
(B) Finite $k$ performance. What happens if we hold $k$ fixed and let $n$ tend to infinity?
(C) The choice of the weight vector $\left(w_{n 1}, \ldots, w_{n n}\right)$. Are equal weights for the $k$ nearest neighbors better than unequal weights in some sense?
(D) The choice of a distance metric. Achieve invariance with respect to a certain family of transformations.
(E) The reduction of the data size. Can we obtain good performance when the data set is edited and/or reduced in size to lessen the storage load?


> Figure 5.1. At every point the decision is the label of the closest data point. The set of points whose nearest neighbor is $X_{i}$ is called the Voronoi cell of $X_{i}$. The partition induced by the Voronoi cells is a Voronoi partition. A Voronoi partition of 15 random points is shown here.

In the first couple of sections, we will be concerned with convergence issues for $k$ nearest neighbor rules when $k$ does not change with $n$. In particular, we will see that for all distributions, the expected error probability $\mathbf{E}\left\{L_{n}\right\}$ tends to a limit $L_{k \mathrm{NN}}$ that is in general close to but larger than $L^{*}$. The methodology for obtaining this result is interesting in its own right. The expression for $L_{k \mathrm{NN}}$ is then studied, and several key inequalities such as $L_{\mathrm{NN}} \leq 2 L^{*}$ (Cover and Hart (1967)) and $L_{k_{N N}} \leq L^{*}(1+\sqrt{2 / k})$ are proved and applied. The other issues mentioned above are dealt with in the remaining sections. For surveys of various aspects of the nearest neighbor or related methods, see Dasarathy (1991), Devijver (1980), or Devroye and Wagner (1982).

Remark. Computational concerns. Storing the $n$ data pairs in an array and searching for the $k$ nearest neighbors may take time proportional to $n k d$ if done in a naive manner-the " $d$ " accounts for the cost of one distance computation. This complexity may be reduced in terms of one or more of the three factors involved. Typically, with $k$ and $d$ fixed, $O\left(n^{1 / d}\right)$ worst-case time (Papadimitriou and Bentley (1980)) and $O(\log n)$ expected time (Friedman, Bentley, and Finkel (1977)) may be achieved. Multidimensional search trees that partition the space and guide the search are invaluable-for this approach, see Fukunaga and Narendra (1975), Friedman, Bentley, and Finkel (1977), Niemann and Goppert (1988), Kim and Park (1986), and Broder (1990). We refer to a survey in Dasarathy (1991) for more references. Other approaches are described by Yunck (1976), Friedman, Baskett, and Shustek (1975), Vidal (1986), Sethi (1981), and Faragó, Linder, and Lugosi (1993). Generally, with preprocessing, one may considerably reduce the overall complexity in terms of $n$ and $d$.

### 5.2 Notation and Simple Asymptotics

We fix $x \in \mathcal{R}^{d}$, and reorder the data ( $X_{1}, Y_{1}$ ), $\ldots,\left(X_{n}, Y_{n}\right)$ according to increasing values of $\left\|X_{i}-x\right\|$. The reordered data sequence is denoted by

$$
\left(X_{(1)}(x), Y_{(1)}(x)\right), \ldots,\left(X_{(n)}(x), Y_{(n)}(x)\right) \text { or by }\left(X_{(1)}, Y_{(1)}\right), \ldots,\left(X_{(n)}, Y_{(n)}\right)
$$

if no confusion is possible. $X_{(k)}(x)$ is the $k$-th nearest neighbor of $x$.
Remark. We note here that, rather arbitrarily, we defined neighbors in terms of the Euclidean distance $\|x-y\|$. Surprisingly, the asymptotic properties derived in this chapter remain valid to a wide variety of metrics-the asymptotic probability of error is independent of the distance measure (see Problem 5.1).

Denote the probability measure for $X$ by $\mu$, and let $S_{x, \epsilon}$ be the closed ball centered at $x$ of radius $\epsilon>0$. The collection of all $x$ with $\mu\left(S_{x, \epsilon}\right)>0$ for all $\epsilon>0$ is called the support of $X$ or $\mu$. This set plays a key role because of the following property.

Lemma 5.1. If $x \in \operatorname{support}(\mu)$ and $\lim _{n \rightarrow \infty} k / n=0$, then $\left\|X_{(k)}(x)-x\right\| \rightarrow 0$ with probability one. If $X$ is independent of the data and has probability measure $\mu$, then $\left\|X_{(k)}(X)-X\right\| \rightarrow 0$ with probability one whenever $k / n \rightarrow 0$.

Proof. Take $\epsilon>0$. By definition, $x \in \operatorname{support}(\mu)$ implies that $\mu\left(S_{x, \epsilon}\right)>0$. Observe that $\left\|X_{(k)}(x)-x\right\|>\epsilon$ if and only if

$$
\frac{1}{n} \sum_{i=1}^{n} I_{\left\{X_{i} \in S_{x, t}\right\}}<\frac{k}{n} .
$$

By the strong law of large numbers, the left-hand side converges to $\mu\left(S_{x, \epsilon}\right)>0$ with probability one, while, by assumption, the right-hand side tends to zero. Therefore, $\left\|X_{(k)}(x)-x\right\| \rightarrow 0$ with probability one.
The second statement follows from the previous argument as well. First note that by Lemma A. 1 in the Appendix, $\mathbf{P}\{X \in \operatorname{support}(\mu)\}=1$, therefore for every $\epsilon>0$,

$$
\begin{aligned}
\mathbf{P} & \left\{\left\|X_{(k)}(X)-X\right\|>\epsilon\right\} \\
& =\mathbf{E}\left\{I_{\{X \in \operatorname{support}(\mu)\}} \mathbf{P}\left\{\left\|X_{(k)}(X)-X\right\|>\epsilon \mid X \in \operatorname{support}(\mu)\right\}\right\},
\end{aligned}
$$

which converges to zero by the dominated convergence theorem, proving convergence in probability. If $k$ does not change with $n$, then $\left\|X_{(k)}(X)-X\right\|$ is monotone decreasing for $n \geq k$; therefore, it converges with probability one as well. If $k=k_{n}$ is allowed to grow with $n$ such that $k / n \rightarrow 0$, then using the notation $X_{\left(k_{n}, n\right)}(X)=X_{(k)}(X)$, we see by a similar argument that the sequence of monotone decreasing random variables

$$
\sup _{m \geq n}\left\|X_{\left(k_{m}, m\right)}(X)-X\right\| \geq\left\|X_{\left(k_{n}, n\right)}(X)-X\right\|
$$

converges to zero in probability, and therefore, with probability one as well. This completes the proof.

Because $\eta$ is measurable (and thus well-behaved in a general sense) and $\left\|X_{(k)}(x)-x\right\|$ is small, the values $\eta\left(X_{(i)}(x)\right)$ should be close to $\eta(x)$ for all $i$ small enough. We now introduce a proof method that exploits this fact, and will make subsequent analyses very simple-it suffices to look at data samples in a new way via embedding. The basic idea is to define an auxiliary rule $g_{n}^{\prime}(x)$ in which the $Y_{(i)}(x)$ 's are replaced by $k$ i.i.d. Bernoulli random variables with parameter $\eta(x)$-locally, the $Y_{(i)}(x)$ 's behave in such a way. It is easy to show that the error probabilities of the two rules are close, and analyzing the behavior of the auxiliary rule is much more convenient.

To make things more precise, we assume that we are given i.i.d. data pairs $\left(X_{1}, U_{1}\right), \ldots,\left(X_{n}, U_{n}\right)$, all distributed as $(X, U)$, where $X$ is as before (and has probability measure $\mu$ on the Borel sets of $\mathcal{R}^{d}$ ), and $U$ is uniformly distributed on [ 0,1$]$ and independent of $X$. If we set $Y_{i}=I_{\left\{U_{i} \leq \eta\left(X_{i}\right)\right\}}$, then $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ are i.i.d. and distributed as the prototype pair ( $X, Y$ ). So why bother with the $U_{i}$ 's? In embedding arguments, we will use the same $U_{i}$ 's to construct a second data sequence that is heavily correlated (coupled) with the original data sequence, and is more convenient to analyze. For example, for fixed $x \in \mathcal{R}^{d}$, we may define

$$
Y_{i}^{\prime}(x)=I_{\left\{U_{i} \leq \eta(x)\right\}} .
$$

We now have an i.i.d. sequence with $i$-th vector given by $X_{i}, Y_{i}, Y_{i}^{\prime}(x), U_{i}$. Reordering the data sequence according to increasing values of $\left\|X_{i}-x\right\|$ yields a new sequence with the $i$-th vector denoted by $X_{(i)}(x), Y_{(i)}(x), Y_{(i)}^{\prime}(x), U_{(i)}(x)$. If no confusion is possible, the argument $x$ will be dropped. A rule is called $k$-local if for $n \geq k, g_{n}$ is of the form

$$
g_{n}(x)= \begin{cases}1 & \text { if } \psi\left(x, Y_{(1)}(x), \ldots, Y_{(k)}(x)\right)>0  \tag{5.1}\\ 0 & \text { otherwise }\end{cases}
$$

for some function $\psi$. For the $k$-NN rule, we have, for example,

$$
\psi\left(x, Y_{(1)}, \ldots, Y_{(k)}\right)=\sum_{i=1}^{k} Y_{(i)}(x)-\frac{k}{2}
$$

In other words, $g_{n}$ takes a majority vote over the $k$ nearest neighbors of $x$ and breaks ties in favor of class 0 .

To study $g_{n}$ turns out to be almost equivalent to studying the approximate rule $g_{n}^{\prime}$ :

$$
g_{n}^{\prime}(x)= \begin{cases}1 & \text { if } \psi\left(x, Y_{(1)}^{\prime}(x), \ldots, Y_{(k)}^{\prime}(x)\right)>0 \\ 0 & \text { otherwise }\end{cases}
$$

The latter rule is of no practical value because it requires the knowledge of $\eta(x)$. Interestingly however, it is easier to study, as $Y_{(1)}^{\prime}(x), \ldots, Y_{(k)}^{\prime}(x)$ are i.i.d., whereas $Y_{(1)}(x), \ldots, Y_{(k)}(x)$ are not. Note, in particular, the following:

Lemma 5.2. For all $x, n \geq k$,

$$
\begin{aligned}
\mathbf{P}\left\{\psi\left(x, Y_{(1)}(x), \ldots, Y_{(k)}(x)\right)\right. & \left.\neq \psi\left(x, Y_{(1)}^{\prime}(x), \ldots, Y_{(k)}^{\prime}(x)\right)\right\} \\
& \leq \sum_{i=1}^{k} \mathbf{E}\left\{\left|\eta(x)-\eta\left(X_{(i)}(x)\right)\right|\right\}
\end{aligned}
$$

and

$$
\mathbf{P}\left\{g_{n}(x) \neq g_{n}^{\prime}(x)\right\} \leq \sum_{i=1}^{k} \mathbf{E}\left\{\left|\eta(x)-\eta\left(X_{(i)}(x)\right)\right|\right\}
$$

Proof. Both statements follow directly from the observation that

$$
\begin{aligned}
& \left\{\psi\left(x, Y_{(1)}(x), \ldots, Y_{(k)}(x)\right) \neq \psi\left(x, Y_{(1)}^{\prime}(x), \ldots, Y_{(k)}^{\prime}(x)\right)\right\} \\
& \quad \subseteq\left\{\left(Y_{(1)}(x), \ldots, Y_{(k)}(x)\right) \neq\left(Y_{(1)}^{\prime}(x), \ldots, Y_{(k)}^{\prime}(x)\right)\right\} \\
& \quad \subseteq \bigcup_{i=1}^{k}\left\{\eta\left(X_{(i)}(x)\right) \leq U_{(i)}(x) \leq \eta(x)\right\} \cup \bigcup_{i=1}^{k}\left\{\eta(x) \leq U_{(i)}(x) \leq \eta\left(X_{(i)}(x)\right)\right\},
\end{aligned}
$$

and using the union bound and the fact that the $U_{(i)}(x)$ 's are uniform $[0,1]$.

We need the following result, in which $X$ is distributed as $X_{1}$, but independent of the data sequence:

Lemma 5.3. (STONE (1977)). For any integrable function $f$, any $n$, and any $k \leq n$,

$$
\sum_{i=1}^{k} \mathbf{E}\left\{\left|f\left(X_{(i)}(X)\right)\right|\right\} \leq k \gamma_{d} \mathbf{E}\{|f(X)|\}
$$

where $\gamma_{d} \leq(1+2 / \sqrt{2-\sqrt{3}})^{d}-1$ depends upon the dimension only.
The proof of this lemma is beautiful but a bit technical-it is given in a separate section. Here is how it is applied, and why, for fixed $k$, we may think of $f\left(X_{(k)}(X)\right)$ as $f(X)$ for all practical purposes.

Lemma 5.4. For any integrable function $f$,

$$
\frac{1}{k} \sum_{i=1}^{k} \mathbf{E}\left\{\left|f(X)-f\left(X_{(i)}(X)\right)\right|\right\} \rightarrow 0
$$

as $n \rightarrow \infty$ whenever $k / n \rightarrow 0$.

Proof. Given $\epsilon>0$, find a uniformly continuous function $g$ vanishing off a bounded set $A$, such that $\mathbf{E}\{|g(X)-f(X)|\}<\epsilon$ (see Theorem A. 8 in the Appendix).

Then for each $\epsilon>0$, there is a $\delta>0$ such that $\|x-z\|<\delta$ implies $|g(x)-g(z)|<$ $\epsilon$. Thus,

$$
\begin{aligned}
\frac{1}{k} & \sum_{i=1}^{k} \mathbf{E}\left\{\left|f(X)-f\left(X_{(i)}(X)\right)\right|\right\} \\
\leq & \mathbf{E}\{|f(X)-g(X)|\}+\frac{1}{k} \sum_{i=1}^{k} \mathbf{E}\left\{\left|g(X)-g\left(X_{(i)}(X)\right)\right|\right\} \\
& \quad+\frac{1}{k} \sum_{i=1}^{k} \mathbf{E}\left\{\left|g\left(X_{(i)}(X)\right)-f\left(X_{(i)}(X)\right)\right|\right\} \\
\leq & \left(1+\gamma_{d}\right) \mathbf{E}\{|f(X)-g(X)|\}+\epsilon+\|g\|_{\infty} \mathbf{P}\left\{\left\|X-X_{(k)}(X)\right\|>\delta\right\} \\
& (\text { by Lemma 5.3, where } \delta \text { depends on } \epsilon \text { only) } \\
\leq & \left(2+\gamma_{d}\right) \epsilon+o(1) \quad \text { (by Lemma 5.1). }
\end{aligned}
$$

### 5.3 Proof of Stone's Lemma

In this section we prove Lemma 5.3. For $\theta \in(0, \pi / 2)$, a cone $C(x, \theta)$ is the collection of all $y \in \mathcal{R}^{d}$ for which angle $(x, y) \leq \theta$. Equivalently, in vector notation, $x^{T} y /\|x\|\|y\| \geq \cos \theta$. The set $z+C(x, \theta)$ is the translation of $C(x, \theta)$ by $z$.


Figure 5.2. A cone of angle $\theta$.

If $y, z \in C(x, \pi / 6)$, and $\|y\|<\|z\|$, then $\|y-z\|<\|z\|$, as we will now show. Indeed,

$$
\begin{aligned}
\|y-z\|^{2} & =\|y\|^{2}+\|z\|^{2}-2\|y\|\|z\| \frac{y^{T} z}{\|y\|\|z\|} \\
& \leq\|y\|^{2}+\|z\|^{2}-2\|y\|\|z\| \cos (\pi / 3) \\
& =\|z\|^{2}\left(1+\frac{\|y\|^{2}}{\|z\|^{2}}-\frac{\|y\|}{\|z\|}\right) \\
& <\|z\|^{2} \quad \quad \quad \text { (see Figure 5.3) } .
\end{aligned}
$$



Figure 5.3. The key geometrical property of cones of angle $<\pi / 2$.

The following covering lemma is needed in what follows:

Lemma 5.5. Let $\theta \in(0, \pi / 2)$ be fixed. Then there exists a set $\left\{x_{1}, \ldots, x_{\gamma_{d}}\right\} \subset \mathcal{R}^{d}$ such that

$$
\mathcal{R}^{d}=\bigcup_{i=1}^{\gamma_{d}} C\left(x_{i}, \theta\right)
$$

Furthermore, it is always possible to take

$$
\gamma_{d} \leq\left(1+\frac{1}{\sin (\theta / 2)}\right)^{d}-1
$$

For $\theta=\pi / 6$, we have

$$
\gamma_{d} \leq\left(1+\frac{2}{\sqrt{2-\sqrt{3}}}\right)^{d}-1
$$

Proof. We assume without loss of generality that $\left\|x_{i}\right\|=1$ for all $i$. Each $x_{i}$ is the center of a sphere $S_{i}$ of radius $r=2 \sin (\theta / 2) . S_{i}$ has the property that

$$
\{x:\|x\|=1\} \cap S_{i}=\{x:\|x\|=1\} \cap C\left(x_{i}, \theta\right)
$$

Let us only look at $x_{i}$ 's such that $\left\|x_{i}-x_{j}\right\| \geq r$ for all $j \neq i$. In that case, $\cup C\left(x_{i}, \theta\right)$ covers $\mathcal{R}^{d}$ if and only if $\bigcup S_{i}$ covers $\{x:\|x\|=1\}$. Then the spheres $S_{i}^{\prime}$ of radius $r / 2$ centered at the $x_{i}$ 's are disjoint and $\bigcup S_{i}^{\prime} \subseteq S_{0,1+r / 2}-S_{0, r / 2}$ (see Figure 5.4).


Figure 5.4. Bounding $\gamma_{d}$.

Thus, if $v_{d}=\operatorname{volume}\left(S_{0,1}\right)$,

$$
\gamma_{d} v_{d}\left(\frac{r}{2}\right)^{d} \leq v_{d}\left(1+\frac{r}{2}\right)^{d}-v_{d}\left(\frac{r}{2}\right)^{d}
$$

or

$$
\gamma_{d} \leq\left(1+\frac{2}{r}\right)^{d}-1=\left(1+\frac{1}{\sin (\theta / 2)}\right)^{d}-1
$$

The last inequality follows from the fact that

$$
\sin \frac{\pi}{12}=\sqrt{\frac{1-\cos (\pi / 6)}{2}}=\sqrt{\frac{2-\sqrt{3}}{4}}
$$



Figure 5.5. Covering the space by cones.

With the preliminary results out of the way, we cover $\mathcal{R}^{d}$ by $\gamma_{d}$ cones $X+$ $C\left(x_{j}, \pi / 6\right), 1 \leq j \leq \gamma_{d}$, and mark in each cone the $X_{i}$ that is nearest to
$X$, if such an $X_{i}$ exists. If $X_{i}$ belongs to $X+C\left(x_{j}, \pi / 6\right)$ and is not marked, then $X$ cannot be the nearest neighbor of $X_{i}$ in $\left\{X_{1}, \ldots, X_{i-1}, X, X_{i+1}, \ldots, X_{n}\right\}$. Similarly, we might mark all $k$ nearest neighbors of $X$ in each cone (if there are less than $k$ points in a cone, mark all of them). By a similar argument, if $X_{i} \in X+C\left(x_{j}, \pi / 6\right)$ is not marked, then $X$ cannot be among the $k$ nearest neighbors of $X_{i}$ in $\left\{X_{1}, \ldots, X_{i-1}, X, X_{i+1}, \ldots, X_{n}\right\}$. (The order of this set of points is important if distance ties occur with positive probability, and they are broken by comparing indices.) Therefore, if $f$ is a nonnegative function,

$$
\begin{aligned}
& \sum_{i=1}^{k} \mathbf{E}\left\{f\left(X_{(i)}(X)\right)\right\} \\
&= \mathbf{E}\left\{\sum_{i=1}^{n} I_{\left\{X_{i} \text { is among the } \mathrm{k} \text { nearest neighbors of } X_{\text {in }}\left\{X_{1}, \ldots, X_{n}\right\}\right\}} f\left(X_{i}\right)\right\} \\
&= \mathbf{E}\left\{f(X) \sum_{i=1}^{n}\right. \\
&\left.I_{\left\{X \text { is among the } \mathrm{k} \text { nearest neighbors of } X_{i} \text { in }\left\{X_{1}, \ldots, X_{i-1}, X, X_{i+1}, \ldots, X_{n}\right\}\right\}}\right\} \\
&\left(\text { by exchanging } X \text { and } X_{i}\right) \\
& \leq \mathbf{E}\left\{f(X) \sum_{i=1}^{n} I_{\left\{X_{i} \text { is marked }\right\}}\right\} \\
& \leq k \gamma_{d} \mathbf{E}\{f(X)\},
\end{aligned}
$$

as we can mark at most $k$ nodes in each cone, and the number of cones is at most $\gamma_{d}$-see Lemma 5.5. This concludes the proof of Stone's lemma.

### 5.4 The Asymptotic Probability of Error

We return to $k$-local rules (and in particular, to $k$-nearest neighbor rules). Let $D_{n}^{\prime}=\left(\left(X_{1}, Y_{1}, U_{1}\right), \ldots,\left(X_{n}, Y_{n}, U_{n}\right)\right)$ be the i.i.d. data augmented by the uniform random variables $U_{1}, \ldots, U_{n}$ as described earlier. For a decision $g_{n}$ based on $D_{n}$, we have the probability of error

$$
\begin{aligned}
L_{n} & =\mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}^{\prime}\right\} \\
& =\mathbf{P}\left\{\operatorname{sign}\left(\psi\left(X, Y_{(1)}(X), \ldots, Y_{(k)}(X)\right)\right) \neq \operatorname{sign}(2 Y-1) \mid D_{n}^{\prime}\right\}
\end{aligned}
$$

where $\psi$ is the function whose sign determines $g_{n}$; see (5.1). Define the random variables $Y_{(1)}^{\prime}(X), \ldots, Y_{(k)}^{\prime}(X)$ as we did earlier, and set

$$
L_{n}^{\prime}=\mathbf{P}\left\{\operatorname{sign}\left(\psi\left(X, Y_{(1)}^{\prime}(X), \ldots, Y_{(k)}^{\prime}(X)\right)\right) \neq \operatorname{sign}(2 Y-1) \mid D_{n}^{\prime}\right\}
$$

By Lemmas 5.2 and 5.4,

$$
\begin{aligned}
& \mathbf{E}\left\{\left|L_{n}-L_{n}^{\prime}\right|\right\} \\
& \quad \leq \mathbf{P}\left\{\psi\left(X, Y_{(1)}(X), \ldots, Y_{(k)}(X)\right) \neq \psi\left(X, Y_{(1)}^{\prime}(X), \ldots, Y_{(k)}^{\prime}(X)\right)\right\} \\
& \quad \leq \sum_{i=1}^{k} \mathbf{E}\left\{\left|\eta(X)-\eta\left(X_{(i)}(X)\right)\right|\right\} \\
& \quad=o(1) .
\end{aligned}
$$

Because $\lim _{n \rightarrow \infty}\left(\mathbf{E} L_{n}^{\prime}-\mathbf{E} L_{n}\right)=0$, we need only study the rule $g_{n}^{\prime}$

$$
g_{n}^{\prime}(x)= \begin{cases}1 & \text { if } \psi\left(x, Z_{1}, \ldots, Z_{k}\right)>0 \\ 0 & \text { otherwise }\end{cases}
$$

$\left(Z_{1}, \ldots, Z_{k}\right.$ are i.i.d. Bernoulli $\left.(\eta(x))\right)$ unless we are concerned with the closeness of $L_{n}$ to $\mathbf{E} L_{n}$ as well.

We now illustrate this important time-saving device on the 1-nearest neighbor rule. Clearly, $\psi\left(x, Z_{1}\right)=2 Z_{1}-1$, and therefore

$$
\mathbf{E}\left\{L_{n}^{\prime}\right\}=\mathbf{P}\left\{Z_{1} \neq Y\right\}=\mathbf{E}\{2 \eta(X)(1-\eta(X))\}
$$

We have, without further work:
Theorem 5.1. For the nearest neighbor rule, for any distribution of $(X, Y)$,

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}=\mathbf{E}\{2 \eta(X)(1-\eta(X))\}=L_{\mathrm{NN}}
$$

Under various continuity conditions ( $X$ has a density $f$ and both $f$ and $\eta$ are almost everywhere continuous); this result is due to Cover and Hart (1967). In the present generality, it essentially appears in Stone (1977). See also Devroye (1981c). Elsewhere (Chapter 3), we show that

$$
L^{*} \leq L_{\mathrm{NN}} \leq 2 L^{*}\left(1-L^{*}\right) \leq 2 L^{*}
$$

Hence, the previous result says that the nearest neighbor rule is asymptotically at most twice as bad as the Bayes rule-especially for small $L^{*}$, this property should be useful.

We formally define the quantity, when $k$ is odd,

$$
\begin{aligned}
& L_{k \mathrm{NN}} \\
& \quad=\mathbf{E}\left\{\sum_{j=0}^{k}\binom{k}{j} \eta^{j}(X)(1-\eta(X))^{k-j}\left(\eta(X) I_{\{j<k / 2\}}+(1-\eta(X)) I_{\{j>k / 2\}}\right)\right\} .
\end{aligned}
$$

We have the following result:
Theorem 5.2. Let $k$ be odd and fixed. Then, for the $k$-NN rule,

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}=L_{k \mathrm{NN}} .
$$

Proof. We note that it suffices to show that $\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}^{\prime}\right\}=L_{k \mathrm{NN}}$ (in the previously introduced notation). But for every $n$,

$$
\begin{aligned}
& \mathbf{E}\left\{L_{n}^{\prime}\right\} \\
& \quad=\mathbf{P}\left\{Z_{1}+\cdots+Z_{k}>\frac{k}{2}, Y=0\right\}+\mathbf{P}\left\{Z_{1}+\cdots+Z_{k}<\frac{k}{2}, Y=1\right\} \\
& \quad=\mathbf{P}\left\{Z_{1}+\cdots+Z_{k}>\frac{k}{2}, Z_{0}=0\right\}+\mathbf{P}\left\{Z_{1}+\cdots+Z_{k}<\frac{k}{2}, Z_{0}=1\right\}
\end{aligned}
$$

(where $Z_{0}, \ldots, Z_{k}$ are i.i.d. Bernoulli $(\eta(X))$ random variables),
which leads directly to the sought result.

Several representations of $L_{k N N}$ will be useful for later analysis. For example, we have

$$
\begin{aligned}
& L_{k N \mathrm{~N}} \\
&= \mathbf{E}\left\{\eta(X) \mathbf{P}\left\{\left.\operatorname{Binomial}(k, \eta(X))<\frac{k}{2} \right\rvert\, X\right\}\right\} \\
&+\mathbf{E}\left\{(1-\eta(X)) \mathbf{P}\left\{\left.\operatorname{Binomial}(k, \eta(X))>\frac{k}{2} \right\rvert\, X\right\}\right\} \\
&= \mathbf{E}\{\min (\eta(X), 1-\eta(X))\} \\
&+\mathbf{E}\left\{(1-2 \min (\eta(X), 1-\eta(X))) \mathbf{P}\left\{\left.\operatorname{Binomial}(k, \eta(X))>\frac{k}{2} \right\rvert\, X\right\}\right\} .
\end{aligned}
$$

It should be stressed that the limit result in Theorem 5.2 is distribution-free. The limit $L_{k N N}$ depends upon $\eta(X)$ (or $\min (\eta(X), 1-\eta(X))$ ) only. The continuity or lack of smoothness of $\eta$ is immaterial-it only matters for the speed with which $\mathbf{E}\left\{L_{n}\right\}$ approaches the limit $L_{k N \mathrm{~N}}$.

### 5.5 The Asymptotic Error Probability of Weighted Nearest Neighbor Rules

Following Royall (1966), a weighted nearest neighbor rule with weights $w_{1}, \ldots$, $w_{k}$ makes a decision according to

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i: Y_{(i)}(x)=1}^{k} w_{i}>\sum_{i: Y_{(i)}(x)=0}^{k} w_{i} \\ 0 & \text { otherwise }\end{cases}
$$

In case of a voting tie, this rule is not symmetric. We may modify it so that $g_{n}(x) \stackrel{\text { def }}{=}$ -1 if we have a voting tie. The " -1 " should be considered as an indecision. By
previous arguments the asymptotic probability of error is a function of $w_{1}, \ldots, w_{k}$ given by $L\left(w_{1}, \ldots, w_{k}\right)=\mathbf{E}\{\alpha(\eta(X))\}$, where

$$
\begin{aligned}
\alpha(p)= & \mathbf{P}\left\{\sum_{i=1}^{k} w_{i} Y_{i}^{\prime}>\sum_{i=1}^{k} w_{i}\left(1-Y_{i}^{\prime}\right)\right\}(1-p) \\
& +\mathbf{P}\left\{\sum_{i=1}^{k} w_{i} Y_{i}^{\prime} \leq \sum_{i=1}^{k} w_{i}\left(1-Y_{i}^{\prime}\right)\right\} p
\end{aligned}
$$

where now $Y_{1}^{\prime}, \ldots, Y_{k}^{\prime}$ are i.i.d. Bernoulli $(p)$. Equivalently, with $Z_{i}=2 Y_{i}^{\prime}-1 \in$ $\{-1,1\}$,

$$
\alpha(p)=(1-p) \mathbf{P}\left\{\sum_{i=1}^{k} w_{i} Z_{i}>0\right\}+p \mathbf{P}\left\{\sum_{i=1}^{k} w_{i} Z_{i} \leq 0\right\}
$$

Assume that $\mathbf{P}\left\{\sum_{i=1}^{k} w_{i} Z_{i}=0\right\}=0$ for now. Then, if $p<1 / 2$,

$$
\alpha(p)=p+(1-2 p) \mathbf{P}\left\{\sum_{i=1}^{k} w_{i} Z_{i}>0\right\}
$$

and an antisymmetric expression is valid when $p>1 / 2$. Note next the following. If we let $N_{l}$ be the number of vectors $z=\left(z_{1}, \ldots, z_{k}\right) \in\{-1,1\}^{k}$ with $\sum I_{\left\{z_{i}=1\right\}}=l$ and $\sum w_{i} z_{i}>0$, then $N_{l}+N_{k-l}=\binom{k}{l}$. Thus,

$$
\begin{aligned}
& \mathbf{P}\left\{\sum_{i=1}^{k} w_{i} Z_{i}>0\right\} \\
& = \\
& =\sum_{l=0}^{k} N_{l} p^{l}(1-p)^{k-l} \\
& =\sum_{l<k / 2}\binom{k}{l} p^{k-l}(1-p)^{l}+\sum_{l<k / 2} N_{l}\left(p^{l}(1-p)^{k-l}-p^{k-l}(1-p)^{l}\right) \\
& \\
& \quad+\frac{1}{2}\binom{k}{k / 2} p^{k / 2}(1-p)^{k / 2} I_{\{k \text { even }\}} \\
& = \\
& I+I I+I I I .
\end{aligned}
$$

Note that $I+I I I$ does not depend on the vector of weights, and represents

$$
\mathbf{P}\{\operatorname{Binomial}(k, 1-p) \leq k / 2\}=\mathbf{P}\{\operatorname{Binomial}(k, p) \geq k / 2\}
$$

Finally, since $p \leq 1 / 2$,

$$
\begin{aligned}
I I & =\sum_{l<k / 2} N_{l}\left(p^{l}(1-p)^{k-l}-p^{k-l}(1-p)^{l}\right) \\
& =\sum_{l<k / 2} N_{l} p^{l}(1-p)^{l}\left((1-p)^{k-2 l}-p^{k-2 l}\right) \\
& \geq 0
\end{aligned}
$$

This term is zero if and only if $N_{l}=0$ for all $l<k / 2$. In other words, it vanishes if and only if no numerical minority of $w_{i}$ 's can sum to a majority (as in the case ( $0.7,0.2,0.1$ ), where 0.7 alone, a numerical minority, outweighs the others). But such cases are equivalent to ordinary $k$-nearest neighbor rules if $k$ is odd. When $k$ is even, and we add a tiny weight to one $w_{i}$, as in

$$
\left(\frac{1+\epsilon}{k}, \frac{1-\epsilon /(k-1)}{k}, \ldots, \frac{1-\epsilon /(k-1)}{k}\right)
$$

for small $\epsilon>0$, then no numerical minority can win either, and we have an optimal rule $(I I=0)$. We have thus shown the following:

Theorem 5.3. (Bailey and Jain (1978)). Let $L\left(w_{1}, \ldots, w_{k}\right)$ be the asymptotic probability of error of the weighted $k$-NN rule with weights $w_{1}, \ldots, w_{k}$. Let the $k$-NN rule be defined by $(1 / k, 1 / k, \ldots, 1 / k)$ if $k$ is odd, and by

$$
(1 / k, 1 / k, \ldots, 1 / k)+\epsilon(1,-1 /(k-1),-1 /(k-1), \ldots,-1 /(k-1))
$$

for $0<\epsilon<1 / k$ when $k$ is even. Denoting the asymptotic probability of error by $L_{k \mathrm{NN}}$ for the latter rule, we have

$$
L\left(w_{1}, \ldots, w_{k}\right) \geq L_{k \mathrm{NN}}
$$

If $\mathbf{P}\{\eta(X)=1 / 2\}<1$, then equality occurs if and only if every numerical minority of the $w_{i}$ 's carries less than half of the total weight.

The result states that standard $k$-nearest neighbor rules are to be preferred in an asymptotic sense. This does not mean that for a particular sample size, one should steer clear of nonuniform weights. In fact, if $k$ is allowed to vary with $n$, then nonuniform weights are advantageous (Royall (1966)).

Consider the space $\mathcal{W}$ of all weight vectors $\left(w_{1}, \ldots, w_{k}\right)$ with $w_{i} \geq 0, \sum_{i=1}^{k} w_{i}=$ 1. Is it totally ordered with respect to $L\left(w_{1}, \ldots, w_{k}\right)$ or not? To answer this question, we must return to $\alpha(p)$ once again. The weight vector only influences the term II given there. Consider, for example, the weight vectors

$$
(0.3,0.22,0.13,0.12,0.071,0.071,0.071,0.017)
$$

and $\quad(0.26,0.26,0.13,0.12,0.071,0.071,0.071,0.017)$.
Numerical minorities are made up of one, two, or three components. For both weight vectors, $N_{1}=0, N_{2}=1$. However, $N_{3}=6+4$ in the former case, and $N_{3}=6+2$ in the latter. Thus, the " $I I$ " term is uniformly smaller over all $p<1 / 2$ in the latter case, and we see that for all distributions, the second weight vector is better. When the $N_{l}$ 's are not strictly nested, such a universal comparison becomes impossible, as in the example of Problem 5.8. Hence, $\mathcal{W}$ is only partially ordered.

Unwittingly, we have also shown the following theorem:
Theorem 5.4. For all distributions,

$$
L^{*} \leq \cdots \leq L_{(2 k+1) \mathrm{NN}} \leq L_{(2 k-1) \mathrm{NN}} \leq \cdots \leq L_{3 \mathrm{NN}} \leq L_{\mathrm{NN}} \leq 2 L^{*}
$$

Proof. It suffices once again to look at $\alpha(p)$. Consider the weight vector $w_{1}=$ $\cdots=w_{2 k+1}=1$ (ignoring normalization) as for the $(2 k+1)$-nN rule. The term $I I$ is zero, as $N_{0}=N_{1}=\cdots=N_{k}=0$. However, the $(2 k-1)$-nN rule with vector $w_{1}=\cdots=w_{2 k-1}=1, w_{2 k}=w_{2 k+1}=0$, has a nonzero term $I I$, because $N_{0}=\cdots=N_{k-1}=0$, yet $N_{k}=\binom{2 k-1}{k}>0$. Hence, $L_{(2 k+1) \mathrm{NN}} \leq L_{(2 k-1) \mathrm{NN}}$.

Remark. We have strict inequality $L_{(2 k+1) \mathrm{nN}}<L_{(2 k-1) \mathrm{Nn}}$ whenever $\mathbf{P}\{\eta(X) \notin$ $\{0,1,1 / 2\}\}>0$. When $L^{*}=0$, we have $L_{\mathrm{NN}}=L_{3 \mathrm{NN}}=L_{5 \mathrm{NN}}=\cdots=0$ as well.

## $5.6 k$-Nearest Neighbor Rules: Even $k$

Until now we assumed throughout that $k$ was odd, so that voting ties were avoided. The tie-breaking procedure we follow for the $2 k$-nearest neighbor rule is as follows:

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{2 k} Y_{(i)}(x)>k \\ 0 & \text { if } \sum_{i=1}^{2 k} Y_{(i)}(x)<k \\ Y_{(1)}(x) & \text { if } \sum_{i=1}^{2 k} Y_{(i)}(x)=k\end{cases}
$$

Formally, this is equivalent to a weighted $2 k$-nearest neighbor rule with weight vector $(3,2,2,2, \ldots, 2,2)$. It is easy to check from Theorem 5.3 that this is the asymptotically best weight vector. Even values do not decrease the probability of error. In particular, we have the following:

Theorem 5.5. (DEVIJVER (1978)). For all distributions, and all integers $k$,

$$
L_{(2 k-1) \mathrm{NN}}=L_{(2 k) \mathrm{NN}} .
$$

Proof. Recall that $L_{k N N}$ may be written in the form $L_{k N N}=\mathbf{E}\{\alpha(\eta(X))\}$, where

$$
\alpha(\eta(x))=\lim _{n \rightarrow \infty} \mathbf{P}\left\{g_{n}^{(k)}(X) \neq Y \mid X=x\right\}
$$

is the pointwise asymptotic error probability of the $k$-NN rule $g_{n}^{(k)}$. It is convenient to consider $Z_{1}, \ldots, Z_{2 k}$ i.i.d. $\{-1,1\}$-valued random variables with $\mathbf{P}\left\{Z_{i}=1\right\}=p=$ $\eta(x)$, and to base the decision upon the sign of $\sum_{i=1}^{2 k} Z_{i}$. From the general formula for weighted nearest neighbor rules, the pointwise asymptotic error probability of the $(2 k)$-nN rule is

$$
\begin{aligned}
& \lim _{n \rightarrow \infty} \mathbf{P}\left\{g_{n}^{(2 k)}(X) \neq Y \mid X=x\right\} \\
&= p \mathbf{P}\left\{\sum_{i=1}^{2 k} Z_{i}<0\right\}+p \mathbf{P}\left\{\sum_{i=1}^{2 k} Z_{i}=0, Z_{1}<0\right\} \\
&+(1-p) \mathbf{P}\left\{\sum_{i=1}^{2 k} Z_{i}>0\right\}+(1-p) \mathbf{P}\left\{\sum_{i=1}^{2 k} Z_{i}=0, Z_{1}>0\right\}
\end{aligned}
$$

$$
\begin{aligned}
& =p \mathbf{P}\left\{\sum_{i=2}^{2 k} Z_{i}<0\right\}+(1-p) \mathbf{P}\left\{\sum_{i=2}^{2 k} Z_{i}>0\right\} \\
& =\lim _{n \rightarrow \infty} \mathbf{P}\left\{g_{n}^{(2 k-1)}(X) \neq Y \mid X=x\right\}
\end{aligned}
$$

Therefore, $L_{(2 k) \mathrm{NN}}=L_{(2 k-1) \mathrm{NN}}$.

### 5.7 Inequalities for the Probability of Error

We return to the case when $k$ is odd. Recall that

$$
L_{k \mathrm{NN}}=\mathbf{E}\left\{\alpha_{k}(\eta(X))\right\}
$$

where

$$
\alpha_{k}(p)=\min (p, 1-p)+|2 p-1| \mathbf{P}\left\{\operatorname{Binomial}(k, \min (p, 1-p))>\frac{k}{2}\right\} .
$$

Since $L^{*}=\mathbf{E}\{\min (\eta(X), 1-\eta(X))\}$, we may exploit this representation to obtain a variety of inequalities on $L_{k \mathrm{NN}}-L^{*}$. We begin with one that is very easy to prove but perhaps not the strongest.

Theorem 5.6. For all odd $k$ and all distributions,

$$
L_{k \mathrm{NN}} \leq L^{*}+\frac{1}{\sqrt{k e}}
$$

Proof. By the above representation,

$$
\begin{aligned}
L_{k \mathrm{NN}}-L^{*} \leq & \sup _{0 \leq p \leq 1 / 2}(1-2 p) \mathbf{P}\left\{B>\frac{k}{2}\right\} \\
& (B \text { is Binomial }(k, p)) \\
= & \sup _{0 \leq p \leq 1 / 2}(1-2 p) \mathbf{P}\left\{\frac{B-k p}{k}>\frac{1}{2}-p\right\} \\
\leq & \sup _{0 \leq p \leq 1 / 2}(1-2 p) e^{-2 k(1 / 2-p)^{2}} \\
& (\text { by the Okamoto-Hoeffding inequality-Theorem } 8.1) \\
= & \sup _{0 \leq u \leq 1} u e^{-k u^{2} / 2} \\
= & \frac{1}{\sqrt{k e}} .
\end{aligned}
$$

Theorem 5.7. (GyÖRFI AND GYÖRFI (1978)). For all distributions and all odd $k$,

$$
L_{k \mathrm{NN}} \leq L^{*}+\sqrt{\frac{2 L_{\mathrm{NN}}}{k}}
$$

Proof. We note that for $p \leq 1 / 2$, with $B$ binomial $(k, p)$,

$$
\begin{aligned}
\mathbf{P}\left\{B>\frac{k}{2}\right\} & =\mathbf{P}\left\{B-k p>k\left(\frac{1}{2}-p\right)\right\} \\
& \leq \frac{\mathbf{E}\{|B-k p|\}}{k(1 / 2-p)} \quad \text { (Markov's inequality) } \\
& \leq \frac{\sqrt{\operatorname{Var}\{B\}}}{k(1 / 2-p)} \quad \text { (Cauchy-Schwarz inequality) } \\
& =\frac{2 \sqrt{p(1-p)}}{\sqrt{k}(1-2 p)}
\end{aligned}
$$

Hence,

$$
\begin{aligned}
L_{k \mathrm{NN}}-L^{*} & \leq \mathbf{E}\left\{\frac{2}{\sqrt{k}} \sqrt{\eta(X)(1-\eta(X))}\right\} \\
& \leq \frac{2}{\sqrt{k}} \sqrt{\mathbf{E}\{\eta(X)(1-\eta(X))\}} \quad \text { (Jensen's inequality) } \\
& =\frac{2}{\sqrt{k}} \sqrt{\frac{L_{\mathrm{NN}}}{2}} \\
& =\sqrt{\frac{2 L_{\mathrm{NN}}}{k}} .
\end{aligned}
$$

REMARK. For large $k, B$ is approximately normal $(k, p(1-p)$ ), and thus $\mathbf{E}\{\mid B-$ $k p \mid\} \approx \sqrt{k p(1-p)} \sqrt{2 / \pi}$, as the first absolute moment of a normal random variable is $\sqrt{2 / \pi}$ (see Problem 5.11). Working this through yields an approximate bound of $\sqrt{L_{\mathrm{NN}} /(\pi k)}$. The bound is proportional to $\sqrt{L_{\mathrm{NN}}}$. This can be improved to $L^{*}$ if, instead of bounding it from above by Markov's inequality, we directly approximate $\mathbf{P}\{B-k p>k(1 / 2-p)\}$ as shown below.

Theorem 5.8. (Devroye (1981B)). For all distributions and $k \geq 3$ odd,

$$
L_{k \mathrm{NN}} \leq L^{*}\left(1+\frac{\gamma}{\sqrt{k}}\left(1+O\left(k^{-1 / 6}\right)\right)\right)
$$

where $\gamma=\sup _{r>0} 2 r \mathbf{P}\{N>r\}=0.33994241 \ldots, N$ is normal $(0,1)$, and $O(\cdot)$ refers to $k \rightarrow \infty$. (Explicit constants are given in the proof.)

The constant $\gamma$ in the proof cannot be improved. A slightly weaker bound was obtained by Devijver (1979):

$$
\begin{aligned}
L_{k \mathrm{NN}} & \leq L^{*}+\frac{1}{2^{2 k^{\prime}}}\binom{2 k^{\prime}}{k^{\prime}} L_{\mathrm{NN}} \quad\left(\text { where } k^{\prime}=\lceil k / 2\rceil\right) \\
& =L^{*}+L_{\mathrm{NN}} \sqrt{\frac{2}{\pi k^{\prime}}}(1+o(1)) \quad(\text { as } k \rightarrow \infty, \text { see Lemma A.3 })
\end{aligned}
$$

See also Devijver and Kittler (1982, p.102).
Lemma 5.6. (Devroye (1981B)). For $p \leq 1 / 2$ and with $k>3$ odd,

$$
\begin{aligned}
\mathbf{P}\left\{\operatorname{Binomial}(k, p)>\frac{k}{2}\right\} & =\frac{k!}{\left(\frac{k-1}{2}\right)!\left(\frac{k-1}{2}\right)!} \int_{0}^{p}(x(1-x))^{(k-1) / 2} d x \\
& \leq A \int_{(1-2 p) \sqrt{k-1}}^{\sqrt{k-1}} e^{-z^{2} / 2} d z
\end{aligned}
$$

where $A \leq \frac{1}{\sqrt{2 \pi}}\left(1+\frac{2}{k}+\frac{3}{4 k^{2}}\right)$.

Proof. Consider $k$ i.i.d. uniform random variables on [ 0,1 ]. The number of values in $[0, p]$ is binomial $(k, p)$. The number exceeds $k / 2$ if and only if the $(k+1) / 2$-th order statistic of the uniform cloud is at most $p$. The latter is beta $((k+1) / 2,(k+1) / 2)$ distributed, explaining the first equality (Problem 5.32). Note that we have written a discrete sum as an integral-in some cases, such tricks pay off handsome rewards. To show the inequality, replace $x$ by $\frac{1}{2}\left(1-\frac{z}{\sqrt{k-1}}\right)$ and use the inequality $1-u \leq$ $e^{-u}$ to obtain a bound as shown with

$$
A=\frac{1}{2^{k} \sqrt{k-1}} \times \frac{k!}{\left(\frac{k-1}{2}\right)!\left(\frac{k-1}{2}\right)!} .
$$

Finally,

$$
\begin{aligned}
A & =\mathbf{P}\left\{B=\frac{k+1}{2}\right\} \frac{k+1}{2 \sqrt{k-3}} \quad(B \text { is binomial }(k, 1 / 2)) \\
& \leq \sqrt{\frac{k}{2 \pi \frac{k+1}{2} \frac{k-1}{2}} \frac{k+1}{2 \sqrt{k-1}} \quad(\text { Problem 5.17) }} \\
& =\frac{1}{\sqrt{2 \pi}} \frac{\sqrt{k(k+1)}}{k-1} \\
& \leq \frac{1}{\sqrt{2 \pi}}\left(1+\frac{2}{k}+\frac{3}{4 k^{2}}\right) \quad \text { (Problem 5.18). }
\end{aligned}
$$

Proof of Theorem 5.8. From earlier remarks,

$$
\begin{aligned}
L_{k \mathrm{NN}}-L^{*} & =\mathbf{E}\left\{\alpha_{k}(\eta(X))-\min (\eta(X), 1-\eta(X))\right\} \\
& =\mathbf{E}\left\{\left(\frac{\alpha_{k}(\eta(X))}{\min (\eta(X), 1-\eta(X))}-1\right) \min (\eta(X), 1-\eta(X))\right\} \\
& \leq\left(\sup _{0<p<1 / 2} \frac{1-2 p}{p} \mathbf{P}\left\{B>\frac{k}{2}\right\}\right) L^{*} \quad(B \text { is } \operatorname{binomial}(k, p)) .
\end{aligned}
$$

We merely bound the factor in brackets. Clearly, by Lemma 5.6,

$$
L_{k N N}-L^{*} \leq L^{*}\left(\sup _{0<p<1 / 2} \frac{1-2 p}{p} A \int_{(1-2 p) \sqrt{k-1}}^{\sqrt{k-1}} e^{-z^{2} / 2} d z\right)
$$

Take $a<1$ as the solution of $\left(3 /\left(e a^{2}\right)\right)^{3 / 2} \frac{1}{\sqrt{2 \pi(k-1)}}=\gamma$, which is possible if $k-1>\frac{1}{2 \pi} \frac{1}{\gamma^{2}}\left(\frac{3}{e}\right)^{6}=2.4886858 \ldots$. Setting $v=(1-2 p) \sqrt{k-1}$, we have

$$
\begin{aligned}
\sup _{0<p<1 / 2} & \frac{1-2 p}{p} \int_{(1-2 p) \sqrt{k-1}}^{\sqrt{k-1}} \frac{e^{-z^{2} / 2}}{\sqrt{2 \pi}} d z \\
\leq & \max \left(\sup _{0<v \leq a \sqrt{k-1}} \frac{2 v / \sqrt{k-1}}{1-v / \sqrt{k-1}} \int_{v}^{\infty} \frac{e^{-z^{2} / 2}}{\sqrt{2 \pi}} d z\right. \\
& \sup _{\left.a \sqrt{k-1 \leq v<\sqrt{k-1}} \frac{2 p \sqrt{k-1}}{1-v / \sqrt{k-1}} \frac{e^{-v^{2} / 2}}{\sqrt{2 \pi}}\right)}^{\leq} \quad \max \left(\frac{\gamma}{(1-a) \sqrt{k-1}}, \frac{\sqrt{k-1}}{\sqrt{2 \pi}} e^{-a^{2}(k-1) / 2}\right) \\
\leq & \max \left(\frac{\gamma}{(1-a) \sqrt{k-1}},\left(\frac{3}{e a^{2}}\right)^{3 / 2} \frac{1}{(k-1) \sqrt{2 \pi}}\right) \\
& \left(\operatorname{use} u^{3 / 2} e^{-c u} \leq(3 /(2 c e))^{3 / 2} \text { for all } u>0\right) \\
= & \frac{\gamma}{(1-a) \sqrt{k-1}} .
\end{aligned}
$$

Collect all bounds and note that $a=O\left(k^{-1 / 6}\right)$.

### 5.8 Behavior When $L^{*}$ Is Small

In this section, we look more closely at $L_{k N N}$ when $L^{*}$ is small. Recalling that $L_{k \mathrm{NN}}=\mathbf{E}\left\{\alpha_{k}(\eta(X))\right\}$ with
$\alpha_{k}(p)=\min (p, 1-p)+|1-2 \min (p, 1-p)| \mathbf{P}\left\{\operatorname{Binomial}(k, \min (p, 1-p))>\frac{k}{2}\right\}$
for odd $k$, it is easily seen that $L_{k N N}=\mathbf{E}\left\{\xi_{k}(\min (\eta(X), 1-\eta(X)))\right\}$ for some function $\xi_{k}$. Because

$$
\min (p, 1-p)=\frac{1-\sqrt{1-4 p(1-p)}}{2}
$$

we also have $L_{k_{\mathrm{NN}}}=\mathbf{E}\left\{\psi_{k}(\eta(X)(1-\eta(X)))\right\}$ for some other function $\psi_{k}$. Workedout forms of $L_{k_{\mathrm{NN}}}$ include

$$
\begin{aligned}
L_{k \mathrm{NN}}= & \mathbf{E}\left\{\sum_{j<k / 2}\binom{k}{j} \eta(X)^{j+1}(1-\eta(X))^{k-j}\right. \\
& \left.+\sum_{j>k / 2}\binom{k}{j} \eta(X)^{j}(1-\eta(X))^{k-j+1}\right\} \\
= & \sum_{j<k / 2}\binom{k}{j} \mathbf{E}\left\{(\eta(X)(1-\eta(X)))^{j+1}\left((1-\eta(X))^{k-2 j-1}+\eta(X)^{k-2 j-1}\right)\right\}
\end{aligned}
$$

As $p^{a}+(1-p)^{a}$ is a function of $p(1-p)$ for integer $a$, this may be further reduced to simplified forms such as

$$
\begin{aligned}
L_{\mathrm{NN}}= & \mathbf{E}\{2 \eta(X)(1-\eta(X))\}, \\
L_{3 \mathrm{NN}}= & \mathbf{E}\{\eta(X)(1-\eta(X))\}+4 \mathbf{E}\left\{(\eta(X)(1-\eta(X)))^{2}\right\}, \\
L_{5_{\mathrm{NN}}}= & \mathbf{E}\{\eta(X)(1-\eta(X))\}+\mathbf{E}\left\{(\eta(X)(1-\eta(X)))^{2}\right\} \\
& +12 \mathbf{E}\left\{(\eta(X)(1-\eta(X)))^{3}\right\} .
\end{aligned}
$$

The behavior of $\alpha_{k}$ near zero is very informative. As $p \downarrow 0$, we have

$$
\begin{aligned}
& \alpha_{1}(p)=2 p(1-p) \sim 2 p, \\
& \alpha_{3}(p)=p(1-p)(1+4 p) \sim p+3 p^{2}, \\
& \alpha_{5}(p) \sim p+10 p^{3},
\end{aligned}
$$

while for the Bayes error, $L^{*}=\mathbf{E}\{\min (\eta(X), 1-\eta(X))\}=\mathbf{E}\left\{\alpha_{\infty}(\eta(X))\right\}$, where $\alpha_{\infty}=\min (p, 1-p) \sim p$ as $p \downarrow 0$. Assume that $\eta(x)=p$ at all $x$. Then, as $p \downarrow 0$,

$$
L_{\mathrm{NN}} \sim 2 L^{*} \text { and } L_{3 \mathrm{NN}} \sim L^{*} .
$$

Moreover, $L_{\mathrm{NN}}-L^{*} \sim L^{*}, L_{3 \mathrm{NN}}-L^{*} \sim 3 L^{* 2}$. Assume that $L^{*}=p=0.01$. Then $L_{1} \approx 0.02$, whereas $L_{3 \mathrm{NN}}-L^{*} \approx 0.0003$. For all practical purposes, the 3-NN rule is virtually perfect. For this reason, the $3-$ nn rule is highly recommended. Little is gained by considering the $5-\mathrm{NN}$ rule when $p$ is small, as $L_{5_{\mathrm{NN}}}-L^{*} \approx 0.00001$.

Let $a_{k}$ be the smallest number such that $\alpha_{k}(p) \leq a_{k} \min (p, 1-p)$ for all $p$ (the tangents in Figure 5.6). Then

$$
\begin{aligned}
L_{k \mathrm{~N} \mathrm{~N}}=\mathbf{E}\left\{\alpha_{k}(\eta(X))\right\} & \leq a_{k} \mathbf{E}\{\min (\eta(X),(1-\eta(X))\} \\
& =a_{k} L^{*} .
\end{aligned}
$$

This is precisely at the basis of the inequalities of Theorems 5.6 through 5.8, where it was shown that $a_{k}=1+O(1 / \sqrt{k})$.

Figure 5.6. $\alpha_{k}(p)$ as a function


### 5.9 Nearest Neighbor Rules When $L^{*}=0$

From Theorem 5.4 we retain that if $L^{*}=0$, then $L_{k \mathrm{NN}}=0$ for all $k$. In fact, then, for every fixed $k$, the $k$-nearest neighbor rule is consistent. Cover has a beautiful example to illustrate this remarkable fact. $L^{*}=0$ implies that $\eta(x) \in\{0,1\}$ for all $x$, and thus, the classes are separated. This does not imply that the support of $X$ given $Y=0$ is different from the support of $X$ given $Y=1$. Take for example a random rational number from $[0,1]$ (e.g., generate $I, J$ independently and at random from the geometric distribution on $\{1,2,3, \ldots\}$, and set $X=\min (I, J) / \max (I, J)$ ). Every rational number on $[0,1]$ has positive probability. Given $Y=1, X$ is as above, and given $Y=0, X$ is uniform on $[0,1]$. Let $\mathbf{P}\{Y=1\}=\mathbf{P}\{Y=0\}=1 / 2$. The support of $X$ is identical in both cases. As

$$
\eta(x)= \begin{cases}1 & \text { if } x \text { is rational } \\ 0 & \text { if } x \text { is irrational }\end{cases}
$$

we see that $L^{*}=0$ and that the nearest neighbor rule is consistent. If someone shows us a number $X$ drawn from the same distribution as the data, then we may decide the rationality of $X$ merely by looking at the rationality of the nearest neighbor of $X$. Although we did not show this, the same is true if we are given any $x \in[0,1]$ :

$$
\begin{aligned}
\lim _{n \rightarrow \infty} \mathbf{P}\{x & \text { is rational } \left.Y_{(1)}(x)=0\left(X_{(1)}(x) \text { is not rational }\right)\right\} \\
& =\lim _{n \rightarrow \infty} \mathbf{P}\left\{x \text { is not rational } Y_{(1)}(x)=1\left(X_{(1)}(x) \text { is rational) }\right\}\right. \\
& =0 \quad \text { (see Problem 5.38). }
\end{aligned}
$$

### 5.10 Admissibility of the Nearest Neighbor Rule

The consistency theorems of Chapter 11 show us that we should take $k=k_{n} \rightarrow \infty$ in the $k-\mathrm{NN}$ rule. The decreasing nature of $L_{k \mathrm{NN}}$ corroborates this. Yet, there exist distributions for which for all $n$, the $1-\mathrm{NN}$ rule is better than the $k$-NN rule for any $k \geq 3$. This observation, due to Cover and Hart (1967), rests on the following class of examples. Let $S_{0}$ and $S_{1}$ be two spheres of radius 1 centered at $a$ and $b$, where $\|a-b\|>2$. Given $Y=1, X$ is uniform on $S_{1}$, while given $Y=0, X$ is uniform on $S_{0}$, whereas $\mathbf{P}\{Y=1\}=\mathbf{P}\{Y=0\}=1 / 2$. We note that given $n$ observations, with the $1-\mathrm{NN}$ rule,

$$
\mathbf{E}\left\{L_{n}\right\}=\mathbf{P}\left\{Y=0, Y_{1}=\cdots=Y_{n}=1\right\}+\mathbf{P}\left\{Y=1, Y_{1}=\cdots=Y_{n}=0\right\}=\frac{1}{2^{n}}
$$

For the $k$-nN rule, $k$ being odd, we have

$$
\begin{aligned}
\mathbf{E}\left\{L_{n}\right\}= & \mathbf{P}\left\{Y=0, \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} \leq\lfloor k / 2\rfloor\right\} \\
& +\mathbf{P}\left\{Y=1, \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} \leq\lfloor k / 2\rfloor\right\} \\
= & \mathbf{P}\{\operatorname{Binomial}(n, 1 / 2) \leq\lfloor k / 2\rfloor\} \\
= & \frac{1}{2^{n}} \sum_{j=0}^{\lfloor k / 2\rfloor}\binom{n}{j}>\frac{1}{2^{n}} \quad \text { when } k \geq 3 .
\end{aligned}
$$

Hence, the $k$-NN rule is worse than the $1-\mathrm{NN}$ rule for every $n$ when the distribution is given above. We refer to the exercises regarding some interesting admissibility questions for $k-\mathrm{NN}$ rules.

### 5.11 The ( $k, l$ )-Nearest Neighbor Rule

In 1970, Hellman (1970) proposed the $(k, l)$-nearest neighbor rule, which is identical to the $k$-nearest neighbor rule, but refuses to make a decision unless at least $l>k / 2$ observations are from the same class. Formally, we set

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{k} Y_{(i)}(x) \geq l \\ 0 & \text { if } \sum_{i=1}^{k} Y_{(i)}(x) \leq k-l \\ -1 & \text { otherwise (no decision) }\end{cases}
$$

Define the pseudoprobability of error by

$$
L_{n}=\mathbf{P}\left\{g_{n}(X) \notin\{-1, Y\} \mid D_{n}\right\},
$$

that is, the probability that we reach a decision and correctly classify $X$. Clearly, $L_{n} \leq \mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}$, our standard probability of error. The latter inequality
is only superficially interesting, as the probability of not reaching a decision is not taken into account in $L_{n}$. We may extend Theorem 5.2 to show the following (Problem 5.35):

Theorem 5.9. For the ( $k, l$ )-nearest neighbor rule, the pseudoprobability of error $L_{n}$ satisfies

$$
\begin{aligned}
\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}= & \mathbf{E}\{\eta(X) \mathbf{P}\{\operatorname{Binomial}(k, \eta(X)) \leq k-l \mid X\} \\
& +(1-\eta(X)) \mathbf{P}\{\operatorname{Binomial}(k, \eta(X)) \geq l \mid X\}\} \\
\stackrel{\text { def }}{=} & L_{k, l} .
\end{aligned}
$$

The above result is distribution-free. Note that the $k$-nearest neighbor rule for odd $k$ corresponds to $L_{k,(k+1) / 2}$. The limit $L_{k, l}$ by itself is not interesting, but it was shown by Devijver (1979) that $L_{k, l}$ holds information regarding the Bayes error $L^{*}$.

Theorem 5.10. (DEVIJVER (1979)). For all distributions and with $k$ odd,

$$
L_{k, k} \leq L_{k, k-1} \leq \cdots \leq L_{k,\lceil k / 2]+1} \leq L^{*} \leq L_{k,[k / 2]}=L_{k \mathrm{NN}} .
$$

Also,

$$
\frac{L_{k N \mathrm{~N}}+L_{k,\lceil k / 2\rceil+1}}{2} \leq L^{*} \leq L_{k \mathrm{NN}} .
$$

This theorem (for which we refer to Problem 5.34) shows that $L^{*}$ is tightly sandwiched between $L_{k \mathrm{NN}}$, the asymptotic probability of error of the $k$-nearest neighbor rule, and the "tennis" rule which requires that the difference of votes between the two classes among the $k$ nearest neighbors be at least two. If $L_{n}$ is close to its limit, and if we can estimate $L_{n}$ (see the chapters on error estimation), then we may be able to use Devijver's inequalities to obtain estimates of the Bayes error $L^{*}$. For additional results, see Loizou and Maybank (1987).

As a corollary of Devijver's inequalities, we note that

$$
L_{k \mathrm{NN}}-L^{*} \leq \frac{L_{k \mathrm{NN}}-L_{k,\lceil k / 2\rceil+1}}{2}
$$

We have

$$
\begin{aligned}
L_{k, l}= & L^{*}+\mathbf{E}\{|1-2 \min (\eta(X), 1-\eta(X))| \\
& \times \mathbf{P}\{\operatorname{Binomial}(k, \min (\eta(X), 1-\eta(X))) \geq l \mid X\}\},
\end{aligned}
$$

and therefore

$$
\begin{aligned}
L_{k, l}- & L_{k, l+1} \\
= & \mathbf{E}\{|1-2 \min (\eta(X), 1-\eta(X))| \\
& \times \mathbf{P}\{\operatorname{Binomial}(k, \min (\eta(X), 1-\eta(X)))=l \mid X\}\}
\end{aligned}
$$

$$
\begin{aligned}
= & \mathbf{E}\{|1-2 \min (\eta(X), 1-\eta(X))| \\
& \left.\times\binom{ k}{l} \min (\eta(X), 1-\eta(X))^{l}(1-\min (\eta(X), 1-\eta(X)))^{k-l}\right\} \\
\leq & \binom{k}{l} \frac{l^{l}(k-l)^{k-l}}{k^{k}} \\
& \left(\text { because } u^{l}(1-u)^{k-l} \text { reaches its maximum on [0, 1] at } u=l / k\right) \\
\leq & \sqrt{\frac{k}{12 \pi l(k-l)}} \\
& \text { (use }\binom{k}{l} \leq \frac{k^{k}}{l^{k}(k-l)^{k-l}} \frac{1}{\sqrt{2 \pi}} \sqrt{\frac{k}{l(k-l)}}, \text { by Stirling's formula). }
\end{aligned}
$$

With $l=\lceil k / 2\rceil$, we thus obtain

$$
\begin{aligned}
L_{k \mathrm{NN}}-L^{*} & \leq \frac{1}{2 \sqrt{2 \pi}} \sqrt{\frac{k}{\lceil k / 2\rceil\lfloor k / 2\rfloor}} \\
& =\sqrt{\frac{k}{2 \pi\left(k^{2}-1\right)} \approx \frac{0.398942}{\sqrt{k}}}
\end{aligned}
$$

improving on Theorem 5.6. Various other inequalities may be derived in this manner as well.

## Problems and Exercises

Problem 5.1. Let $\|\cdot\|$ be an arbitrary norm on $\mathcal{R}^{d}$, and define the $k$-nearest neighbor rule in terms of the distance $\rho(x, z)=\|x-z\|$. Show that Theorems 5.1 and 5.2 remain valid. Hint: Only Stone's lemma needs adjusting. The role of cones $C(x, \pi / 6)$ used in the proof are now played by sets with the following property: $x$ and $z$ belong to the same set if and only if

$$
\left\|\frac{x}{\|x\|}-\frac{z}{\|z\|}\right\|<1
$$

Problem 5.2. Does there exist a distribution for which $\sup _{n \geq 1} \mathbf{E}\left\{L_{n}\right\}>1 / 2$ for the nearest neighbor rule?

Problem 5.3. Show that $L_{3 \mathrm{NN}} \leq 1.32 L^{*}$ and that $L_{5 \mathrm{NN}} \leq 1.22 L^{*}$.
Problem 5.4. Show that if $C^{*}$ is a compact subset of $\mathcal{R}^{d}$ and $C$ is the support set for the probability measure $\mu$,

$$
\sup _{x \in C \cap C^{*}}\left\|X_{(1)}-x\right\| \rightarrow 0
$$

with probability one, where $X_{(1)}$ is the nearest neighbor of $x$ among $X_{1}, \ldots, X_{n}$ (Wagner (1971)).

Problem 5.5. Let $\mu$ be the probability measure of $X$ given $Y=0$, and let $\nu$ be the probability measure of $X$ given $Y=1$. Assume that $X$ is real-valued and that $\mathbf{P}\{Y=0\}=\mathbf{P}\{Y=1\}=$ $1 / 2$. Find a pair $(\nu, \mu)$ such that
(1) $\operatorname{support}(\mu)=\operatorname{support}(\nu)$;
(2) $L^{*}=0$.

Conclude that $L^{*}=0$ does not tell us a lot about the support sets of $\mu$ and $\nu$.
Problem 5.6. Consider the $(2 k+1)$-nearest neighbor rule for distributions ( $X, Y$ ) with $\eta(x) \equiv p$ constant, and $Y$ independent of $X$. This exercise explores the behavior of $L_{(2 k+1) \mathrm{NN}}$ as $p \downarrow 0$.
(1) For fixed integer $l>0$, as $p \downarrow 0$, show that

$$
\mathbf{P}\{\operatorname{Binomial}(2 k, p) \geq l\} \sim\binom{2 k}{l} p^{l}
$$

(2) Use a convenient representation of $L_{(2 k+1) \mathrm{NN}}$ to conclude that as $p \downarrow 0$,

$$
L_{(2 k+1) N N}=p+\left(\binom{2 k}{k}+\binom{2 k}{k+1}\right) p^{k+1}+o\left(p^{k+1}\right) .
$$

Problem 5.7. Das Gupta and Lin (1980) proposed the following rule for data with $X \in \mathcal{R}$. Assume $X$ is nonatomic. First, reorder $X_{1}, \ldots, X_{n}, X$ according to increasing values, and denote the ordered set by $X_{(1)}<X_{(2)}<\cdots<X_{(i)}<X<X_{(i+1)}<\cdots<X_{(n)}$. The $Y_{i}$ 's are permuted so that $Y_{(j)}$ is the label of $X_{(j)}$. Take votes among $\left\{Y_{(i)}, Y_{(i+1)}\right\},\left\{Y_{(i-1)}, Y_{(i+2)}\right\}, \ldots$ until for the first time there is agreement $\left(Y_{(i-j)}=Y_{(i+j+1)}\right)$, at which time we decide that class, that is, $g_{n}(X)=Y_{(i-j)}=Y_{(i+j+1)}$. This rule is invariant under monotone transformations of the $x$-axis.
(1) If $L$ denotes the asymptotic expected probability of error, show that for all nonatomic $X$,

$$
L=\mathbf{E}\left\{\frac{\eta(X)(1-\eta(X)}{1-2 \eta(X)(1-\eta(X)}\right\} .
$$

(2) Show that $L$ is the same as for the rule in which $X \in \mathcal{R}^{d}$ and we consider 2-NN, $4-\mathrm{NN}, 6-\mathrm{NN}$, etc. rules in turn, stopping at the first $2 k$-nN rule for which there is no voting tie. Assume for simplicity that $X$ has a density (with a good distance-tie breaking rule, this may be dropped).
(3) Show that $L-L_{\mathrm{NN}} \geq(1 / 2)\left(L_{3 \mathrm{NN}}-L_{\mathrm{NN}}\right)$, and thus that $L \geq\left(L_{\mathrm{NN}}+L_{3 \mathrm{NN}}\right) / 2$.
(4) Show that $L \leq L_{\mathrm{NN}}$. Hence, the rule performs somewhere in between the $1-\mathrm{NN}$ and 3 -nn rules.

Problem 5.8. Let $Y$ be independent of $X$, and $\eta(x) \equiv p$ constant. Consider a weighted $(2 k+1)$-nearest neighbor rule with weights $(2 m+1,1,1, \ldots, 1)$ (there are $2 k$ "ones"), where $k-1 \geq m \geq 0$. For $m=0$, we obtain the $(2 k+1)$-nN rule. Let $L(k, m)$ be the asymptotic probability of error.
(1) Using results from Problem 5.6, show that

$$
L(k, m)=p+\binom{2 k}{k-m} p^{k-m+1}+\binom{2 k}{k+m+1} p^{k+m+1}+o\left(p^{k-m+1}\right)
$$

as $p \downarrow 0$. Conclude that within this class of rules, for small $p$, the goodness of a rule is measured by $k-m$.
(2) Let $\delta>0$ be small, and set $p=1 / 2-\delta$. Show that if $X$ is binomial $(2 k, 1 / 2-\delta)$ and $Z$ is $\operatorname{binomial}(2 k, 1 / 2)$, then for fixed $l$, as $\delta \downarrow 0$,

$$
\mathbf{P}\{X \geq l\}=\mathbf{P}\{Z \geq l\}-2 k \delta \mathbf{P}\{Z=l\}+o\left(\delta^{2}\right)
$$

and

$$
\mathbf{P}\{X \leq l\}=\mathbf{P}\{Z \leq l\}+2 k \delta \mathbf{P}\{Z=l+1\}+o\left(\delta^{2}\right) .
$$

(3) Conclude that for fixed $k, m$ as $\delta \downarrow 0$,

$$
\begin{aligned}
L(k, m)= & \frac{1}{2}-2 \delta^{2}(k \mathbf{P}\{Z=k+m\}+k \mathbf{P}\{Z=k+m+1\} \\
& +\mathbf{P}\{Z \leq k+m\}-\mathbf{P}\{Z \geq k+m+1\})+o\left(\delta^{2}\right)
\end{aligned}
$$

(4) Take weight vector $w$ with $k$ fixed and $m=\lfloor 10 \sqrt{k}\rfloor$, and compare it with weight vector $w^{\prime}$ with $k / 2$ components and $m=\lfloor\sqrt{k / 2}\rfloor$ as $p \downarrow 0$ and $p \uparrow 1 / 2$. Assume that $k$ is very large but fixed. In particular, show that $w$ is better as $p \downarrow 0$, and $w^{\prime}$ is better as $p \uparrow 1 / 2$. For the last example, note that for fixed $c>0$,

$$
k \mathbf{P}\{Z=k+m\}+k \mathbf{P}\{Z=k+m+1\} \sim 8 \sqrt{k} \frac{1}{\sqrt{2 \pi}} e^{-2 c^{2}} \text { as } k \rightarrow \infty
$$

by the central limit theorem.
(5) Conclude that there exist different weight vectors $w, w^{\prime}$ for which there exists a pair of distributions of $(X, Y)$ such that their asymptotic error probabilities are differently ordered. Thus, $\mathcal{W}$ is not totally ordered with respect to the probability of error.

Problem 5.9. Patrick and Fisher (1970) find the $k$-th nearest neighbor in each of the two classes and classify according to which is nearest. Show that their rule is equivalent to a ( $2 k-1$ )-nearest neighbor rule.

Problem 5.10. Rabiner et al. (1979) generalize the rule of Problem 5.9 so as to classify according to the average distance to the $k$-th nearest neighbor within each class. Assume that $X$ has a density. For fixed $k$, find the asymptotic probability of error.

Problem 5.11. If $N$ is normal $(0,1)$, then $\mathbf{E}\{|N|\}=\sqrt{2 / \pi}$. Prove this.
Problem 5.12. Show that if $L_{9 \mathrm{NN}}=L_{(11) \mathrm{NN}}$, then $L_{(99)_{\mathrm{NN}}}=L_{(111) \mathrm{NN}}$.
Problem 5.13. Show that

$$
L_{3 \mathrm{NN}} \leq\left(\frac{7 \sqrt{7}+17}{27 \sqrt{3}}+1\right) L^{*} \approx 1.3155 \ldots L^{*}
$$

for all distributions (Devroye (1981b)). Hint: Find the smallest constant $a$ such that $L_{\text {3NN }} \leq$ $L^{*}(1+a)$ using the representation of $L_{3 \mathrm{NN}}$ in terms of the binomial tail.

Problem 5.14. Show that if $X$ has a density $f$, then for all $u>0$,

$$
\lim _{n \rightarrow \infty} \mathbf{P}\left\{n^{1 / d}\left\|X_{(1)}(X)-X\right\|>u \mid X\right\}=e^{-f(X) v u^{d}}
$$

with probability one, where $v=\int_{S_{0,1}} d x$ is the volume of the unit ball in $\mathcal{R}^{d}$ (Györfi (1978)).

Problem 5.15. Consider a rule that takes a majority vote over all $Y_{i}$ 's for which $\left\|X_{i}-x\right\| \leq$ $(c / v n)^{1 / d}$, where $v=\int_{S_{0,1}} d x$ is the volume of the unit ball, and $c>0$ is fixed. In case of a tie decide $g_{n}(x)=0$.
(1) If $X$ has a density $f$, show that $\liminf _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\} \geq \mathbf{E}\left\{\eta(X) e^{-c f(X)}\right\}$. Hint: Use the obvious inequality $\mathbf{E}\left\{L_{n}\right\} \geq \mathbf{P}\left\{Y=1, \mu_{n}\left(S_{X, c / v n}\right)=0\right\}$.
(2) If $Y$ is independent of $X$ and $\eta \equiv p>1 / 2$, then

$$
\frac{\mathbf{E}\left\{\eta(X) e^{-c f(X)}\right\}}{L^{*}}=\mathbf{E}\left\{e^{-c f(X)}\right\} \frac{p}{1-p} \uparrow \infty
$$

as $p \uparrow 1$. Show this.
(3) Conclude that

$$
\sup _{(X, Y): L^{*}>0} \frac{\liminf _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}}{L^{*}}=\infty,
$$

and thus that distribution-free bounds of the form $\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\} \leq c^{\prime} L^{*}$ obtained for $k$-nearest neighbor estimates do not exist for these simple rules (Devroye (1981a)).

Problem 5.16. Take an example with $\eta(X) \equiv 1 / 2-1 /(2 \sqrt{k})$, and show that the bound $L_{k \mathrm{NN}}-L^{*} \leq 1 / \sqrt{k e}$ cannot be essentially bettered for large values of $k$, that is, there exists a sequence of distributions (indexed by $k$ ) for which

$$
L_{k \mathrm{NN}}-L^{*} \geq \frac{1-o(1)}{\sqrt{k}} \mathbf{P}\{N \geq 1\}
$$

as $k \rightarrow \infty$, where $N$ is a normal $(0,1)$ random variable.
Problem 5.17. If $B$ is $\operatorname{binomial}(n, p)$, then

$$
\sup _{p} \mathbf{P}\{B=i\} \leq \sqrt{\frac{n}{2 \pi i(n-i)}}, \quad 0<i<n .
$$

Problem 5.18. Show that for $k \geq 3$,

$$
\frac{\sqrt{k(k+1)}}{k-1} \leq\left(1+\frac{1}{2 k}\right)\left(1+\frac{3}{2 k}\right)=1+\frac{2}{k}+\frac{3}{4 k^{2}} .
$$

Problem 5.19. Show that there exists a sequence of distributions of ( $X, Y$ ) (indexed by $k$ ) in which $Y$ is independent of $X$ and $\eta(x) \equiv p$ (with $p$ depending on $k$ only) such that

$$
\liminf _{n \rightarrow \infty}\left(\frac{L_{k \mathrm{NN}}-L^{*}}{L^{*}}\right) \sqrt{k} \geq \gamma=0.339942 \ldots,
$$

where $\gamma$ is the constant of Theorem 5.8 (Devroye (1981b)). Hint: Verify the proof of Theorem 5.8 but bound things from below. Slud's inequality (see Lemma A. 6 in the Appendix) may be of use here.

Problem 5.20. Consider a weighted nearest neighbor rule with weights $1, \rho, \rho^{2}, \rho^{3}, \ldots$ for $\rho<1$. Show that the expected probability of error tends for all distributions to a limit $L(\rho)$. Hint: Truncate at $k$ fixed but large, and argue that the tail has asymptotically negligible weight.

Problem 5.21. Continued. With $L(\rho)$ as in the previous exercise, show that $L(\rho)=L_{\mathrm{NN}}$ whenever $\rho<1 / 2$.

Problem 5.22. Continued. Prove or disprove: as $\rho$ increases from $1 / 2$ to $1, L(\rho)$ decreases monotonically from $L_{\mathrm{NN}}$ to $L^{*}$. (This question is difficult.)

Problem 5.23. Show that in the weighted nN rule with weights ( $1, \rho, \rho^{2}$ ), $0<\rho<1$, the asymptotic probability of error is $L_{\mathrm{NN}}$ if $\rho<(\sqrt{5}-1) / 2$ and is $L_{3 \mathrm{NN}}$ if $\rho>(\sqrt{5}-1) / 2$.

Problem 5.24. Is there any $k$, other than one, for which the $k$-nN rule is admissible, that is, for which there exists a distribution of ( $X, Y$ ) such that $\mathbf{E}\left\{L_{n}\right\}$ for the $k$-NN rule is smaller than $\mathbf{E}\left\{L_{n}\right\}$ for any $k^{\prime}$-NN rule with $k^{\prime} \neq k$, for all $n$ ? HinT: This is difficult. Note that if this is to hold for all $n$, then it must hold for the limits. From this, deduce that with probability one, $\eta(x) \in\{0,1 / 2,1\}$ for any such distribution.

Problem 5.25. For every fixed $n$ and odd $k$ with $n>1000 k$, find a distribution of ( $X, Y$ ) such that $\mathbf{E}\left\{L_{n}\right\}$ for the $k$-NN rule is smaller than $\mathbf{E}\left\{L_{n}\right\}$ for any $k^{\prime}$-nN rule with $k^{\prime} \neq k, k^{\prime}$ odd. Thus, for a given $n$, no $k$ can be a priori discarded from consideration.

Problem 5.26. Let $X$ be uniform on $[0,1], \eta(x) \equiv x$, and $\mathbf{P}\{Y=0\}=\mathbf{P}\{Y=1\}=1 / 2$. Show that for the nearest neighbor rule,

$$
\mathbf{E}\left\{L_{n}\right\}=\frac{1}{3}+\frac{3 n+5}{2(n+1)(n+2)(n+3)}
$$

(Cover and Hart (1967); Peterson (1970)).
Problem 5.27. For the nearest neighbor rule, if $X$ has a density, then

$$
\left|\mathbf{E}\left\{L_{n}\right\}-\mathbf{E}\left\{L_{n+1}\right\}\right| \leq \frac{1}{n+1}
$$

(Cover (1968a)).
Problem 5.28. Let $X$ have a density $f \geq c>0$ on [0, 1], and assume that $f_{0}^{\prime \prime \prime}$ and $f_{1}^{\prime \prime \prime}$ exist and are uniformly bounded. Show that for the nearest neighbor rule, $\mathbf{E}\left\{L_{n}\right\}=L_{\mathrm{NN}}+O\left(1 / n^{2}\right)$ (Cover (1968a)). For $d$-dimensional problems this result was generalized by Psaltis, Snapp, and Venkatesh (1994).

Problem 5.29. Show that $L_{k N \mathrm{~N}} \leq(1+\sqrt{2 / k}) L^{*}$ is the best possible bound of the form $L_{k N N} \leq(1+a / \sqrt{k}) L^{*}$ valid simultaneously for all $k \geq 1$ (Devroye (1981b)).

Problem 5.30. Show that $L_{k N \mathrm{~N}} \leq(1+\sqrt{1 / k}) L^{*}$ for all $k \geq 3$ (Devroye (1981b)).
Problem 5.31. Let $x=(x(1), x(2)) \in \mathcal{R}^{2}$. Consider the nearest neighbor rule based upon vectors with components $\left(x^{3}(1), x^{7}(2), x(1) x(2)\right)$. Show that this is asymptotically not better than if we had used $(x(1), x(2))$. Show by example that $\left(x^{2}(1), x^{3}(2), x^{6}(1) x(2)\right)$ may yield a worse asymptotic error probability than ( $x(1), x(2)$ ).

Problem 5.32. Uniform order statistics. Let $U_{(1)}<\cdots<U_{(n)}$ be order statistics of $n$ i.i.d. uniform $[0,1]$ random variables. Show the following:
(1) $U_{(k)}$ is beta $(k, n+1-k)$, that is, $U_{(k)}$ has density

$$
f(x)=\frac{n!}{(k-1)!(n-k)!} x^{k-1}(1-x)^{n-k}, \quad 0 \leq x \leq 1 .
$$

(2)

$$
\mathbf{E}\left\{U_{(k)}^{a}\right\}=\frac{\Gamma(k+a) \Gamma(n+1)}{\Gamma(k) \Gamma(n+1+a)}, \quad \text { for any } a>0 .
$$

$$
\begin{equation*}
1-\frac{a}{n} \leq \frac{\mathbf{E}\left\{U_{(k)}^{a}\right\}}{(k / n)^{a}} \leq 1+\frac{\psi(a)}{k} \tag{3}
\end{equation*}
$$

for $a \geq 1$, where $\psi(a)$ is a function of $a$ only (Royall (1966)).
Problem 5.33. Dudani's rule. Dudani (1976) proposes a weighted $k$-Nn rule where $Y_{(i)}(x)$ receives weight

$$
\left\|X_{(k)}(x)-x\right\|-\left\|X_{(i)}(x)-x\right\|, \quad 1 \leq i \leq k .
$$

Why is this roughly speaking equivalent to attaching weight $1-(i / k)^{1 / d}$ to the $i$-th nearest neighbor if $X$ has a density? Hint: If $\mu$ is the probability measure of $X$, then

$$
\mu\left(S_{x,\left\|X_{(1)}(x)-x\right\|}\right), \ldots, \mu\left(S_{x,\left\|X_{(k)}(x)-x\right\|}\right)
$$

are distributed like $U_{(1)}, \ldots, U_{(k)}$ where $U_{(1)}<\cdots<U_{(n)}$ are the order statistics of $n$ i.i.d. uniform $[0,1]$ random variables. Replace $\mu$ by a good local approximation, and use results from the previous exercise.

Problem 5.34. Show Devijver's theorem (Theorem 5.10) in two parts: first establish the inequality $L^{*} \geq L_{k,\lceil k / 2\rceil-1}$ for the tennis rule, and then establish the monotonicity.

Problem 5.35. Show Theorem 5.9 for the $(k, l)$ nearest neighbor rule.
Problem 5.36. Let $R$ be the asymptotic error probability of the (2,2)-nearest neighbor rule. Prove that $R=\mathbf{E}\{2 \eta(X)(1-\eta(X))\}=L_{\mathrm{NN}}$.

Problem 5.37. For the nearest neighbor rule, show that for all distributions,

$$
\begin{aligned}
\lim _{n \rightarrow \infty} \mathbf{P}\left\{g_{n}(X)=0, Y=1\right\} & =\lim _{n \rightarrow \infty} \mathbf{P}\left\{g_{n}(X)=1, Y=0\right\} \\
& =\mathbf{E}\{\eta(X)(1-\eta(X))\}
\end{aligned}
$$

(Devijver and Kittler (1982)). Thus, errors of both kinds are equally likely.
Problem 5.38. Let $\mathbf{P}\{Y=1\}=\mathbf{P}\{Y=0\}=1 / 2$ and let $X$ be a random rational if $Y=1$ (as defined in Section 5.9) such that every rational number has positive probability, and let $X$ be uniform [0, 1] if $Y=0$. Show that for every $x \in[0,1]$ not rational, $\mathbf{P}\left\{Y_{(1)}(x)=1\right\} \rightarrow 0$ as $n \rightarrow \infty$, while for every $x \in[0,1]$ rational, $\mathbf{P}\left\{Y_{(1)}(x)=0\right\} \rightarrow 0$ as $n \rightarrow \infty$.

Problem 5.39. Let $X_{1}, \ldots, X_{n}$ be i.i.d. and have a common density. Show that for fixed $k>0$,
$n \mathbf{P}\left\{X_{3}\right.$ is among the $k$ nearest neighbors of $X_{1}$ and $X_{2}$ in $\left.\left\{X_{3}, \ldots, X_{n}\right\}\right\} \rightarrow 0$.
Show that the same result remains valid whenever $k$ varies with $n$ such that $k / \sqrt{n} \rightarrow 0$.

Problem 5.40. Imperfect training. Let $(X, Y, Z),\left(X_{1}, Y_{1}, Z_{1}\right), \ldots,\left(X_{n}, Y_{n}, Z_{n}\right)$ be a sequence of i.i.d. triples in $\mathcal{R}^{d} \times\{0,1\} \times\{0,1\}$ with $\mathbf{P}\{Y=1 \mid X=x\}=\eta(x)$ and $\mathbf{P}\{Z=$ $1 \mid X=x\}=\eta^{\prime}(x)$. Let $Z_{(1)}(X)$ be $Z_{j}$ if $X_{j}$ is the nearest neighbor of $X$ among $X_{1}, \ldots, X_{n}$. Show that

$$
\lim _{n \rightarrow \infty} \mathbf{P}\left\{Z_{(1)}(X) \neq Y\right\}=\mathbf{E}\left\{\eta(X)+\eta^{\prime}(X)-2 \eta(X) \eta^{\prime}(X)\right\}
$$

(Lugosi (1992)).
Problem 5.41. Improve the bound in Lemma 5.3 to $\gamma_{d} \leq 3^{d}-1$.
Problem 5.42. Show that if $\left\{C\left(x_{1}, \pi / 6\right), \ldots, C\left(x_{\gamma_{d}}, \pi / 6\right)\right\}$ is a collection of cones covering $\mathcal{R}^{d}$, then $\gamma_{d} \geq 2^{d}$.

Problem 5.43. Recalling that $L_{k \mathrm{NN}}=\mathbf{E}\left\{\alpha_{k}(\eta(X))\right\}$, where

$$
\alpha_{k}(p)=\min (p, 1-p)+|1-2 \min (p, 1-p)| \mathbf{P}\left\{\operatorname{Binomial}(k, \min (p, 1-p))>\frac{k}{2}\right\}
$$

show that for every fixed $p, \mathbf{P}\left\{\operatorname{Binomial}(k, \min (p, 1-p))>\frac{k}{2}\right\} \downarrow 0$ as $k \rightarrow \infty$ (it is the monotonicity that is harder to show). How would you then prove that $\lim _{k \rightarrow \infty} L_{k \mathrm{~N} \mathrm{~N}}=L^{*}$ ?

Problem 5.44. Show that the asymptotic error probability of the rule that decides $g_{n}(x)=$ $Y_{(8)}(x)$ is identical to that of the rule in which $g_{n}(x)=Y_{(3)}(x)$.

Problem 5.45. Show that for all distributions $L_{5 \mathrm{NN}}=\mathbf{E}\left\{\psi_{5}(\eta(X)(1-\eta(X))\}\right.$, where $\psi_{5}(u)=$ $u+u^{2}+12 u^{3}$.

Problem 5.46. Show that for all distributions,

$$
L_{5 \mathrm{NN}} \geq \frac{L_{\mathrm{NN}}}{2}+\frac{L_{\mathrm{NN}}^{2}}{4}+\frac{3 L_{\mathrm{NN}}^{3}}{2}
$$

and that

$$
L_{3 \mathrm{NN}} \geq \frac{L_{\mathrm{NN}}}{2}+L_{\mathrm{NN}}^{2}
$$

Problem 5.47. Let $X_{(1)}$ be the nearest neighbor of $x$ among $X_{1}, \ldots, X_{n}$. Construct an example for which $\mathbf{E}\left\{\left\|X_{(1)}-x\right\|\right\}=\infty$ for all $x \in \mathcal{R}^{d}$. (Therefore, we have to steer clear of convergence in the mean in Lemma 5.1.) Let $X_{(1)}$ be the nearest neighbor of $X_{1}$ among $X_{2}, \ldots, X_{n}$. Construct a distribution such that $\mathbf{E}\left\{\left\|X_{(1)}-X_{1}\right\|\right\}=\infty$ for all $n$.

Problem 5.48. Consider the weighted nearest neighbor rule with weights ( $w_{1}, \ldots, w_{k}$ ). Define a new weight vector $\left(w_{1}, w_{2}, \ldots, w_{k-1}, v_{1}, \ldots, v_{l}\right)$, where $\sum_{i=1}^{l} v_{i}=w_{k}$. Thus, the weight vectors are partially ordered by the operation "partition." Assume that all weights are nonnegative. Let the asymptotic expected probability of errors be $L$ and $L^{\prime}$, respectively. True or false: for all distributions of $(X, Y), L^{\prime} \leq L$.

Problem 5.49. Gabriel neighbors. Given $X_{1}, \ldots, X_{n} \in \mathcal{R}^{d}$, we say that $X_{i}$ and $X_{j}$ are Gabriel neighbors if the ball centered at $\left(X_{i}+X_{j}\right) / 2$ of radius $\left\|X_{i}-X_{j}\right\| / 2$ contains no $X_{k}$, $k \neq i, j$ (Gabriel and Sokal (1969); Matula and Sokal (1980)). Clearly, if $X_{j}$ is the nearest neighbor of $X_{i}$, then $X_{i}$ and $X_{j}$ are Gabriel neighbors. Show that if $X$ has a density and $X_{1}, \ldots, X_{n}$ are i.i.d. and drawn from the distribution of $X$, then the expected number of Gabriel neighbors of $X_{1}$ tends to $2^{d}$ as $n \rightarrow \infty$ (Devroye (1988c)).


Figure 5.7. The Gabriel graph of 20 points on the plane is shown: Gabriel neighbors are connected by an edge. Note that all circles straddling these edges have no data point in their interior.

Problem 5.50. Gabriel neighbor rule. Define the Gabriel neighbor rule simply as the rule that takes a majority vote over all $Y_{i}$ 's for the Gabriel neighbors of $X$ among $X_{1}, \ldots, X_{n}$. Ties are broken by flipping a coin. Let $L_{n}$ be the conditional probability of error for the Gabriel rule. Using the result of the previous exercise, show that if $L^{*}$ is the Bayes error, then
(1) $\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}=0$ if $L^{*}=0$;
(2) $\lim \sup \mathbf{E}\left\{L_{n}\right\}<L_{\mathrm{NN}}$ if $L^{*}>0, d>1$;
(3) $\lim \sup \mathbf{E}\left\{L_{n}\right\} \leq c L^{*}$ for some $c<2$, if $d>1$.

For (3), determine the best possible value of $c$. Hint: Use Theorem 5.8 and try obtaining, for $d=2$, a lower bound for $\mathbf{P}\left\{N_{X} \geq 3\right\}$, where $N_{X}$ is the number of Gabriel neighbors of $X$ among $X_{1}, \ldots, X_{n}$.

## 6

## Consistency

### 6.1 Universal Consistency

If we are given a sequence $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$ of training data, the best we can expect from a classification function is to achieve the Bayes error probability $L^{*}$. Generally, we cannot hope to obtain a function that exactly achieves the Bayes error probability, but it is possible to construct a sequence of classification functions $\left\{g_{n}\right\}$, that is, a classification rule, such that the error probability

$$
L_{n}=L\left(g_{n}\right)=\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y \mid D_{n}\right\}
$$

gets arbitrarily close to $L^{*}$ with large probability (that is, for "most" $D_{n}$ ). This idea is formulated in the definitions of consistency:

DEFINITION 6.1. (WEAK AND STRONG CONSISTENCY). A classification rule is consistent (or asymptotically Bayes-risk efficient) for a certain distribution of $(X, Y)$ if

$$
\mathbf{E} L_{n}=\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y\right\} \rightarrow L^{*} \text { as } n \rightarrow \infty
$$

and strongly consistent if

$$
\lim _{n \rightarrow \infty} L_{n}=L^{*} \quad \text { with probability } 1
$$

REMARK. Consistency is defined as the convergence of the expected value of $L_{n}$ to $L^{*}$. Since $L_{n}$ is a random variable bounded between $L^{*}$ and 1 , this convergence
is equivalent to the convergence of $L_{n}$ to $L^{*}$ in probability, which means that for every $\epsilon>0$

$$
\lim _{n \rightarrow \infty} \mathbf{P}\left\{L_{n}-L^{*}>\epsilon\right\}=0
$$

Obviously, since almost sure convergence always implies convergence in probability, strong consistency implies consistency.

A consistent rule guarantees that by increasing the amount of data the probability that the error probability is within a very small distance of the optimal achievable gets arbitrarily close to one. Intuitively, the rule can eventually learn the optimal decision from a large amount of training data with high probability. Strong consistency means that by using more data the error probability gets arbitrarily close to the optimum for every training sequence except for a set of sequences that has zero probability altogether.

A decision rule can be consistent for a certain class of distributions of $(X, Y)$, but may not be consistent for others. It is clearly desirable to have a rule that is consistent for a large class of distributions. Since in many situations we do not have any prior information about the distribution, it is essential to have a rule that gives good performance for all distributions. This very strong requirement of universal goodness is formulated as follows:

DEFINITION 6.2. (UNIVERSAL CONSISTENCY). A sequence of decision rules is called universally (strongly) consistent if it is (strongly) consistent for any distribution of the pair $(X, Y)$.

In this chapter we show that such universally consistent classification rules exist. At first, this may seem very surprising, for some distributions are very "strange," and seem hard to learn. For example, let $X$ be uniformly distributed on [0, 1] with probability $1 / 2$, and let $X$ be atomic on the rationals with probability $1 / 2$. For example, if the rationals are enumerated $r_{1}, r_{2}, r_{3}, \ldots$, then $\mathbf{P}\left\{X=r_{i}\right\}=1 / 2^{i+1}$. Let $Y=1$ if $X$ is rational and $Y=0$ if $X$ is irrational. Obviously, $L^{*}=0$. If a classification rule $g_{n}$ is consistent, then the probability of incorrectly guessing the rationality of $X$ tends to zero. Note here that we cannot "check" whether $X$ is rational or not, but we should base our decision solely on the data $D_{n}$ given to us. One consistent rule is the following: $g_{n}\left(x, D_{n}\right)=Y_{k}$ if $X_{k}$ is the closest point to $x$ among $X_{1}, \ldots, X_{n}$. The fact that the rationals are dense in $[0,1]$ makes the statement even more surprising. See Problem 6.3.

### 6.2 Classification and Regression Estimation

In this section we show how consistency of classification rules can be deduced from consistent regression estimation. In many cases the a posteriori probability $\eta(x)$ is estimated from the training data $D_{n}$ by some function $\eta_{n}(x)=\eta_{n}\left(x, D_{n}\right)$.

In this case, the error probability $L\left(g_{n}\right)=\mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}$ of the plug-in rule

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

is a random variable. Then a simple corollary of Theorem 2.2 is as follows:

COROLLARy 6.1. The error probability of the classifier $g_{n}(x)$ defined above satisfies the inequality

$$
L\left(g_{n}\right)-L^{*} \leq 2 \int_{\mathcal{R}^{d}}\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)=2 \mathbf{E}\left\{\left|\eta(X)-\eta_{n}(X)\right| \mid D_{n}\right\}
$$

The next corollary follows from the Cauchy-Schwarz inequality.

Corollary 6.2. If

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

then its error probability satisfies

$$
\mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}-L^{*} \leq 2 \sqrt{\int_{\mathcal{R}^{d}}\left|\eta(x)-\eta_{n}(x)\right|^{2} \mu(d x)}
$$

Clearly, $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}=\mathbf{E}\{Y \mid X=x\}$ is just the regression function of $Y$ on $X$. Therefore, the most interesting consequence of Theorem 2.2 is that the mere existence of a regression function estimate $\eta_{n}(x)$ for which

$$
\int\left|\eta(x)-\eta_{n}(x)\right|^{2} \mu(d x) \rightarrow 0
$$

in probability or with probability one implies that the plug-in decision rule $g_{n}$ is consistent or strongly consistent, respectively.

Clearly, from Theorem 2.3, one can arrive at a conclusion analogous to Corollary 6.1 when the probabilities $\eta_{0}(x)=\mathbf{P}\{Y=0 \mid X=x\}$ and $\eta_{1}(x)=\mathbf{P}\{Y=1 \mid X=x\}$ are estimated from data separately by some $\eta_{0, n}$ and $\eta_{1, n}$, respectively. Usually, a key part of proving consistency of classification rules is writing the rules in one of the plug-in forms, and showing $L_{1}$-convergence of the approximating functions to the a posteriori probabilities. Here we have some freedom, as for any positive function $\tau_{n}(x)$, we have, for example,

$$
g_{n}(x)=\left\{\begin{array}{ll}
0 & \text { if } \eta_{1, n}(x) \leq \eta_{0, n}(x) \\
1 & \text { otherwise },
\end{array}= \begin{cases}0 & \text { if } \frac{\eta_{1, n}(x)}{\tau_{n}(x)} \leq \frac{\eta_{0, n}(x)}{\tau_{n}(x)} \\
1 & \text { otherwise }\end{cases}\right.
$$

### 6.3 Partitioning Rules

Many important classification rules partition $\mathcal{R}^{d}$ into disjoint cells $A_{1}, A_{2}, \ldots$ and classify in each cell according to the majority vote among the labels of the $X_{i}$ 's falling in the same cell. More precisely,

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} I_{\left\{X_{i} \in A(x)\right\}} \leq \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} I_{\left\{X_{i} \in A(x)\right\}} \\ 1 & \text { otherwise },\end{cases}
$$

where $A(x)$ denotes the cell containing $x$. The decision is zero if the number of ones does not exceed the number of zeros in the cell where $X$ falls, and vice versa. The partitions we consider in this section may change with $n$, and they may also depend on the points $X_{1}, \ldots, X_{n}$, but we assume that the labels do not play a role in constructing the partition. The next theorem is a general consistency result for such partitioning rules. It requires two properties of the partition: first, cells should be small enough so that local changes of the distribution can be detected. On the other hand, cells should be large enough to contain a large number of points so that averaging among the labels is effective. $\operatorname{diam}(A)$ denotes the diameter of a set $A$, that is,

$$
\operatorname{diam}(A)=\sup _{x, y \in A}\|x-y\| .
$$

Let

$$
N(x)=n \mu_{n}(A(x))=\sum_{i=1}^{n} I_{\left\{X_{i} \in A(x)\right\}}
$$

denote the number of $X_{i}$ 's falling in the same cell as $x$. The conditions of the theorem below require that a random cell-selected according to the distribution of $X$-has a small diameter, and contains many points with large probability.

Theorem 6.1. Consider a partitioning classification rule as defined above. Then $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ if
(1) $\operatorname{diam}(A(X)) \rightarrow 0$ in probability,
(2) $N(X) \rightarrow \infty$ in probability.

Proof. Define $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$. From Corollary 6.1 we recall that we need only show $\mathbf{E}\left\{\left|\widehat{\eta}_{n}(X)-\eta(X)\right|\right\} \rightarrow 0$, where

$$
\widehat{\eta}_{n}(x)=\frac{1}{N(x)} \sum_{i: X_{i} \in A(x)} Y_{i} .
$$

Introduce $\bar{\eta}(x)=\mathbf{E}\{\eta(X) \mid X \in A(x)\}$. By the triangle inequality,

$$
\mathbf{E}\left\{\left|\widehat{\eta}_{n}(X)-\eta(X)\right|\right\} \leq \mathbf{E}\left\{\left|\widehat{\eta}_{n}(X)-\bar{\eta}(X)\right|\right\}+\mathbf{E}\{|\bar{\eta}(X)-\eta(X)|\} .
$$

By conditioning on the random variable $N(x)$, it is easy to see that $N(x) \widehat{\eta}_{n}(x)$ is distributed as $B(N(x), \bar{\eta}(x))$, a binomial random variable with parameters $N(x)$
and $\bar{\eta}(x)$. Thus,

$$
\begin{aligned}
& \mathbf{E}\left\{\left|\widehat{\eta}_{n}(X)-\bar{\eta}(X)\right| \mid X, I_{\left\{X_{1} \in A(X)\right\}}, \ldots, I_{\left\{X_{n} \in A(X)\right\}}\right\} \\
& \leq \\
& \leq \\
& +\quad \mathbf{E}\left\{\left.\left|\frac{B(N(X), \bar{\eta}(X))}{N(X)}-\bar{\eta}(X)\right| I_{\{N(X)>0\}} \right\rvert\, X, I_{\left\{X_{1} \in A(X)\right\}}, \ldots, I_{\left\{X_{n} \in A(X)\right\}}\right\} \\
& \leq \\
& \mathbf{E}\left\{\left.\sqrt{\frac{\bar{\eta}(X)(1-\bar{\eta}(X))}{N(X)}} I_{\{N(X)>0\}} \right\rvert\, X, I_{\left\{X_{1} \in A(X)\right\}}, \ldots, I_{\left\{X_{n} \in A(X)\right\}}\right\} \\
& \\
& \quad+I_{\{N(X)=0\}}
\end{aligned}
$$

by the Cauchy-Schwarz inequality. Taking expectations, we see that

$$
\begin{aligned}
\mathbf{E}\left\{\left|\widehat{\eta}_{n}(X)-\bar{\eta}(X)\right|\right\} & \leq \mathbf{E}\left\{\frac{1}{2 \sqrt{N(X)}} I_{\{N(X)>0\}}\right\}+\mathbf{P}\{N(X)=0\} \\
& \leq \frac{1}{2} \mathbf{P}\{N(X) \leq k\}+\frac{1}{2 \sqrt{k}}+\mathbf{P}\{N(X)=0\}
\end{aligned}
$$

for any $k$, and this can be made small, first by choosing $k$ large enough and then by using condition (2).

For $\epsilon>0$, find a uniformly continuous [0,1]-valued function $\eta_{\epsilon}$ on a bounded set $C$ and vanishing off $C$ so that $\mathbf{E}\left\{\left|\eta_{\epsilon}(X)-\eta(X)\right|\right\}<\epsilon$. Next, we employ the triangle inequality:

$$
\begin{aligned}
\mathbf{E}\{|\bar{\eta}(X)-\eta(X)|\} \leq & \mathbf{E}\left\{\left|\bar{\eta}(X)-\bar{\eta}_{\epsilon}(X)\right|\right\} \\
& +\mathbf{E}\left\{\left|\bar{\eta}_{\epsilon}(X)-\eta_{\epsilon}(X)\right|\right\} \\
& +\mathbf{E}\left\{\left|\eta_{\epsilon}(X)-\eta(X)\right|\right\} \\
= & I+I I+I I I
\end{aligned}
$$

where $\bar{\eta}_{\epsilon}(x)=\mathbf{E}\left\{\eta_{\epsilon}(X) \mid X \in A(x)\right\}$. Clearly, $I I I<\epsilon$ by choice of $\eta_{\epsilon}$. Since $\eta_{\epsilon}$ is uniformly continuous, we can find a $\theta=\theta(\epsilon)>0$ such that

$$
I I \leq \epsilon+\mathbf{P}\{\operatorname{diam}(A(X))>\theta\}
$$

Therefore, $I I<2 \epsilon$ for $n$ large enough, by condition (1). Finally, $I \leq I I I<\epsilon$. Taken together these steps prove the theorem.

### 6.4 The Histogram Rule

In this section we describe the cubic histogram rule and show its universal consistency by checking the conditions of Theorem 6.1. The rule partitions $\mathcal{R}^{d}$ into
cubes of the same size, and makes a decision according to the majority vote among the $Y_{i}$ 's such that the corresponding $X_{i}$ falls in the same cube as $X$. Formally, let $\mathcal{P}_{n}=\left\{A_{n 1}, A_{n 2}, \ldots\right\}$ be a partition of $\mathcal{R}^{d}$ into cubes of size $h_{n}>0$, that is, into sets of the type $\prod_{i=1}^{d}\left[k_{i} h_{n},\left(k_{i}+1\right) h_{n}\right)$, where the $k_{i}$ 's are integers. For every $x \in \mathcal{R}^{d}$ let $A_{n}(x)=A_{n i}$ if $x \in A_{n i}$. The histogram rule is defined by

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} I_{\left\{X_{i} \in A_{n}(x)\right\}} \leq \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} I_{\left\{X_{i} \in A_{n}(x)\right\}} \\ 1 & \text { otherwise. }\end{cases}
$$



Consistency of the histogram rule was established under some additional conditions by Glick (1973). Universal consistency follows from the results of Gordon and Olshen (1978), (1980). A direct proof of strong universal consistency is given in Chapter 9.

The next theorem establishes universal consistency of certain cubic histogram rules.

Theorem 6.2. If $h_{n} \rightarrow 0$ and $n h_{n}^{d} \rightarrow \infty$ as $n \rightarrow \infty$, then the cubic histogram rule is universally consistent.

Proof. We check the two simple conditions of Theorem 6.1. Clearly, the diameter of each cell is $\sqrt{d} h^{d}$. Therefore condition (1) follows trivially. To show condition (2), we need to prove that for any $M<\infty, \mathbf{P}\{N(X) \leq M\} \rightarrow 0$. Let $S$ be an arbitrary ball centered at the origin. Then the number of cells intersecting $S$ is not more than $c_{1}+c_{2} / h^{d}$ for some positive constants $c_{1}, c_{2}$. Then

$$
\begin{aligned}
& \mathbf{P}\{N(X) \leq M\} \\
& \quad \leq \sum_{j: A_{n j} \cap S \neq \neq \emptyset} \mathbf{P}\left\{X \in A_{n j}, N(X) \leq M\right\}+\mathbf{P}\left\{X \in S^{c}\right\} \\
& \quad \leq \sum_{\substack{j: A_{n j} \cap \cap \neq \varnothing \\
\mu\left(A_{n j}\right) \leq 2 M / n}} \mu\left(A_{n} j\right)+\sum_{\substack{j: A_{n j} \cap \cap \neq \neq \emptyset \\
\mu\left(A_{n j}\right)>2 M / n}} \mu\left(A_{n j}\right) \mathbf{P}\left\{n \mu_{n}\left(A_{n j}\right) \leq M\right\}+\mu\left(S^{c}\right)
\end{aligned}
$$

$$
\begin{aligned}
\leq & \frac{2 M}{n}\left(c_{1}+\frac{c_{2}}{h^{d}}\right) \\
& +\sum_{\substack{j: A_{n j} \cap \cap \neq \nmid \emptyset \\
\mu\left(A_{n j}\right)>2 M / n}} \mu\left(A_{n j}\right) \mathbf{P}\left\{\mu_{n}\left(A_{n j}\right)-\mathbf{E}\left\{\mu_{n}\left(A_{n j}\right)\right\} \leq M / n-\mu\left(A_{n j}\right)\right\}+\mu\left(S^{c}\right) \\
\leq & \frac{2 M}{n}\left(c_{1}+\frac{c_{2}}{h^{d}}\right) \\
& +\sum_{\substack{j: A_{n j} \cap \cap\{\neq \neq 0}} \mu\left(A_{n j}\right) \mathbf{P}\left\{\mu_{n}\left(A_{n j}\right)-\mathbf{E}\left\{\mu_{n}\left(A_{n j}\right)\right\} \leq \frac{-\mu\left(A_{n j}\right)}{2}\right\}+\mu\left(S^{c}\right) \\
\leq & \frac{2 M}{n}\left(c_{1}+\frac{c_{2}}{h^{d}}\right)+\sum_{\substack{j: A_{n j} \cap \cap \neq \neq \emptyset / n \\
\mu\left(A_{n j}\right)>2 M / n}} 4 \mu\left(A_{n j}\right) \frac{\operatorname{Var}\left\{\mu_{n}\left(A_{n j}\right)\right\}}{\left(\mu\left(A_{n j}\right)\right)^{2}}+\mu\left(S^{c}\right)
\end{aligned}
$$

(by Chebyshev's inequality)
$\leq \frac{2 M}{n}\left(c_{1}+\frac{c_{2}}{h^{d}}\right)+\sum_{\substack{j: A_{n j} \cap\left\{\neq \bar{\emptyset} \\ \mu\left(A_{n j}\right)>2 M / n\right.}} 4 \mu\left(A_{n j}\right) \frac{1}{n \mu\left(A_{n j}\right)}+\mu\left(S^{c}\right)$
$\leq \frac{2 M+4}{n}\left(c_{1}+\frac{c_{2}}{h^{d}}\right)+\mu\left(S^{c}\right)$
$\rightarrow \mu\left(S^{c}\right)$,
because $n h^{d} \rightarrow \infty$. Since $S$ is arbitrary, the proof of the theorem is complete.

### 6.5 Stone's Theorem

A general theorem by Stone (1977) allows us to deduce universal consistency of several classification rules. Consider a rule based on an estimate of the a posteriori probability $\eta$ of the form

$$
\eta_{n}(x)=\sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} W_{n i}(x)=\sum_{i=1}^{n} Y_{i} W_{n i}(x)
$$

where the weights $W_{n i}(x)=W_{n i}\left(x, X_{1}, \ldots, X_{n}\right)$ are nonnegative and sum to one:

$$
\sum_{i=1}^{n} W_{n i}(x)=1
$$

The classification rule is defined as

$$
\begin{aligned}
g_{n}(x) & = \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} W_{n i}(x) \leq \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} W_{n i}(x) \\
1 & \text { otherwise },\end{cases} \\
& = \begin{cases}0 & \text { if } \sum_{i=1}^{n} Y_{i} W_{n i}(x) \leq 1 / 2 \\
1 & \text { otherwise } .\end{cases}
\end{aligned}
$$

$\eta_{n}$ is a weighted average estimator of $\eta$. It is intuitively clear that pairs ( $X_{i}, Y_{i}$ ) such that $X_{i}$ is close to $x$ should provide more information about $\eta(x)$ than those far from $x$. Thus, the weights are typically much larger in the neighborhood of $X$, so $\eta_{n}$ is roughly a (weighted) relative frequency of the $X_{i}$ 's that have label 1 among points in the neighborhood of $X$. Thus, $\eta_{n}$ might be viewed as a local average estimator, and $g_{n}$ a local (weighted) majority vote. Examples of such rules include the histogram, kernel, and nearest neighbor rules. These rules will be studied in depth later.

Theorem 6.3. (STONE (1977)). Assume that for any distribution of $X$, the weights satisfy the following three conditions:
(i) There is a constant c such that, for every nonnegative measurable function $f$ satisfying $\mathbf{E} f(X)<\infty$,

$$
\mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X) f\left(X_{i}\right)\right\} \leq c \mathbf{E} f(X)
$$

(ii) For all $a>0$,

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X) I_{\left\{\left\|X_{i}-X\right\|>a\right\}}\right\}=0
$$

(iii)

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{\max _{1 \leq i \leq n} W_{n i}(X)\right\}=0
$$

Then $g_{n}$ is universally consistent.

REMARK. Condition (ii) requires that the overall weight of $X_{i}$ 's outside of any ball of a fixed radius centered at $X$ must go to zero. In other words, only points in a shrinking neighborhood of $X$ should be taken into account in the averaging. Condition (iii) requires that no single $X_{i}$ has too large a contribution to the estimate. Hence, the number of points encountered in the averaging must tend to infinity. Condition (i) is technical.

Proof. By Corollary 6.2 it suffices to show that for every distribution of $(X, Y)$

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{\left(\eta(X)-\eta_{n}(X)\right)^{2}\right\}=0
$$

Introduce the notation

$$
\widehat{\eta}_{n}(x)=\sum_{i=1}^{n} \eta\left(X_{i}\right) W_{n i}(x) .
$$

Then by the simple inequality $(a+b)^{2} \leq 2\left(a^{2}+b^{2}\right)$ we have

$$
\begin{align*}
\mathbf{E} & \left\{\left(\eta(X)-\eta_{n}(X)\right)^{2}\right\} \\
& =\mathbf{E}\left\{\left(\left(\eta(X)-\widehat{\eta}_{n}(X)\right)+\left(\widehat{\eta}_{n}(X)-\eta_{n}(X)\right)\right)^{2}\right\} \\
& \leq 2\left(\mathbf{E}\left\{\left(\eta(X)-\widehat{\eta}_{n}(X)\right)^{2}\right\}+\mathbf{E}\left\{\left(\widehat{\eta}_{n}(X)-\eta_{n}(X)\right)^{2}\right\}\right) . \tag{6.1}
\end{align*}
$$

Therefore, it is enough to show that both terms on the right-hand side tend to zero. Since the $W_{n i}$ 's are nonnegative and sum to one, by Jensen's inequality, the first term is

$$
\begin{aligned}
\mathbf{E}\left\{\left(\eta(X)-\widehat{\eta}_{n}(X)\right)^{2}\right\} & =\mathbf{E}\left\{\left(\sum_{i=1}^{n} W_{n i}(X)\left(\eta(X)-\eta\left(X_{i}\right)\right)\right)^{2}\right\} \\
& \leq \mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X)\left(\eta(X)-\eta\left(X_{i}\right)\right)^{2}\right\}
\end{aligned}
$$

If the function $0 \leq \eta^{*} \leq 1$ is continuous with bounded support, then it is uniformly continuous as well: for every $\epsilon>0$, there is an $a>0$ such that for $\left\|x_{1}-x\right\|<a$, $\left|\eta^{*}\left(x_{1}\right)-\eta^{*}(x)\right|^{2}<\epsilon$. Recall here that $\|x\|$ denotes the Euclidean norm of a vector $x \in \mathcal{R}^{d}$. Thus, since $\left|\eta^{*}\left(x_{1}\right)-\eta^{*}(x)\right| \leq 1$,

$$
\begin{aligned}
& \mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X)\left(\eta^{*}(X)-\eta^{*}\left(X_{i}\right)\right)^{2}\right\} \\
& \quad \leq \mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X) I_{\left\{\left\|X-X_{i}\right\| \geq a\right\}}\right\}+\mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X) \epsilon\right\} \rightarrow \epsilon,
\end{aligned}
$$

by (ii). Since the set of continuous functions with bounded support is dense in $L_{2}(\mu)$, for every $\epsilon>0$ we can choose $\eta^{*}$ such that

$$
\mathbf{E}\left\{\left(\eta(X)-\eta^{*}(X)\right)^{2}\right\}<\epsilon
$$

By this choice, using the inequality $(a+b+c)^{2} \leq 3\left(a^{2}+b^{2}+c^{2}\right)$ (which follows from the Cauchy-Schwarz inequality),

$$
\begin{aligned}
& \mathbf{E}\left\{\left(\eta(X)-\widehat{\eta}_{n}(X)\right)^{2}\right\} \\
& \leq \\
& \leq \quad \mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X)\left(\eta(X)-\eta\left(X_{i}\right)\right)^{2}\right\} \\
& \leq \\
& \\
& \quad 3 \mathbf{E}\left\{\sum _ { i = 1 } ^ { n } W _ { n i } ( X ) \left(\left(\eta(X)-\eta^{*}(X)\right)^{2}+\left(\eta^{*}(X)-\eta^{*}\left(X_{i}\right)\right)^{2}\right.\right. \\
& \leq \\
& \quad 3 \mathbf{E}\left\{\left(\eta(X)-\eta^{*}(X)\right)^{2}\right\} \\
& \\
& \quad+3 \mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X)\left(\eta^{*}(X)-\eta^{*}\left(X_{i}\right)\right)^{2}\right\}+3 c \mathbf{E}\left\{\left(\eta(X)-\eta^{*}(X)\right)^{2}\right\},
\end{aligned}
$$

where we used (i). Therefore,

$$
\limsup _{n \rightarrow \infty} \mathbf{E}\left\{\left(\eta(X)-\widehat{\eta}_{n}(X)\right)^{2}\right\} \leq 3 \epsilon(1+1+c)
$$

To handle the second term of the right-hand side of (6.1), observe that

$$
\mathbf{E}\left\{\left(Y_{i}-\eta\left(X_{i}\right)\right)\left(Y_{j}-\eta\left(X_{j}\right)\right) \mid X, X_{1}, \ldots, X_{n}\right\}=0 \text { for all } i \neq j
$$

by independence. Therefore,

$$
\begin{aligned}
\mathbf{E} & \left\{\left(\widehat{\eta}_{n}(X)-\eta_{n}(X)\right)^{2}\right\} \\
& =\mathbf{E}\left\{\left(\sum_{i=1}^{n} W_{n i}(X)\left(\eta\left(X_{i}\right)-Y_{i}\right)\right)^{2}\right\} \\
& =\sum_{i=1}^{n} \sum_{j=1}^{n} \mathbf{E}\left\{W_{n i}(X)\left(\eta\left(X_{i}\right)-Y_{i}\right) W_{n j}(X)\left(\eta\left(X_{j}\right)-Y_{j}\right)\right\} \\
& =\sum_{i=1}^{n} \mathbf{E}\left\{W_{n i}^{2}(X)\left(\eta\left(X_{i}\right)-Y_{i}\right)^{2}\right\} \\
& \leq \mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}^{2}(X)\right\} \leq \mathbf{E}\left\{\max _{1 \leq i \leq n} W_{n i}(X) \sum_{j=1}^{n} W_{n j}(X)\right\} \\
& =\mathbf{E}\left\{\max _{1 \leq i \leq n} W_{n i}(X)\right\} \rightarrow 0
\end{aligned}
$$

by (iii), and the theorem is proved.

### 6.6 The $k$-Nearest Neighbor Rule

In Chapter 5 we discussed asymptotic properties of the $k$-nearest neighbor rule when $k$ remains fixed as the sample size $n$ grows. In such cases the expected probability of error converges to a number between $L^{*}$ and $2 L^{*}$. In this section we show that if $k$ is allowed to grow with $n$ such that $k / n \rightarrow 0$, the rule is weakly universally consistent. The proof is a very simple application of Stone's theorem. This result, appearing in Stone's paper (1977), was the first universal consistency result for any rule. Strong consistency, and many other different aspects of the $k$-NN rule are studied in Chapters 11 and 26.

Recall the definition of the $k$-nearest neighbor rule: first the data are ordered according to increasing Euclidean distances of the $X_{j}$ 's to $x$ :

$$
\left(X_{(1)}(x), Y_{(1)}(x)\right), \ldots,\left(X_{(n)}(x), Y_{(n)}(x)\right)
$$

that is, $X_{(i)}(x)$ is the $i$-th nearest neighbor of $x$ among the points $X_{1}, \ldots, X_{n}$. Distance ties are broken by comparing indices, that is, in case of $\left\|X_{i}-x\right\|=$ $\left\|X_{j}-x\right\|, X_{i}$ is considered to be "closer" to $x$ if $i<j$.

The $k$-nN classification rule is defined as

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{k} I_{\left\{Y_{i()}(x)=1\right\}} \leq \sum_{i=1}^{k} I_{\left\{Y_{i(i)}(x)=0\right\}} \\ 1 & \text { otherwise. }\end{cases}
$$

In other words, $g_{n}(x)$ is a majority vote among the labels of the $k$ nearest neighbors of $x$.

Theorem 6.4. (Stone (1977)). If $k \rightarrow \infty$ and $k / n \rightarrow 0$, then for all distributions $\mathrm{E} L_{n} \rightarrow L^{*}$.

Proof. We proceed by checking the conditions of Stone's weak convergence theorem (Theorem 6.3). The weight $W_{n i}(X)$ in Theorem 6.3 equals $1 / k$ iff $X_{i}$ is among the $k$ nearest neighbors of $X$, and equals 0 otherwise.

Condition (iii) is obvious since $k \rightarrow \infty$. For condition (ii) observe that

$$
\mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X) I_{\left\{\left\|X_{i}-X\right\|>\epsilon\right\}}\right\} \rightarrow 0
$$

holds whenever

$$
\mathbf{P}\left\{\left\|X_{(k)}(X)-X\right\|>\epsilon\right\} \rightarrow 0
$$

where $X_{(k)}(x)$ denotes the $k$-th nearest neighbor of $x$ among $X_{1}, \ldots, X_{n}$. But we know from Lemma 5.1 that this is true for all $\epsilon>0$ whenever $k / n \rightarrow 0$.

Finally, we consider condition (i). We have to show that for any nonnegative measurable function $f$ with $\mathbf{E}\{f(X)\}<\infty$,

$$
\mathbf{E}\left\{\sum_{i=1}^{n} \frac{1}{k} I_{\left\{X_{i} \text { is among the } k \text { nearest neighbors of } X\right\}} f\left(X_{i}\right)\right\} \leq \mathbf{E}\{c f(X)\}
$$

for some constant $c$. But we have shown in Lemma 5.3 that this inequality always holds with $c=\gamma_{d}$. Thus, condition (i) is verified.

### 6.7 Classification Is Easier Than Regression Function Estimation

Once again assume that our decision is based on some estimate $\eta_{n}$ of the a posteriori probability function $\eta$, that is,

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

The bounds of Theorems 2.2, 2.3, and Corollary 6.2 point out that if $\eta_{n}$ is a consistent estimate of $\eta$, then the resulting rule is also consistent. For example, writing $L_{n}=\mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}$, we have

$$
\mathbf{E} L_{n}-L^{*} \leq 2 \sqrt{\mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\}}
$$

that is, $L_{2}$-consistent estimation of the regression function $\eta$ leads to consistent classification, and in fact, this is the main tool used in the proof of Theorem 6.3. While the said bounds are useful for proving consistency, they are almost useless when it comes to studying rates of convergence. As Theorem 6.5 below shows, for consistent rules rates of convergence of $\mathbf{P}\left\{g_{n}(X) \neq Y\right\}$ to $L^{*}$ are always orders of magnitude better than rates of convergence of $\sqrt{\mathbf{E}\left\{\left(\eta(X)-\eta_{n}(X)\right)^{2}\right\}}$ to zero.


FIGURE 6.2. The difference between the error probabilities grows roughly in proportion to the shaded area. Elsewhere $\eta_{n}(x)$ does not need to be close $\eta(x)$.
Pattern recognition is thus easier than regression function estimation. This will be a recurring theme-to achieve acceptable results in pattern recognition, we can do more with smaller sample sizes than in regression function estimation. This is really just a consequence of the fact that less is required in pattern recognition. It also corroborates our belief that pattern recognition is dramatically different from regression function estimation, and that it deserves separate treatment in the statistical community.

Theorem 6.5. Let $\eta_{n}$ be a weakly consistent regression estimate, that is,

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\}=0
$$

Define

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

Then

$$
\lim _{n \rightarrow \infty} \frac{\mathbf{E} L_{n}-L^{*}}{\sqrt{\mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\}}}=0
$$

that is, $\mathbf{E} L_{n}-L^{*}$ converges to zero faster than the $L_{2}$-error of the regression estimate.

Proof. We start with the equality of Theorem 2.2:

$$
\mathbf{E} L_{n}-L^{*}=2 \mathbf{E}\left\{|\eta(X)-1 / 2| I_{\left\{g_{n}(X) \neq g^{*}(X)\right\}}\right\}
$$

Fix $\epsilon>0$. We may bound the last factor by

$$
\begin{aligned}
& \mathbf{E}\left\{|\eta(X)-1 / 2| I_{\left\{g_{n}(X) \neq g^{*}(X)\right\}}\right\} \\
& \leq \mathbf{E}\left\{I_{\{\eta(X) \neq 1 / 2\}}\left|\eta(X)-\eta_{n}(X)\right| I_{\left\{g_{n}(X) \neq g^{*}(X)\right\}}\right\} \\
&= \mathbf{E}\left\{\left|\eta(X)-\eta_{n}(X)\right| I_{\left\{g_{n}(X) \neq g^{*}(X)\right\}} I_{\{|\eta(X)-1 / 2| \leq \epsilon\}} I_{\{\eta(X) \neq 1 / 2\}}\right\} \\
&+\mathbf{E}\left\{\left|\eta(X)-\eta_{n}(X)\right| I_{\left\{g_{n}(X) \neq g^{*}(X)\right\}} I_{\{|\eta(X)-1 / 2|>\epsilon\}}\right\} \\
& \leq \sqrt{\mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\}} \\
& \times(\sqrt{\mathbf{P}\{|\eta(X)-1 / 2| \leq \epsilon, \eta(X) \neq 1 / 2\}} \\
&\left.\quad+\sqrt{\mathbf{P}\left\{g_{n}(X) \neq g^{*}(X),|\eta(X)-1 / 2|>\epsilon\right\}}\right)
\end{aligned}
$$

(by the Cauchy-Schwarz inequality).
Since $g_{n}(X) \neq g^{*}(X)$ and $|\eta(X)-1 / 2|>\epsilon$ imply that $\left|\eta_{n}(X)-\eta(X)\right|>\epsilon$, consistency of the regression estimate implies that for any fixed $\epsilon>0$,

$$
\lim _{n \rightarrow \infty} \mathbf{P}\left\{g_{n}(X) \neq g^{*}(X),|\eta(X)-1 / 2|>\epsilon\right\}=0
$$

On the other hand,

$$
\mathbf{P}\{|\eta(X)-1 / 2| \leq \epsilon, \eta(X) \neq 1 / 2\} \rightarrow 0 \text { as } \epsilon \rightarrow 0
$$

which completes the proof.

The actual value of the ratio

$$
\rho_{n}=\frac{\mathbf{E} L_{n}-L^{*}}{\sqrt{\mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\}}}
$$

cannot be universally bounded. In fact, $\rho_{n}$ may tend to zero arbitrarily slowly (see Problem 6.5). On the other hand, $\rho_{n}$ may tend to zero extremely quickly. In Problems 6.6 and 6.7 and in the theorem below, upper bounds on $\rho_{n}$ are given that
may be used in deducing rate-of-convergence results. Theorem 6.6, in particular, states that $\mathbf{E} L_{n}-L^{*}$ tends to zero as fast as the square of the $L_{2}$ error of the regression estimate, i.e., $\mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\}$, whenever $L^{*}=0$. Just how slowly $\rho_{n}$ tends to zero depends upon two things, basically: (1) the rate of convergence of $\eta_{n}$ to $\eta$, and (2) the behavior of $\mathbf{P}\{|\eta(X)-1 / 2| \leq \epsilon, \eta(X) \neq 1 / 2\}$ as a function of $\epsilon$ when $\epsilon \downarrow 0$ (i.e., the behavior of $\eta(x)$ at those $x$ 's where $\eta(x) \approx 1 / 2$ ).

Theorem 6.6. Assume that $L^{*}=0$, and consider the decision

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise. }\end{cases}
$$

Then

$$
\mathbf{P}\left\{g_{n}(X) \neq Y\right\} \leq 4 \mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\} .
$$

Proof. By Theorem 2.2,

$$
\begin{aligned}
\mathbf{P}\left\{g_{n}(X) \neq Y\right\}= & 2 \mathbf{E}\left\{|\eta(X)-1 / 2| I_{\left\{g_{n}(X) \neq \exists^{*}(X)\right\}}\right\} \\
= & 2 \mathbf{E}\left\{|\eta(X)-1 / 2| I_{\left\{g_{n}(X) \neq Y\right\}}\right\} \\
& \left(\text { since } g^{*}(X)=Y \text { by the assumption } L^{*}=0\right) \\
\leq & 2 \sqrt{\mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\}} \sqrt{\mathbf{P}\left\{g_{n}(X) \neq Y\right\}} \\
& \text { (by the Cauchy-Schwarz inequality). }
\end{aligned}
$$

Dividing both sides by $\sqrt{\mathbf{P}\left\{g_{n}(X) \neq Y\right\}}$ yields the result.
The results above show that the bounds of Theorems 2.2, 2.3, and Corollary 6.2 may be arbitrarily loose, and the error probability converges to $L^{*}$ faster than the $L_{2}$-error of the regression estimate converges to zero. In some cases, consistency may even occur without convergence of $\mathbf{E}\left|\eta_{n}(X)-\eta(X)\right|$ to zero. Consider for example a strictly separable distribution, that is, a distribution such that there exist two sets $A, B \subset \mathcal{R}^{d}$ with

$$
\inf _{x \in A, y \in B}\|x-y\| \geq \delta>0
$$

for some $\delta>0$, and having the property that

$$
\mathbf{P}\{X \in A \mid Y=1\}=\mathbf{P}\{X \in B \mid Y=0\}=1 .
$$

In such cases, there is a version of $\eta$ that has $\eta(x)=1$ on $A$ and $\eta(x)=0$ on $B$. We say version because $\eta$ is not defined on sets of measure zero. For such strictly separable distributions, $L^{*}=0$. Let $\tilde{\eta}$ be $1 / 2-\epsilon$ on $B$ and $1 / 2+\epsilon$ on $A$. Then, with

$$
g(x)=\left\{\begin{array}{ll}
0 & \text { if } \tilde{\eta}(x) \leq 1 / 2 \\
1 & \text { otherwise, }
\end{array}= \begin{cases}0 & \text { if } x \in B \\
1 & \text { if } x \in A,\end{cases}\right.
$$

we have $\mathbf{P}\{g(X) \neq Y\}=L^{*}=0$. Yet,

$$
2 \mathbf{E}|\eta(X)-\tilde{\eta}(X)|=1-2 \epsilon
$$

is arbitrarily close to one.
In a more realistic example, we consider the kernel rule (see Chapter 10),

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

in which

$$
\eta_{n}(x)=\frac{\sum_{i=1}^{n} Y_{i} K\left(x-X_{i}\right)}{\sum_{i=1}^{n} K\left(x-X_{i}\right)}
$$

where $K$ is the standard normal density in $\mathcal{R}^{d}$ :

$$
K(u)=\frac{1}{(2 \pi)^{d / 2}} e^{-\|u\|^{2} / 2}
$$

Assume that $A$ and $B$ consist of one point each, at distance $\delta$ from each other-that is, the distribution of $X$ is concentrated on two points. If $\mathbf{P}\{Y=0\}=\mathbf{P}\{Y=1\}=$ $1 / 2$, we see that

$$
\lim _{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^{n} K\left(x-X_{i}\right)=\frac{K(0)+K(\delta)}{2} \quad \text { with probability one }
$$

at $x \in A \cup B$, by the law of large numbers. Also,

$$
\lim _{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^{n} Y_{i} K\left(x-X_{i}\right)=\left\{\begin{array}{ll}
K(0) / 2 & \text { if } x \in A \\
K(\delta) / 2 & \text { if } x \in B
\end{array} \quad\right. \text { with probability one. }
$$

Thus,

$$
\lim _{n \rightarrow \infty} \eta_{n}(x)=\left\{\begin{array}{ll}
\frac{K(0)}{K(0)+K(\delta)} & \text { if } x \in A \\
\frac{K(\delta)}{K(0)+K(\delta)} & \text { if } x \in B
\end{array} \quad\right. \text { with probability one. }
$$

Hence, as $\eta(x)=1$ on $A$ and $\eta(x)=0$ on $B$,

$$
\begin{aligned}
& \lim _{n \rightarrow \infty} 2 \mathbf{E}\left|\eta(X)-\eta_{n}(X)\right| \\
& \quad=2 \frac{1}{2}\left(\left|1-\frac{K(0)}{K(0)+K(\delta)}\right|+\left|0-\frac{K(\delta)}{K(0)+K(\delta)}\right|\right) \\
& \quad=\frac{2 K(\delta)}{K(0)+K(\delta)} .
\end{aligned}
$$

Yet, $L^{*}=0$ and $\mathbf{P}\left\{g_{n}(X) \neq Y\right\} \rightarrow 0$. In fact, if $D_{n}$ denotes the training data,

$$
\lim _{n \rightarrow \infty} \mathbb{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}=L^{*} \quad \text { with probability one }
$$

and

$$
\lim _{n \rightarrow \infty} 2 \mathbf{E}\left\{\left|\eta_{n}(X)-\eta(X)\right| \mid D_{n}\right\}=\frac{2 K(\delta)}{K(0)+K(\delta)} \quad \text { with probability one. }
$$

This shows very strongly that for any $\delta>0$, for many practical classification rules, we do not need convergence of $\eta_{n}$ to $\eta$ at all!

As all the consistency proofs in Chapters 6 through 11 rely on the convergence of $\eta_{n}$ to $\eta$, we will create unnecessary conditions for some distributions, although it will always be possible to find distributions of $(X, Y)$ for which the conditions are needed-in the latter sense, the conditions of these universal consistency results are not improvable.

### 6.8 Smart Rules

A rule is a sequence of mappings $g_{n}: \mathcal{R}^{d} \times\left(\mathcal{R}^{d} \times\{0,1\}\right)^{n} \rightarrow\{0,1\}$. Most rules are expected to perform better when $n$ increases. So, we say that a rule is smart if for all distributions of $(X, Y), \mathbf{E}\left\{L\left(g_{n}\right)\right\}$ is nonincreasing, where

$$
L\left(g_{n}\right)=\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y \mid D_{n}\right\}
$$

Some dumb rules are smart, such as the (useless) rule that, for each $n$, takes a majority over all $Y_{i}$ 's, ignoring the $X_{i}$ 's. This follows from the fact that

$$
\mathbf{P}\left\{\sum_{i=1}^{n}\left(2 Y_{i}-1\right)>0, Y=0 \text { or } \sum_{i=1}^{n}\left(2 Y_{i}-1\right) \leq 0, Y=1\right\}
$$

is monotone in $n$. This is a property of the binomial distribution (see Problem 6.12). A histogram rule with a fixed partition is smart (Problem 6.13). The 1nearest neighbor rule is not smart. To see this, let $(X, Y)$ be $(0,1)$ and $(Z, 0)$ with probabilities $p$ and $1-p$, respectively, where $Z$ is uniform on $[-1000,1000]$. Verify that for $n=1, \mathbf{E} L_{n}=2 p(1-p)$, while for $n=2$,

$$
\begin{aligned}
\mathbf{E} L_{n} & =2 p(1-p)^{2}\left(\frac{1}{2}+\frac{\mathbf{E}|Z|}{4000}\right)+p^{2}(1-p)+(1-p)^{2} p \\
& =2 p(1-p)\left(\frac{5(1-p)}{8}+\frac{1}{2}\right)
\end{aligned}
$$

which is larger than $2 p(1-p)$ whenever $p \in(0,1 / 5)$. This shows that in all these cases it is better to have $n=1$ than $n=2$. Similarly, the standard kernel rule-discussed in Chapter 10-with fixed $h$ is not smart (see Problems 6.14, 6.15).

The error probabilities of the above examples of smart rules do not change dramatically with $n$. However, change is necessary to guarantee Bayes risk consistency. At the places of change-for example when $h_{n}$ jumps to a new value in the histogram rule-the monotonicity may be lost. This leads to the conjecture that no universally consistent rule can be smart.

## Problems and Exercises

Problem 6.1. Let the i.i.d. random variables $X_{1}, \ldots, X_{n}$ be distributed on $\mathcal{R}^{d}$ according to the density $f$. Estimate $f$ by $f_{n}$, a function of $x$ and $X_{1}, \ldots, X_{n}$, and assume that $\int\left|f_{n}(x)-f(x)\right| d x \rightarrow 0$ in probability (or with probability one). Then show that there exists a consistent (or strongly consistent) classification rule whenever the conditional densities $f_{0}$ and $f_{1}$ exist.

Problem 6.2. Histogram density estimation. Let $X_{1}, \ldots, X_{n}$ be i.i.d. random variables in $\mathcal{R}^{d}$ with density $f$. Let $\mathcal{P}_{n}$ be a partition of $\mathcal{R}^{d}$ into cubes of size $h_{n}$, and define the histogram density estimate by

$$
f_{n}(x)=\frac{1}{n h_{n}^{d}} \sum_{i=1}^{n} I_{\left\{X_{i} \in A_{n}(x)\right\}},
$$

where $A_{n}(x)$ is the set in $\mathcal{P}_{n}$ that contains $x$. Prove that the estimate is universally consistent in $L_{1}$ if $h_{n} \rightarrow 0$ and $n h_{n}^{d} \rightarrow \infty$ as $n \rightarrow \infty$, that is, for any $f$ the $L_{1}$ error of the estimate $\int\left|f_{n}(x)-f(x)\right| d x$ converges to zero in probability, or equivalently,

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{\int\left|f_{n}(x)-f(x)\right| d x\right\}=0
$$

Hint: The following suggestions may be helpful
(1) $\mathbf{E}\left\{\int\left|f_{n}-f\right|\right\} \leq \mathbf{E}\left\{\int\left|f_{n}-\mathbf{E} f_{n}\right|\right\}+\int\left|\mathbf{E} f_{n}-f\right|$.
(2) $\mathbf{E}\left\{\int\left|f_{n}-\mathbf{E} f_{n}\right|\right\}=\sum_{j}\left|\mu\left(A_{n j}\right)-\mu_{n}\left(A_{n j}\right)\right|$.
(3) First show $\int\left|\mathbf{E} f_{n}-f\right| \rightarrow 0$ for uniformly continuous $f$, and then extend it to arbitrary densities.

Problem 6.3. Let $X$ be uniformly distributed on [0, 1] with probability $1 / 2$, and let $X$ be atomic on the rationals with probability $1 / 2$ (e.g., if the rationals are enumerated $r_{1}, r_{2}, r_{3}, \ldots$, then $\mathbf{P}\left\{X=r_{i}\right\}=1 / 2^{i+1}$ ). Let $Y=1$ if $X$ is rational and $Y=0$ if $X$ is irrational. Give a direct proof of consistency of the 1-nearest neighbor rule. Hint: Given $Y=1$, the conditional distribution of $X$ is discrete. Thus, for every $\epsilon>0$, there is an integer $k$ such that given $Y=1, X$ equals one of $k$ rationals with probability at least $1-\epsilon$. Now, if $n$ is large enough, every point in this set captures data points with label 1 with large probability. Also, for large $n$, the space between these points is filled with data points labeled with zeros.

Problem 6.4. Prove the consistency of the cubic histogram rule by checking the conditions of Stone's theorem. Hint: To check (i), first bound $W_{n i}(x)$ by

$$
I_{\left\{X_{i} \in A_{n}(x)\right\}} / \sum_{j=1}^{n} I_{\left\{X_{j} \in A_{n}(x)\right\}}+1 / n .
$$

Since

$$
\mathbf{E}\left\{\sum_{i=1}^{n} \frac{1}{n} f\left(X_{i}\right)\right\}=\mathbf{E} f(X)
$$

it suffices to show that there is a constant $c^{\prime}>0$ such that for any nonnegative function $f$ with $\mathbf{E} f(X)<\infty$,

$$
\mathbf{E}\left\{\sum_{i=1}^{n} f\left(X_{i}\right) \frac{I_{\left\{X_{i} \in A_{n}(X)\right\}}}{\sum_{j=1}^{n} I_{\left\{X_{j} \in A_{n}(X)\right\}}}\right\} \leq c^{\prime} \mathbf{E} f(X) .
$$

In doing so, you may need to use Lemma A. 2 (i). To prove that condition (iii) holds, write

$$
\begin{aligned}
\mathbf{E} & \left\{\max _{1 \leq i \leq n} W_{n i}(X)\right\} \\
& \leq \frac{1}{n}+\mathbf{P}\left\{X \in S^{c}\right\}+\sum_{j: A_{n j} \cap S \neq \emptyset} \mathbf{E}\left\{I_{\left\{X \in A_{n j}\right\}} \frac{1}{n \mu_{n}\left(A_{n j}\right)} I_{\left\{\mu_{n}\left(A_{n j}\right)>0\right\}}\right\}
\end{aligned}
$$

and use Lemma A. 2 (ii).
Problem 6.5. Let $\left\{a_{n}\right\}$ be a sequence of positive numbers converging to zero. Give an example of an a posteriori probability function $\eta$, and a sequence of functions $\left\{\eta_{n}\right\}$ approximating $\eta$ such that

$$
\liminf _{n \rightarrow \infty} \frac{\mathbf{P}\left\{g_{n}(X) \neq Y\right\}-L^{*}}{a_{n} \sqrt{\mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\}}}>0,
$$

where

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

Thus, the rate of convergence in Theorem 6.5 may be arbitrarily slow. Hint: Define $\eta=$ $1 / 2+b(x)$, where $b(x)$ is a very slowly increasing nonnegative function.

Problem 6.6. Let $\delta>0$, and assume that $|\eta(x)-1 / 2| \geq \delta$ for all $x$. Consider the decision

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

Prove that

$$
\mathbf{P}\left\{g_{n}(X) \neq Y\right\}-L^{*} \leq \frac{2 \mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\}}{\delta} .
$$

This shows that the rate of convergence implied by the inequality of Theorem 6.6 may be preserved for very general classes of distributions.

Problem 6.7. Assume that $L^{*}=0$, and consider the decision

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

Show that for all $1 \leq p<\infty$,

$$
\mathbf{P}\left\{g_{n}(X) \neq Y\right\} \leq 2^{p} \mathbf{E}\left\{\left|\eta_{n}(X)-\eta(X)\right|^{p}\right\} .
$$

Hint: Proceed as in the proof of Theorem 6.6, but use Hölder's inequality.
Problem 6.8. Theorem 6.5 cannot be generalized to the $L_{1}$ error. In particular, show by example that it is not always true that

$$
\lim _{n \rightarrow \infty} \frac{\mathbf{E} L_{n}-L^{*}}{\mathbf{E}\left\{\left|\eta_{n}(X)-\eta(X)\right|\right\}}=0
$$

when $\mathbf{E}\left\{\left|\eta_{n}(X)-\eta(X)\right|\right\} \rightarrow 0$ as $n \rightarrow \infty$ for some regression function estimate $\eta_{n}$. Thus, the inequality (Corollary 6.2)

$$
\mathbf{E} L_{n}-L^{*} \leq 2 \mathbf{E}\left\{\left|\eta_{n}(X)-\eta(X)\right|\right\}
$$

cannot be universally improved.

Problem 6.9. Let $\eta^{\prime}: \mathcal{R}^{d} \rightarrow[0,1]$, and define $g(x)=I_{\left\{\eta^{\prime}(x)>1 / 2\right\}}$. Assume that the random variable $X$ satisfies that $\mathbf{P}\left\{\eta^{\prime}(X)=1 / 2\right\}=0$. Let $\eta_{1}, \eta_{2}, \ldots$ be a sequence of functions such that $\mathbf{E}\left\{\left|\eta^{\prime}(X)-\eta_{n}(X)\right|\right\} \rightarrow 0$ as $n \rightarrow \infty$. Prove that $L\left(g_{n}\right) \rightarrow L(g)$ for all distributions of $(X, Y)$ satisfying the condition on $X$ above, where $g_{n}(x)=I_{\left\{\eta_{n}(x)>1 / 2\right\}}$ (Lugosi (1992)).

Problem 6.10. A lying teacher. Sometimes the training labels $Y_{1}, \ldots, Y_{n}$ are not available, but can only be observed through a noisy binary channel. Still, we want to decide on $Y$. Consider the following model. Assume that the $Y_{i}$ 's in the training data are replaced by the i.i.d. binary-valued random variables $Z_{i}$, whose distribution is given by

$$
\begin{aligned}
& \mathbf{P}\left\{Z_{i}=1 \mid Y_{i}=0, X_{i}=x\right\}=\mathbf{P}\left\{Z_{i}=1 \mid Y_{i}=0\right\}=p<1 / 2, \\
& \mathbf{P}\left\{Z_{i}=0 \mid Y_{i}=1, X_{i}=x\right\}=\mathbf{P}\left\{Z_{i}=0 \mid Y_{i}=1\right\}=q<1 / 2 .
\end{aligned}
$$

Consider the decision

$$
g(x)= \begin{cases}0 & \text { if } \eta^{\prime}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

where $\eta^{\prime}(x)=\mathbf{P}\left\{Z_{1}=1 \mid X_{1}=x\right\}$. Show that

$$
\mathbf{P}\{g(X) \neq Y\} \leq L^{*}\left(1+\frac{2|p-q|}{1-2 \max (p, q)}\right) .
$$

Use Problem 6.9 to conclude that if the binary channel is symmetric (i.e., $p=q<1 / 2$ ), and $\mathbf{P}\left\{\eta^{\prime}(X)=1 / 2\right\}=0$, then $L_{1}$-consistent estimation leads to a consistent rule, in spite of the fact that the labels $Y_{i}$ were not available in the training sequence (Lugosi (1992)).

Problem 6.11. Develop a discrimination rule which has the property

$$
\lim _{n \rightarrow \infty} \mathbf{E} L_{n}=\rho=\mathbf{E}\{\sqrt{\eta(X)(1-\eta(X))}\},
$$

for all distributions such that $X$ has a density. Note: clearly, since $\rho \geq L^{*}$, this rule is not universally consistent, but it will aid you in "visualizing" the Matushita error!

Problem 6.12. If $Z_{n}$ is binomial ( $n, p$ ) and $Z$ is Bernoulli ( $p$ ), independent of $Z_{n}$, then show that $\mathbf{P}\left\{Z_{n}>n / 2, Z=0\right\}+\mathbf{P}\left\{Z_{n} \leq n / 2, Z=1\right\}$ is nonincreasing in $n$.

Problem 6.13. Let $g_{n}$ be the histogram rule based on a fixed partition $\mathcal{P}$. Show that $g_{n}$ is smart.

Problem 6.14. Show that the kernel rule with gaussian kernel and $h=1, d=1$, is not smart (kernel rules are discussed in Chapter 10). Hint: Consider $n=1$ and $n=2$ only.

Problem 6.15. Show that the kernel rule on $\mathcal{R}$, with $K(x)=I_{[-1,1]}(x)$, and $h \downarrow 0$, such that $n h \rightarrow \infty$, is not smart.

Problem 6.16. Conjecture: no universally consistent rule is smart.
Problem 6.17. A rule $g_{n}: \mathcal{R}^{d} \times\left(\mathcal{R}^{d} \times\{0,1\}\right)^{n}$ is called symmetric if $g_{n}\left(x, D_{n}\right)=$ $g_{n}\left(x, D_{n}^{\prime}\right)$ for every $x$, and every training sequence $D_{n}$, where $D_{n}^{\prime}$ is an arbitrary permutation of the pairs ( $X_{i}, Y_{i}$ ) in $D_{n}$. Any nonsymmetric rule $g_{n}$ may be symmetrized by taking a majority vote at every $x \in \mathcal{R}^{d}$ over all $g_{n}\left(x, D_{n}^{\prime}\right)$, obtained by the $n!$ permutations of $D_{n}$. It may intuitively be expected that symmetrized rules perform better. Prove that this is false, that is, exhibit a distribution and a nonsymmetric classifier $g_{n}$ such that its expected probability of error is smaller than that of the symmetrized version of $g_{n}$. Hint: Take $g_{3}\left(x, D_{3}\right)=1-Y_{1}$.

## 7

## Slow Rates of Convergence

In this chapter we consider the general pattern recognition problem: Given the observation $X$ and the training data $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$ of independent identically distributed random variable pairs, we estimate the label $Y$ by the decision

$$
g_{n}(X)=g_{n}\left(X, D_{n}\right)
$$

The error probability is

$$
L_{n}=\mathbf{P}\left\{Y \neq g_{n}(X) \mid D_{n}\right\}
$$

Obviously, the average error probability $\mathbf{E} L_{n}=\mathbf{P}\left\{Y \neq g_{n}(X)\right\}$ is completely determined by the distribution of the pair $(X, Y)$, and the classifier $g_{n}$. We have seen in Chapter 6 that there exist classification rules such as the cubic histogram rule with properly chosen cube sizes such that $\lim _{n \rightarrow \infty} \mathbf{E} L_{n}=L^{*}$ for all possible distributions. The next question is whether there are classification rules with $\mathbf{E} L_{n}$ tending to the Bayes risk at a specified rate for all distributions. Disappointingly, such rules do not exist.

### 7.1 Finite Training Sequence

The first negative result shows that for any classification rule and for any fixed $n$, there exists a distribution such that the difference between the error probability of the rule and $L^{*}$ is larger than $1 / 4$. To explain this, note that for fixed $n$, we can find a sufficiently complex distribution for which the sample size $n$ is hopelessly small.

Theorem 7.1. (Devroye (1982B)). Let $\epsilon>0$ be an arbitrarily small number. For any integer $n$ and classification rule $g_{n}$, there exists a distribution of $(X, Y)$ with Bayes risk $L^{*}=0$ such that

$$
\mathbf{E} L_{n} \geq 1 / 2-\epsilon
$$

Proof. First we construct a family of distributions of $(X, Y)$. Then we show that the error probability of any classifier is large for at least one member of the family. For every member of the family, $X$ is uniformly distributed on the set $\{1,2, \ldots, K\}$ of positive integers

$$
p_{i}=\mathbf{P}\{X=i\}= \begin{cases}1 / K & \text { if } i \in\{1, \ldots, K\} \\ 0 & \text { otherwise }\end{cases}
$$

where $K$ is a large integer specified later. Now, the family of distributions of $(X, Y)$ is parameterized by a number $b \in[0,1)$, that is, every $b$ determines a distribution as follows. Let $b \in[0,1)$ have binary expansion $b=0 . b_{0} b_{1} b_{2} \ldots$, and define $Y=b_{X}$. As the label $Y$ is a function of $X$, there exists a perfect decision, and thus $L^{*}=0$. We show that for any decision rule $g_{n}$ there is a $b$ such that if $Y=b_{X}$, then $g_{n}$ has very poor performance. Denote the average error probability corresponding to the distribution determined by $b$, by $R_{n}(b)=\mathbf{E} L_{n}$.

The proof of the existence of a bad distribution is based on the so-called probabilistic method. Here the key trick is the randomization of $b$. Define a random variable $B$ which is uniformly distributed in $[0,1)$ and independent of $X$ and $X_{1}, \ldots, X_{n}$. Then we may compute the expected value of the random variable $R_{n}(B)$. Since for any decision rule $g_{n}$,

$$
\sup _{b \in[0,1)} R_{n}(b) \geq \mathbf{E}\left\{R_{n}(B)\right\},
$$

a lower bound for $\mathbf{E}\left\{R_{n}(B)\right\}$ proves the existence of a $b \in[0,1)$ whose corresponding error probability exceeds the lower bound.

Since $B$ is uniformly distributed in $[0,1)$, its binary extension $B=0 . B_{1} B_{2} \ldots$ is a sequence of independent binary random variables with $\mathbf{P}\left\{B_{i}=0\right\}=\mathbf{P}\left\{B_{i}=\right.$ $1\}=1 / 2$. But

$$
\begin{aligned}
\mathbf{E} & \left\{R_{n}(B)\right\} \\
& =\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y\right\} \\
& =\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq B_{X}\right\} \\
& =\mathbf{P}\left\{g_{n}\left(X, X_{1}, B_{X_{1}}, \ldots, X_{n}, B_{X_{n}}\right) \neq B_{X}\right\} \\
& =\mathbf{E}\left\{\mathbf{P}\left\{g_{n}\left(X, X_{1}, B_{X_{1}}, \ldots, X_{n}, B_{X_{n}}\right) \neq B_{X} \mid X, X_{1}, \ldots, X_{n}\right\}\right\} \\
& \geq \frac{1}{2} \mathbf{P}\left\{X \neq X_{1}, X \neq X_{2}, \ldots, X \neq X_{n}\right\},
\end{aligned}
$$

since if $X \neq X_{i}$ for all $i=1,2, \ldots, n$, then given $X, X_{1}, \ldots, X_{n}, Y=B_{X}$ is conditionally independent of $g_{n}\left(X, D_{n}\right)$ and $Y$ takes values 0 and 1 with probability
$1 / 2$. But clearly,

$$
\mathbf{P}\left\{X \neq X_{1}, X \neq X_{2}, \ldots, X \neq X_{n} \mid X\right\}=\mathbf{P}\left\{X \neq X_{1} \mid X\right\}^{n}=(1-1 / K)^{n}
$$

In summary,

$$
\sup _{b \in[0,1)} R_{n}(b) \geq \frac{1}{2}(1-1 / K)^{n} .
$$

The lower bound tends to $1 / 2$ as $K \rightarrow \infty$. $\square$

Theorem 7.1 states that even though we have rules that are universally consistent, that is, they asymptotically provide the optimal performance for any distribution, their finite sample performance is always extremely bad for some distributions. This means that no classifier guarantees that with a sample size of (say) $n=10^{8}$ we get within $1 / 4$ of the Bayes error probability for all distributions. However, as the bad distribution depends upon $n$, Theorem 7.1 does not allow us to conclude that there is one distribution for which the error probability is more than $L^{*}+1 / 4$ for all $n$. Indeed, that would contradict the very existence of universally consistent rules.

### 7.2 Slow Rates

The next question is whether a certain universal rate of convergence to $L^{*}$ is achievable for some classifier. For example, Theorem 7.1 does not exclude the existence of a classifier such that for every $n, \mathbf{E} L_{n}-L^{*} \leq c / n$ for all distributions, for some constant $c$ depending upon the actual distribution. The next negative result is that this cannot be the case. Theorem 7.2 below states that the error probability $\mathbf{E} L_{n}$ of any classifier is larger than (say) $L^{*}+c /(\log \log \log n)$ for every $n$ for some distribution, even if $c$ depends on the distribution. (This can be seen by considering that by Theorem 7.2, there exists a distribution of $(X, Y)$ such that $\mathbb{E} L_{n} \geq L^{*}+1 / \sqrt{\log \log \log n}$ for every $n$.) Moreover, there is no sequence of numbers $a_{n}$ converging to zero such that there is a classification rule with error probability below $L^{*}$ plus $a_{n}$ for all distributions.

Thus, in practice, no classifier assures us that its error probability is close to $L^{*}$, unless the actual distribution is known to be a member of a restricted class of distributions. Now, it is easily seen that in the proof of both theorems we could take $X$ to have uniform distribution on $[0,1]$, or any other density (see Problem 7.2). Therefore, putting restrictions on the distribution of $X$ alone does not suffice to obtain rate-of-convergence results. For such results, one needs conditions on the a posteriori probability $\eta(x)$ as well. However, if only training data give information about the joint distribution, then theorems with extra conditions on the distribution have little practical value, as it is impossible to detect whether, for example, the a posteriori probability $\eta(x)$ is twice differentiable or not.

Now, the situation may look hopeless, but this is not so. Simply put, the Bayes error is too difficult a target to shoot at.

Weaker versions of Theorem 7.2 appeared earlier in the literature. First Cover (1968b) showed that for any sequence of classification rules, for sequences $\left\{a_{n}\right\}$ converging to zero at arbitrarily slow algebraic rates (i.e., as $1 / n^{\delta}$ for arbitrarily small $\delta>0$ ), there exists a distribution such that $\mathbf{E} L_{n} \geq L^{*}+a_{n}$ infinitely often. Devroye (1982b) strengthened Cover's result allowing sequences tending to zero arbitrarily slowly. The next result asserts that $\mathbf{E} L_{n}>L^{*}+a_{n}$ for every $n$.

Theorem 7.2. Let $\left\{a_{n}\right\}$ be a sequence of positive numbers converging to zero with $1 / 16 \geq a_{1} \geq a_{2} \geq \ldots$. For every sequence of classification rules, there exists $a$ distribution of $(X, Y)$ with $L^{*}=0$, such that

$$
\mathbf{E} L_{n} \geq a_{n}
$$

for all $n$.
This result shows that universally good classification rules do not exist. Rate of convergence studies for particular rules must necessarily be accompanied by conditions on $(X, Y)$. That these conditions too are necessarily restrictive follows from examples suggested in Problem 7.2. Under certain regularity conditions it is possible to obtain upper bounds for the rates of convergence for the probability of error of certain rules to $L^{*}$. Then it is natural to ask what the fastest achievable rate is for the given class of distributions. A theory for regression function estimation was worked out by Stone (1982). Related results for classification were obtained by Marron (1983). In the proof of Theorem 7.2 we will need the following simple lemma:

Lemma 7.1. For any monotone decreasing sequence $\left\{a_{n}\right\}$ of positive numbers converging to zero with $a_{1} \leq 1 / 16$, a probability distribution $\left(p_{1}, p_{2}, \ldots\right)$ may be found such that $p_{1} \geq p_{2} \geq \ldots$, and for all $n$

$$
\sum_{i=n+1}^{\infty} p_{i} \geq \max \left(8 a_{n}, 32 n p_{n+1}\right)
$$

Proof. It suffices to look for $p_{i}$ 's such that

$$
\sum_{i=n+1}^{\infty} p_{i} \geq \max \left(8 a_{n}, 32 n p_{n}\right)
$$

These conditions are easily satisfied. For positive integers $u<v$, define the function $H(v, u)=\sum_{i=u}^{v-1} 1 / i$. First we find a sequence $1=n_{1}<n_{2}<\ldots$ of integers with the following properties:
(a) $\quad H\left(n_{k+1}, n_{k}\right)$ is monotonically increasing,
(b) $\quad H\left(n_{2}, n_{1}\right) \geq 32$,
(c) $\quad 8 a_{n_{k}} \leq 1 / 2^{k}$ for all $k \geq 1$.

Note that (c) may only be satisfied if $a_{n_{1}}=a_{1} \leq 1 / 16$. To this end, define constants $c_{1}, c_{2}, \ldots$ by

$$
c_{k}=\frac{32}{2^{k} H\left(n_{k+1}, n_{k}\right)}, \quad k \geq 1
$$

so that the $c_{k}$ 's are decreasing in $k$, and

$$
\frac{1}{32} \sum_{k=1}^{\infty} c_{k} H\left(n_{k+1}, n_{k}\right)=\sum_{k=1}^{\infty} \frac{1}{2^{k}}=1
$$

For $n \in\left[n_{k}, n_{k+1}\right)$, we define $p_{n}=c_{k} /(32 n)$. We claim that these numbers have the required properties. Indeed, $\left\{p_{n}\right\}$ is decreasing, and

$$
\sum_{n=1}^{\infty} p_{n}=\sum_{k=1}^{\infty} \frac{c_{k}}{32} H\left(n_{k+1}, n_{k}\right)=1
$$

Finally, if $n \in\left[n_{k}, n_{k+1}\right)$, then

$$
\sum_{i=n+1}^{\infty} p_{i} \geq \sum_{j=k+1}^{\infty} \frac{c_{j}}{32} H\left(n_{j+1}, n_{j}\right)=\sum_{j=k+1}^{\infty} \frac{1}{2^{j}}=\frac{1}{2^{k}}
$$

Clearly, on the one hand, by the monotonicity of $H\left(n_{k+1}, n_{k}\right), 1 / 2^{k} \geq c_{k}=32 n p_{n}$. On the other hand, $1 / 2^{k} \geq 8 a_{n_{k}} \geq 8 a_{n}$. This concludes the proof.

Proof of Theorem 7.2. We introduce some notation. Let $b=0 . b_{1} b_{2} b_{3} \ldots$ be a real number on $[0,1]$ with the shown binary expansion, and let $B$ be a random variable uniformly distributed on $[0,1]$ with expansion $B=0 . B_{1} B_{2} B_{3} \ldots$ Let us restrict ourselves to a random variable $X$ with

$$
\mathbf{P}\{X=i\}=p_{i}, \quad i \geq 1,
$$

where $p_{1} \geq p_{2} \geq \ldots>0$, and $\sum_{i=n+1}^{\infty} p_{i} \geq \max \left(8 a_{n}, 32 n p_{n+1}\right)$ for every $n$. That such $p_{i}$ 's exist follows from Lemma 7.1. Set $Y=b_{X}$. As $Y$ is a function of $X$, we see that $L^{*}=0$. Each $b \in[0,1$ ) however describes a different distribution. With $b$ replaced by $B$ we have a random distribution. Introduce the short notation $\Delta_{n}=\left(\left(X_{1}, B_{X_{1}}\right), \ldots,\left(X_{n}, B_{X_{n}}\right)\right)$, and define $G_{n i}=g_{n}\left(i, \Delta_{n}\right)$. If $L_{n}(B)$ denotes the probability of error $\mathbf{P}\left\{g_{n}\left(X, \Delta_{n}\right) \neq Y \mid B, X_{1}, \ldots, X_{n}\right\}$ for the random distribution, then we note that we may write

$$
L_{n}(B)=\sum_{i=1}^{\infty} p_{i} I_{\left\{G_{n i} \neq B_{i}\right\}} .
$$

If $L_{n}(b)$ is the probability of error for a distribution parametrized by $b$, then

$$
\begin{aligned}
\sup _{b} \inf _{n} \mathbf{E} \frac{L_{n}(b)}{2 a_{n}} & \geq \sup _{b} \mathbf{E}\left\{\inf _{n} \frac{L_{n}(b)}{2 a_{n}}\right\} \\
& \geq \mathbf{E}\left\{\inf _{n} \frac{L_{n}(B)}{2 a_{n}}\right\} \\
& \left.=\mathbf{E}\left\{\mathbf{E}\left\{\inf _{n} \frac{L_{n}(B)}{2 a_{n}}\right\} X_{1}, X_{2}, \ldots\right\}\right\}
\end{aligned}
$$

We consider only the conditional expectation for now. We have

$$
\begin{aligned}
& \mathbf{E}\left\{\left.\inf _{n} \frac{L_{n}(B)}{2 a_{n}} \right\rvert\, X_{1}, X_{2}, \ldots\right\} \\
& \geq \mathbf{P}\left\{\bigcap_{n=1}^{\infty}\left\{L_{n}(B) \geq 2 a_{n}\right\} \mid X_{1}, X_{2}, \ldots\right\} \\
& \geq 1-\sum_{n=1}^{\infty} \mathbf{P}\left\{L_{n}(B)<2 a_{n} \mid X_{1}, X_{2}, \ldots\right\} \\
& \\
& =1-\sum_{n=1}^{\infty} \mathbf{P}\left\{L_{n}(B)<2 a_{n} \mid X_{1}, X_{2}, \ldots, X_{n}\right\} \\
& \\
& =1-\sum_{n=1}^{\infty} \mathbf{E}\left\{\mathbf{P}\left\{L_{n}(B)<2 a_{n} \mid \Delta_{n}\right\} \mid X_{1}, X_{2}, \ldots, X_{n}\right\}
\end{aligned}
$$

We bound the conditional probabilities inside the sum:

$$
\begin{aligned}
& \mathbf{P}\left\{L_{n}(B)<2 a_{n} \mid \Delta_{n}\right\} \\
& \leq \mathbf{P}\left\{\sum_{i \notin\left\{X_{1}, \ldots, X_{n}\right\}} p_{i} I_{\left\{G_{n i} \neq B_{i}\right\}}<2 a_{n} \mid \Delta_{n}\right\} \\
& \text { (and, noting that } G_{n i}, X_{1}, \ldots, X_{n} \\
& \text { are all functions of } \Delta_{n}, \text { we have:) } \\
&= \mathbf{P}\left\{\sum_{i \notin\left\{X_{1}, \ldots, X_{n}\right\}} p_{i} I_{\left\{B_{i}=1\right\}}<2 a_{n} \mid \Delta_{n}\right\} \\
& \leq \mathbf{P}\left\{\sum_{i=n+1}^{\infty} p_{i} I_{\left\{B_{i}=1\right\}}<2 a_{n}\right\} \\
&\left(\text { since the } p_{i}, s\right. \text { are decreasing, by stochastic dominance) } \\
&= \mathbf{P}\left\{\sum_{i=n+1}^{\infty} p_{i} B_{i}<2 a_{n}\right\} .
\end{aligned}
$$

Now everything boils down to bounding these probabilities from above. We proceed by Chernoff's bounding technique. The idea is the following: For any random variable $X$, and $s>0$, by Markov's inequality,

$$
\mathbf{P}\{X \geq \epsilon\}=\mathbf{P}\left\{e^{s X} \geq e^{s \epsilon}\right\} \leq \frac{\mathbf{E}\left\{e^{s X}\right\}}{e^{s \epsilon}}
$$

By cleverly choosing $s$ one can often obtain very sharp bounds. For more discussion and examples of Chernoff's method, refer to Chapter 8. In our case,

$$
\begin{aligned}
& \mathbf{P}\left\{\sum_{i=n+1}^{\infty} p_{i} B_{i}<2 a_{n}\right\} \\
& \leq \mathbf{E}\left\{e^{2 s a_{n}-s \sum_{i=n+1}^{\infty} p_{i} B_{i}}\right\} \\
&= e^{2 s a_{n}} \prod_{i=n+1}^{\infty}\left(\frac{1}{2}+\frac{1}{2} e^{-s p_{i}}\right) \\
& \leq e^{2 s a_{n}} \prod_{i=n+1}^{\infty} \frac{1}{2}\left(2-s p_{i}+\frac{s^{2} p_{i}^{2}}{2}\right) \\
&\left(\text { since } e^{-x} \leq 1-x+x^{2} / 2 \text { for } x \geq 0\right) \\
& \leq \exp \left(2 s a_{n}+\sum_{i=n+1}^{\infty}\left(-\frac{s p_{i}}{2}+\frac{s^{2} p_{i}^{2}}{4}\right)\right) \\
&\left(\operatorname{since} 1-x \leq e^{-x}\right) \\
& \leq \exp \left(2 s a_{n}-\frac{s \Sigma}{2}+\frac{s^{2} p_{n+1} \Sigma}{4}\right) \\
&\left(\text { where } \Sigma=\sum_{i=n+1}^{\infty} p_{i}\right) \\
&= \exp \left(-\frac{1}{4} \frac{\left(4 a_{n}-\Sigma\right)^{2}}{\Sigma p_{n+1}}\right) \\
& \leq \exp \left(-\frac{1}{16} \frac{\Sigma}{p_{n+1}}\right) \quad\left(\text { since } \Sigma \geq 8 a_{n}\right) \\
& \leq e^{-2 n}\left(\operatorname{since} \Sigma \geq 32 p_{n+1} n\right) .
\end{aligned}
$$

Thus, we conclude that

$$
\sup _{b} \inf _{n} \mathbf{E} \frac{L_{n}(b)}{2 a_{n}} \geq 1-\sum_{n=1}^{\infty} e^{-2 n}=\frac{e^{2}-2}{e^{2}-1}>\frac{1}{2}
$$

so that there exists a $b$ for which $\mathbf{E} L_{n}(b) \geq a_{n}$ for all $n$.

## Problems and Exercises

Problem 7.1. Extend Theorem 7.2 for distributions with $0<L^{*}<1 / 2$ : show that if $a_{n}$ is a sequence of positive numbers as in Theorem 7.2, then for any classification rule there is a distribution such that $\mathbf{E} L_{n}-L^{*} \geq a_{n}$ for every $n$ for which $L^{*}+a_{n}<1 / 2$.

Problem 7.2. Prove Theorems 7.1 and 7.2 , under one of the following additional assumptions, which make the case that one will need very restrictive conditions indeed to study rates of convergence.
(1) $X$ has a uniform density on $[0,1)$.
(2) $X$ has a uniform density on $[0,1)$ and $\eta$ is infinitely many times continuously differentiable on $[0,1)$.
(3) $\eta$ is unimodal in $x \in \mathcal{R}^{2}$, that is, $\eta(\lambda x)$ decreases as $\lambda>0$ increases for any $x \in \mathcal{R}^{2}$.
(4) $\eta$ is $\{0,1\}$-valued, $X$ is $\mathcal{R}^{2}$-valued, and the set $\{x: \eta(x)=1\}$ is a compact convex set containing the origin.

Problem 7.3. There is no super-classifier. Show that for every sequence of classification rules $\left\{g_{n}\right\}$ there is a universally consistent sequence of rules $\left\{g_{n}^{\prime}\right\}$, such that for some distribution of ( $X, Y$ ),

$$
\mathbf{P}\left\{g_{n}(X) \neq Y\right\}>\mathbf{P}\left\{g_{n}^{\prime}(X) \neq Y\right\}
$$

for all $n$.
Problem 7.4. The next two exercises are intended to demonstrate that the weaponry of pattern recognition can often be successfully used for attacking other statistical problems. For example, a consequence of Theorem 7.2 is that estimating infinite discrete distributions is hard. Consider the problem of estimating a distribution ( $p_{1}, p_{2}, \ldots$ ) on the positive integers $\{1,2,3, \ldots\}$ from a sample $X_{1}, \ldots, X_{n}$ of i.i.d. random variables with $\mathbf{P}\left\{X_{1}=i\right\}=p_{i}$, $i \geq 1$. Show that for any decreasing sequence $\left\{a_{n}\right\}$ of positive numbers converging to zero with $a_{1} \leq 1 / 16$, and any estimate $\left\{p_{i, n}\right\}$, there exists a distribution such that

$$
\mathbf{E}\left\{\sum_{i=1}^{\infty}\left|p_{i}-p_{i, n}\right|\right\} \geq a_{n}
$$

Hint: Consider a classification problem with $L^{*}=0, \mathbf{P}\{Y=0\}=1 / 2$, and $X$ concentrated on $\{1,2, \ldots\}$. Assume that the class-conditional probabilities $p_{i}^{(0)}=\mathbf{P}\{X=i \mid Y=0\}$ and $p_{i}^{(1)}=\mathbf{P}\{X=i \mid Y=1\}$ are estimated from two i.i.d. samples $X_{1}^{(0)}, \ldots, X_{n}^{(0)}$ and $X_{1}^{(1)}, \ldots, X_{n}^{(1)}$, distributed according to $\left\{p_{i}^{(0)}\right\}$ and $\left\{p_{i}^{(1)}\right\}$, respectively. Use Theorem 2.3 to show that for the classification rule obtained from these estimates in a natural way,

$$
\mathbf{E} L_{n} \leq \frac{1}{2} \mathbf{E}\left\{\sum_{i=1}^{\infty}\left|p_{i}^{(0)}-p_{i, n}^{(0)}\right|+\sum_{i=1}^{\infty}\left|p_{i}^{(1)}-p_{i, n}^{(1)}\right|\right\},
$$

therefore the lower bound of Theorem 7.2 can be applied.
Problem 7.5. A similar slow-rate result appears in density estimation. Consider the problem of estimating a density $f$ on $\mathcal{R}$, from an i.i.d. sample $X_{1}, \ldots, X_{n}$ having density $f$. Show that for any decreasing sequence $\left\{a_{n}\right\}$ of positive numbers converging to zero with $a_{1} \leq 1 / 16$, and any density estimate $f_{n}$, there exists a distribution such that

$$
\mathbf{E}\left\{\int\left|f(x)-f_{n}(x)\right| d x\right\} \geq a_{n}
$$

This result was proved by Birgé (1986) using a different-and in our view much more complicated-argument. HINT: Put $p_{i}=\int_{i}^{i+1} f(x) d x$ and $p_{i, n}=\int_{i}^{i+1} f_{n}(x) d x$ and apply Problem 7.4.

## 8

## Error Estimation

### 8.1 Error Counting

Estimating the error probability $L_{n}=\mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}$ of a classification function $g_{n}$ is of essential importance. The designer always wants to know what performance can be expected from a classifier. As the designer does not know the distribution of the data-otherwise there would not be any need to design a classifier-it is important to find error estimation methods that work well without any condition on the distribution of ( $X, Y$ ). This motivates us to search for distribution-free performance bounds for error estimation methods.

Suppose that we want to estimate the error probability of a classifier $g_{n}$ designed from the training sequence $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$. Assume first that a testing sequence

$$
T_{m}=\left(\left(X_{n+1}, Y_{n+1}\right), \ldots,\left(X_{n+m}, Y_{n+m}\right)\right)
$$

is available, which is a sequence of i.i.d. pairs that are independent of $(X, Y)$ and $D_{n}$, and that are distributed as $(X, Y)$. An obvious way to estimate $L_{n}$ is to count the number of errors that $g_{n}$ commits on $T_{m}$. The error-counting estimator $\widehat{L}_{n, m}$ is defined by the relative frequency

$$
\widehat{L}_{n, m}=\frac{1}{m} \sum_{j=1}^{m} I_{\left\{g_{n}\left(X_{n+j}\right) \neq Y_{n+j}\right\}} .
$$

The estimator is clearly unbiased in the sense that

$$
\mathbf{E}\left\{\widehat{L}_{n, m} \mid D_{n}\right\}=L_{n}
$$

and the conditional distribution of $m \widehat{L}_{n, m}$, given the training data $D_{n}$, is binomial with parameters $m$ and $L_{n}$. This makes analysis easy, for properties of the binomial distribution are well known. One main tool in the analysis is Hoeffding's inequality, which we will use many many times throughout this book.

### 8.2 Hoeffding's Inequality

The following inequality was proved for binomial random variables by Chernoff (1952) and Okamoto (1958). The general format is due to Hoeffding (1963):

Theorem 8.1. (Hoeffding (1963)). Let $X_{1}, \ldots, X_{n}$ be independent bounded random variables such that $X_{i}$ falls in the interval $\left[a_{i}, b_{i}\right]$ with probability one. Denote their sum by $S_{n}=\sum_{i=1}^{n} X_{i}$. Then for any $\epsilon>0$ we have

$$
\mathbf{P}\left\{S_{n}-\mathbf{E} S_{n} \geq \epsilon\right\} \leq e^{-2 \epsilon^{2} / \sum_{i=1}^{n}\left(b_{i}-a_{i}\right)^{2}}
$$

and

$$
\mathbf{P}\left\{S_{n}-\mathbf{E} S_{n} \leq-\epsilon\right\} \leq e^{-2 \epsilon^{2} / \sum_{i=1}^{n}\left(b_{i}-a_{i}\right)^{2}}
$$

The proof uses a simple auxiliary inequality:
Lemma 8.1. Let $X$ be a random variable with $\mathbf{E} X=0, a \leq X \leq b$. Then for $s>0$,

$$
\mathbf{E}\left\{e^{s X}\right\} \leq e^{s^{2}(b-a)^{2} / 8}
$$

Proof. Note that by convexity of the exponential function

$$
e^{s x} \leq \frac{x-a}{b-a} e^{s b}+\frac{b-x}{b-a} e^{s a} \quad \text { for } a \leq x \leq b
$$

Exploiting $\mathbf{E} X=0$, and introducing the notation $p=-a /(b-a)$ we get

$$
\begin{aligned}
\mathbf{E} e^{s X} & \leq \frac{b}{b-a} e^{s a}-\frac{a}{b-a} e^{s b} \\
& =\left(1-p+p e^{s(b-a)}\right) e^{-p s(b-a)} \\
& \stackrel{\operatorname{def}}{=} e^{\phi(u)},
\end{aligned}
$$

where $u=s(b-a)$, and $\phi(u)=-p u+\log \left(1-p+p e^{u}\right)$. But by straightforward calculation it is easy to see that the derivative of $\phi$ is

$$
\phi^{\prime}(u)=-p+\frac{p}{p+(1-p) e^{-u}}
$$

therefore $\phi(0)=\phi^{\prime}(0)=0$. Moreover,

$$
\phi^{\prime \prime}(u)=\frac{p(1-p) e^{-u}}{\left(p+(1-p) e^{-u}\right)^{2}} \leq \frac{1}{4} .
$$

Thus, by Taylor series expansion with remainder, for some $\theta \in[0, u]$,

$$
\phi(u)=\phi(0)+u \phi^{\prime}(0)+\frac{u^{2}}{2} \phi^{\prime \prime}(\theta) \leq \frac{u^{2}}{8}=\frac{s^{2}(b-a)^{2}}{8} .
$$

Proof of Theorem 8.1. The proof is based on Chernoff's bounding method (Chernoff (1952)): by Markov's inequality, for any nonnegative random variable $X$, and any $\epsilon>0$,

$$
\mathbf{P}\{X \geq \epsilon\} \leq \frac{\mathbf{E} X}{\epsilon} .
$$

Therefore, if $s$ is an arbitrary positive number, then for any random variable $X$,

$$
\mathbf{P}\{X \geq \epsilon\}=\mathbf{P}\left\{e^{s X} \geq e^{s \epsilon}\right\} \leq \frac{\mathbf{E} e^{s X}}{e^{s \epsilon}} .
$$

In Chernoff's method, we find an $s>0$ that minimizes the upper bound or makes the upper bound small. In our case, we have

$$
\begin{aligned}
& \mathbf{P}\left\{S_{n}-\mathbf{E} S_{n} \geq \epsilon\right\} \\
& \leq e^{-s \epsilon} \mathbf{E}\left\{\exp \left(s \sum_{i=1}^{n}\left(X_{i}-\mathbf{E} X_{i}\right)\right)\right\} \\
&=e^{-s \epsilon} \prod_{i=1}^{n} \mathbf{E}\left\{e^{s\left(X_{i}-\mathbf{E} X_{i}\right)}\right\} \quad \text { (by independence) } \\
& \leq e^{-s \epsilon} \prod_{i=1}^{n} e^{s^{2}\left(b_{i}-a_{i}\right)^{2} / 8} \quad \text { (by Lemma 8.1) } \\
&=e^{-s \epsilon} e^{s^{2} \sum_{i=1}^{n}\left(b_{i}-a_{i}\right)^{2} / 8} \\
&\left.=e^{-2 \epsilon^{2} / \sum_{i=1}^{n}\left(b_{i}-a_{i}\right)^{2}} \quad \text { (by choosing } s=4 \epsilon / \sum_{i=1}^{n}\left(b_{i}-a_{i}\right)^{2}\right) .
\end{aligned}
$$

The second inequality is proved analogously.
The two inequalities in Theorem 8.1 may be combined to get

$$
\mathbf{P}\left\{\left|S_{n}-\mathbf{E} S_{n}\right| \geq \epsilon\right\} \leq 2 e^{-2 \epsilon^{2} / \sum_{i=1}^{n}\left(b_{i}-a_{i}\right)^{2}} .
$$

Now, we can apply this inequality to get a distribution-free performance bound for the counting error estimate:

Corollary 8.1. For every $\epsilon>0$,

$$
\mathbf{P}\left\{\left|\widehat{L}_{n, m}-L_{n}\right|>\epsilon \mid D_{n}\right\} \leq 2 e^{-2 m \epsilon^{2}} .
$$

The variance of the estimate can easily be computed using the fact that, conditioned on the data $D_{n}, m \widehat{L}_{n, m}$ is binomially distributed:

$$
\mathbf{E}\left\{\left(\widehat{L}_{n, m}-L_{n}\right)^{2} \mid D_{n}\right\}=\frac{L_{n}\left(1-L_{n}\right)}{m} \leq \frac{1}{4 m}
$$

These are just the types of inequalities we want, for these are valid for any distribution and data size, and the bounds do not even depend on $g_{n}$.

Consider a special case in which all the $X_{i}$ 's take values on $[-c, c]$ and have zero mean. Then Hoeffding's inequality states that

$$
\mathbf{P}\left\{S_{n} / n>\epsilon\right\} \leq e^{-n \epsilon^{2} /\left(2 c^{2}\right)}
$$

This bound, while useful for $\epsilon$ larger than $c / \sqrt{n}$, ignores variance information. When $\operatorname{Var}\left\{X_{i}\right\} \ll c^{2}$, it is indeed possible to outperform Hoeffding's inequality. In particular, we have:

Theorem 8.2. (Bennett (1962) and Bernstein (1946)). Let $X_{1}, \ldots, X_{n}$ be independent real-valued random variables with zero mean, and assume that $X_{i} \leq c$ with probability one. Let

$$
\sigma^{2}=\frac{1}{n} \sum_{i=1}^{n} \operatorname{Var}\left\{X_{i}\right\}
$$

Then, for any $\epsilon>0$,

$$
\mathbf{P}\left\{\frac{1}{n} \sum_{i=1}^{n} X_{i}>\epsilon\right\} \leq \exp \left(-\frac{n \epsilon}{2 c}\left(\left(1+\frac{\sigma^{2}}{2 c \epsilon}\right) \log \left(1+\frac{2 c \epsilon}{\sigma^{2}}\right)-1\right)\right)
$$

(Bennett (1962)), and

$$
\mathbf{P}\left\{\frac{1}{n} \sum_{i=1}^{n} X_{i}>\epsilon\right\} \leq \exp \left(-\frac{n \epsilon^{2}}{2 \sigma^{2}+2 c \epsilon / 3}\right)
$$

(Bernstein (1946)).
The proofs are left as exercises (Problem 8.2). We note that Bernstein's inequality kicks in when $\epsilon$ is larger than about $\max (\sigma / \sqrt{n}, c / \sqrt{n})$. It is typically better than Hoeffding's inequality when $\sigma \ll c$.

### 8.3 Error Estimation Without Testing Data

A serious problem concerning the practical applicability of the estimate introduced above is that it requires a large, independent testing sequence. In practice, however, an additional sample is rarely available. One usually wants to incorporate all available $\left(X_{i}, Y_{i}\right)$ pairs in the decision function. In such cases, to estimate $L_{n}$, we have to rely on the training data only. There are well-known methods that we
will discuss later that are based on cross-validation (or leave-one-out) (Lunts and Brailovsky (1967); Stone (1974)); and holdout, resubstitution, rotation, smoothing, and bootstrapping (Efron (1979), (1983)), which may be employed to construct an empirical risk from the training sequence, thus obviating the need for a testing sequence. (See Kanal (1974), Cover and Wagner (1975), Toussaint (1974a), Glick (1978), Hand (1986), Jain, Dubes, and Chen (1987), and McLachlan (1992) for surveys, discussion, and empirical comparison.)

Analysis of these methods, in general, is clearly a much harder problem, as $g_{n}$ can depend on $D_{n}$ in a rather complicated way. If we construct some estimator $\widehat{L}_{n}$ from $D_{n}$, then it would be desirable to obtain distribution-free bounds on

$$
\mathbf{P}\left\{\left|\widehat{L}_{n}-L_{n}\right|>\epsilon\right\}
$$

or on

$$
\mathbf{E}\left\{\left|\widehat{L}_{n}-L_{n}\right|^{q}\right\}
$$

for some $q \geq 1$. Conditional probabilities and expectations given $D_{n}$ are meaningless, since everything is a function of $D_{n}$. Here, however, we have to be much more careful as we do not want $\widehat{L}_{n}$ to be optimistically biased because the same data are used both for training and testing.

Distribution-free bounds for the above quantities would be extremely helpful, as we usually do not know the distribution of $(X, Y)$. While for some rules such estimates exist-we will exhibit several avenues in Chapters 22, 23, 24, 25, 26, and 31 -it is disappointing that a single error estimation method cannot possibly work for all discrimination rules. It is therefore important to point out that we have to consider ( $g_{n}, \widehat{L}_{n}$ ) pairs-for every rule one or more error estimates must be found if possible, and vice versa, for every error estimate, its limitations have to be stated. Secondly, rules for which no good error estimates are known should be avoided. Luckily, most popular rules do not fall into this category. On the other hand, proven distribution-free performance guarantees are rarely available-see Chapters 23 and 24 for examples.

### 8.4 Selecting Classifiers

Probably the most important application of error estimation is in the selection of a classification function from a class $\mathcal{C}$ of functions. If a class $\mathcal{C}$ of classifiers is given, then it is tempting to pick the one that minimizes an estimate of the error probability over the class. A good method should pick a classifier with an error probability that is close to the minimal error probability in the class. Here we require much more than distribution-free performance bounds of the error estimator for each of the classifiers in the class. Problem 8.8 demonstrates that it is not sufficient to be able to estimate the error probability of all classifiers in the class. Intuitively, if we can estimate the error probability for the classifiers in $\mathcal{C}$ uniformly well, then the classification function that minimizes the estimated error probability is likely to have an error probability that is close to the best in the class. To certify this
intuition, consider the following situation: Let $\mathcal{C}$ be a class of classifiers, that is, a class of mappings of the form $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$. Assume that the error count

$$
\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{j=1}^{n} I_{\left\{\phi\left(X_{j}\right) \neq Y_{j}\right\}}
$$

is used to estimate the error probability $L(\phi)=\mathbf{P}\{\phi(X) \neq Y\}$ of each classifier $\phi \in \mathcal{C}$. Denote by $\phi_{n}^{*}$ the classifier that minimizes the estimated error probability over the class:

$$
\widehat{L}_{n}\left(\phi_{n}^{*}\right) \leq \widehat{L}_{n}(\phi) \quad \text { for all } \phi \in \mathcal{C}
$$

Then for the error probability

$$
L\left(\phi_{n}^{*}\right)=\mathbf{P}\left\{\phi_{n}^{*}(X) \neq Y \mid D_{n}\right\}
$$

of the selected rule we have:
Lemma 8.2. (Vapnik and Chervonenkis (1974C); SEE also Devroye (1988b)).

$$
\begin{aligned}
& L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi) \leq 2 \sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|, \\
& \left|\widehat{L}_{n}\left(\phi_{n}^{*}\right)-L\left(\phi_{n}^{*}\right)\right| \leq \sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|
\end{aligned}
$$

Proof.

$$
\begin{aligned}
L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi) & =L\left(\phi_{n}^{*}\right)-\widehat{L}_{n}\left(\phi_{n}^{*}\right)+\widehat{L}_{n}\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi) \\
& \leq L\left(\phi_{n}^{*}\right)-\widehat{L}_{n}\left(\phi_{n}^{*}\right)+\sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right| \\
& \leq 2 \sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|
\end{aligned}
$$

The second inequality is trivially true.
We see that upper bounds for $\sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|$ provide us with upper bounds for two things simultaneously:
(1) An upper bound for the suboptimality of $\phi_{n}^{*}$ within $\mathcal{C}$, that is, a bound for $L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)$.
(2) An upper bound for the error $\left|\widehat{L}_{n}\left(\phi_{n}^{*}\right)-L\left(\phi_{n}^{*}\right)\right|$ committed when $\widehat{L}_{n}\left(\phi_{n}^{*}\right)$ is used to estimate the probability of error $L\left(\phi_{n}^{*}\right)$ of the selected rule.

In other words, by bounding $\sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|$, we kill two flies at once. It is particularly useful to know that even though $\widehat{L}_{n}\left(\phi_{n}^{*}\right)$ is usually optimistically biased, it is within given bounds of the unknown probability of error with $\phi_{n}^{*}$, and that no other test sample is needed to estimate this probability of error. Whenever
our bounds indicate that we are close to the optimum in $\mathcal{C}$, we must at the same time have a good estimate of the probability of error, and vice versa.

As a simple, but interesting application of Lemma 8.2 we consider the case when the class $\mathcal{C}$ contains finitely many classifiers.

Theorem 8.3. Assume that the cardinality of $\mathcal{C}$ is bounded by $N$. Then we have for all $\epsilon>0$,

$$
\mathbf{P}\left\{\sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|>\epsilon\right\} \leq 2 N e^{-2 n \epsilon^{2}}
$$

Proof.

$$
\begin{aligned}
\mathbf{P}\left\{\sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|>\epsilon\right\} & \leq \sum_{\phi \in \mathcal{C}} \mathbf{P}\left\{\left|\widehat{L}_{n}(\phi)-L(\phi)\right|>\epsilon\right\} \\
& \leq 2 N e^{-2 n \epsilon^{2}}
\end{aligned}
$$

where we used Hoeffding's inequality, and the fact that the random variable $n \widehat{L}_{n}(\phi)$ is binomially distributed with parameters $n$ and $L(\phi)$.

Remark. Distribution-free properties. Theorem 8.3 shows that the problem studied here is purely combinatorial. The actual distribution of the data does not play a role at all in the upper bounds.

REMARK. Without testing data. Very often, a class of rules $\mathcal{C}$ of the form $\phi_{n}(x)=$ $\phi_{n}\left(x, D_{n}\right)$ is given, and the same data $D_{n}$ are used to select a rule by minimizing some estimates $\widehat{L}_{n}\left(\phi_{n}\right)$ of the error probabilities $L\left(\phi_{n}\right)=\mathbf{P}\left\{\phi_{n}(X) \neq Y \mid D_{n}\right\}$. A similar analysis can be carried out in this case. In particular, if $\phi_{n}^{*}$ denotes the selected rule, then we have similar to Lemma 8.2:

## Theorem 8.4.

$$
L\left(\phi_{n}^{*}\right)-\inf _{\phi_{n} \in \mathcal{C}} L\left(\phi_{n}\right) \leq 2 \sup _{\phi_{n} \in \mathcal{C}}\left|\widehat{L}_{n}\left(\phi_{n}\right)-L\left(\phi_{n}\right)\right|,
$$

and

$$
\left|\widehat{L}_{n}\left(\phi_{n}^{*}\right)-L\left(\phi_{n}^{*}\right)\right| \leq \sup _{\phi_{n} \in \mathcal{C}}\left|\widehat{L}_{n}\left(\phi_{n}\right)-L\left(\phi_{n}\right)\right|
$$

If $\mathcal{C}$ is finite, then again, similar to Theorem 8.3 , we have for example

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-\inf _{\phi_{n} \in \mathcal{C}} L\left(\phi_{n}\right)>\epsilon\right\} \leq N \sup _{\phi_{n} \in \mathcal{C}} \mathbf{P}\left\{\left|\widehat{L}_{n}\left(\phi_{n}\right)-L\left(\phi_{n}\right)\right|>\epsilon / 2\right\}
$$

### 8.5 Estimating the Bayes Error

It is also important to have a good estimate of the optimal error probability $L^{*}$. First of all, if $L^{*}$ is large, we would know beforehand that any rule is going to perform poorly. Then perhaps the information might be used to return to the feature selection stage. Also, a comparison of estimates of $L_{n}$ and $L^{*}$ gives us an idea how much room is left for improvement. Typically, $L^{*}$ is estimated by an estimate of the error probability of some consistent classification rule (see Fukunaga and Kessel (1971), Chen and Fu (1973), Fukunaga and Hummels (1987), and Garnett and Yau (1977)). Clearly, if the estimate $\widehat{L}_{n}$ we use is consistent in the sense that $\widehat{L}_{n}-L_{n} \rightarrow 0$ with probability one as $n \rightarrow \infty$, and the rule is strongly consistent, then

$$
\widehat{L}_{n} \rightarrow L^{*}
$$

with probability one. In other words, we have a consistent estimate of the Bayes error probability. There are two problems with this approach. The first problem is that if our purpose is comparing $L^{*}$ with $L_{n}$, then using the same estimate for both of them does not give any information. The other problem is that even though for many classifiers, $L_{n}-L_{n}$ can be guaranteed to converge to zero rapidly, regardless what the distribution of ( $X, Y$ ) is (see Chapters 23 and 24), in view of the results of Chapter 7 , the rate of convergence of $\widehat{L}_{n}$ to $L^{*}$ using such a method may be arbitrarily slow. Thus, we cannot expect good performance for all distributions from such a method. The question is whether it is possible to come up with a method of estimating $L^{*}$ such that the difference $\widehat{L}_{n}-L^{*}$ converges to zero rapidly for all distributions. Unfortunately, there is no method that guarantees a certain finite sample performance for all distributions. This disappointing fact is reflected in the following negative result:

Theorem 8.5. For every n, for any estimate $\widehat{L}_{n}$ of the Bayes error probability $L^{*}$, and for every $\epsilon>0$, there exists a distribution of $(X, Y)$, such that

$$
\mathbf{E}\left\{\left|\widehat{L}_{n}-L^{*}\right|\right\} \geq \frac{1}{4}-\epsilon
$$

Proof. For a fixed $n$, we construct a family $\mathcal{F}$ of distributions, and show that for at least one member of the family, $\mathbf{E}\left\{\left|\widehat{L}_{n}-L^{*}\right|\right\} \geq \frac{1}{4}-\epsilon$. The family contains $2^{m}+1$ distributions, where $m$ is a large integer specified later. In all cases, $X_{1}, \ldots, X_{n}$ are drawn independently by a uniform distribution from the set $\{1, \ldots, m\}$. Let $B_{0}, B_{1}, B_{2}, \ldots, B_{n}$ be i.i.d. Bernoulli random variables, independent of the $X_{i}$ 's, with $\mathbf{P}\left\{B_{i}=0\right\}=\mathbf{P}\left\{B_{i}=1\right\}=1 / 2$. For the first member of the family $\mathcal{F}$, let $Y_{i}=B_{i}$ for $i=1, \ldots, n$. Thus, for this distribution, $L^{*}=1 / 2$. The Bayes error for the other $2^{m}$ members of the family is zero. These distributions are determined by $m$ binary parameters $a_{1}, a_{2}, \ldots, a_{m} \in\{0,1\}$ as follows:

$$
\eta(i)=\mathbf{P}\{Y=1 \mid X=i\}=a_{i}
$$

In other words, $Y_{i}=a_{X_{i}}$ for every $i=1, \ldots, n$. Clearly, $L^{*}=0$ for these distributions. Note also that all distributions with $X$ distributed uniformly on $\{1, \ldots, m\}$
and $L^{*}=0$ are members of the family. Just as in the proofs of Theorems 7.1 and 7.2, we randomize over the family $\mathcal{F}$ of distributions. However, the way of randomization is different here. The trick is to use $B_{0}, B_{1}, B_{2}, \ldots, B_{n}$ in randomly picking a distribution. (Recall that these random variables are just the labels $Y_{1}, \ldots, Y_{n}$ in the training sequence for the first distribution in the family.) We choose a distribution randomly, as follows: If $B_{0}=0$, then we choose the first member of $\mathcal{F}$ (the one with $L^{*}=1 / 2$ ). If $B_{0}=1$, then the labels of the training sequence are given by

$$
Y_{i}= \begin{cases}B_{i} & \text { if } X_{i} \neq X_{1}, X_{i} \neq X_{2}, \ldots, X_{i} \neq X_{i-1} \\ B_{j} & \text { if } j<i \text { is the smallest index such that } X_{i}=X_{j}\end{cases}
$$

Note that in case of $B_{0}=1$, for any fixed realization $b_{1}, \ldots, b_{n} \in\{0,1\}$ of $B_{1}, \ldots, B_{n}$, the Bayes risk is zero. Therefore, the distribution is in the family $\mathcal{F}$.

Now, let $A$ be the event that all the $X_{i}$ 's are different. Observe that under $A, \widehat{L}_{n}$ is a function of $X_{1}, \ldots, X_{n}, B_{1}, \ldots, B_{n}$ only, but not $B_{0}$. Therefore,

$$
\begin{aligned}
\sup _{\mathcal{F}} \mathbf{E}\left\{\left|\widehat{L}_{n}-L^{*}\right|\right\} \geq & \mathbf{E}\left\{\left|\widehat{L}_{n}-L^{*}\right|\right\} \\
& \left(\text { with } B_{0}, B_{1}, \ldots, B_{n} \text { random }\right) \\
\geq & \mathbf{E}\left\{I_{A}\left|\widehat{L}_{n}-L^{*}\right|\right\} \\
= & \mathbf{E}\left\{I_{A}\left(I_{\left\{B_{0}=0\right\}}\left|\widehat{L}_{n}-\frac{1}{2}\right|+I_{\left\{B_{0}=1\right\}}\left|\widehat{L}_{n}-0\right|\right)\right\} \\
= & \mathbf{E}\left\{I_{A} \frac{1}{2}\left(\left|\widehat{L}_{n}-\frac{1}{2}\right|+\left|\widehat{L}_{n}-0\right|\right)\right\} \\
\geq & \mathbf{E}\left\{I_{A} \frac{1}{4}\right\} \\
= & \frac{1}{4} \mathbf{P}\{A\}
\end{aligned}
$$

Now, if we pick $m$ large enough, $\mathbf{P}\{A\}$ can be as close to 1 as desired. Hence,

$$
\sup \mathbf{E}\left\{\left|\widehat{L}_{n}-L^{*}\right|\right\} \geq \frac{1}{4}
$$

where the supremum is taken over all distributions of $(X, Y)$.

## Problems and Exercises

Problem 8.1. Let $B$ be a binomial random variable with parameters $n$ and $p$. Show that

$$
\mathbf{P}\{B>\epsilon\} \leq e^{\epsilon-n p-\epsilon \log (\epsilon / n p)}(\epsilon>n p)
$$

and

$$
\mathbf{P}\{B<\epsilon\} \leq e^{\epsilon-n p-\epsilon \log (\epsilon / n p)}(\epsilon<n p) .
$$

(Chernoff (1952).) Hint: Proceed by Chernoff's bounding method.

Problem 8.2. Prove the inequalities of Bennett and Bernstein given in Theorem 8.2. To help you, we will guide you through different stages:
(1) Show that for any $s>0$, and any random variable $X$ with $\mathbf{E} X=0, \mathbf{E} X^{2}=\sigma^{2}$, $X \leq c$,

$$
\mathbf{E}\left\{e^{s X}\right\} \leq e^{f\left(\sigma^{2} / c^{2}\right)}
$$

where

$$
f(u)=\log \left(\frac{1}{1+u} e^{-c s u}+\frac{u}{1+u} e^{c s}\right) .
$$

(2) Show that $f^{\prime \prime}(u) \leq 0$ for $u \geq 0$.
(3) By Chernoff's bounding method, show that

$$
\mathbf{P}\left\{\frac{1}{n} \sum_{i=1}^{n} X_{i} \geq \epsilon\right\} \leq e^{-s n \epsilon+\sum_{i=1}^{n} f\left(\operatorname{Var}\left\{X_{i}\right\} / c^{2}\right)}
$$

(4) Show that $f(u) \leq f(0)+u f^{\prime}(0)=\left(e^{c s}-1-c s\right) u$.
(5) Using the bound of (4), find the optimal value of $s$ and derive Bennett's inequality.

Problem 8.3. Use Bernstein's inequality to show that if $B$ is a binomial ( $n, p$ ) random variable, then for $\epsilon>0$,

$$
\mathbf{P}\left\{\frac{B}{n}-p>\epsilon p\right\} \leq e^{-3 n p \epsilon^{2} / 8}
$$

and

$$
\mathbf{P}\left\{\frac{B}{n}-p<-\epsilon p\right\} \leq e^{-3 n p \epsilon^{2} / 8} .
$$

Problem 8.4. Let $X_{1}, \ldots, X_{n}$ be independent binary-valued random variables with $\mathbf{P}\left\{X_{i}=\right.$ $1\}=1-\mathbf{P}\left\{X_{i}=0\right\}=p_{i}$. Set $p=(1 / n) \sum_{i=1}^{n} p_{i}$ and $S_{n}=\sum_{i=1}^{n} X_{i}$. Prove that

$$
\mathbf{P}\left\{S_{n}-n p \geq n \epsilon\right\} \leq e^{-n p \epsilon^{2} / 3} \quad \text { and } \quad \mathbf{P}\left\{S_{n}-n p \leq-n \epsilon\right\} \leq e^{-n p \epsilon^{2} / 2}
$$

(Angluin and Valiant (1979), see also Hagerup and Rüb (1990)). Compare the results with Bernstein's inequality for this case. Hint: Put $s=\log (1+\epsilon)$ and $s=-\log (1-\epsilon)$ in the Chernoff bounding argument. Prove and exploit the elementary inequalities

$$
-\frac{\epsilon^{2}}{2} \leq \epsilon-(1+\epsilon) \log (1+\epsilon) \leq-\frac{\epsilon^{2}}{3}, \quad \epsilon \in[0,1]
$$

and

$$
-\frac{\epsilon^{2}}{2} \geq \epsilon-(1+\epsilon) \log (1+\epsilon), \quad \epsilon \in(-1,0] .
$$

Problem 8.5. Let $B$ be a Binomial ( $n, p$ ) random variable. Show that for $p \leq a<1$,

$$
\mathbf{P}\{B>a n\} \leq\left(\left(\frac{p}{a}\right)^{a}\left(\frac{1-p}{1-a}\right)^{1-a}\right)^{n} \leq\left(\left(\frac{p}{a}\right)^{a} e^{a-p}\right)^{n}
$$

Show that for $0<a<p$ the same upper bounds hold for $\mathbf{P}\{B \leq a n\}$ (Karp (1988), see also Hagerup and Rüb (1990)). Hint: Use Chernoff's method with parameter $s$ and set $s=\log (a(1-p) /(p(1-a)))$.

Problem 8.6. Let $B$ be a $\operatorname{Binomial}(n, p)$ random variable. Show that if $p \geq 1 / 2$,

$$
\mathbf{P}\{B-n p \geq n \epsilon\}<e^{-\frac{n \epsilon^{2}}{2 p(1-p)}},
$$

and if $p \leq 1 / 2$,

$$
\mathbf{P}\{B-n p \leq-n \epsilon\}<e^{-\frac{n \epsilon^{2}}{2 p(1-p)}}
$$

(Okamoto (1958)). Hint: Use Chernoff's method, and the inequality

$$
x \log \frac{x}{p}+(1-x) \log \frac{1-x}{1-p} \geq \frac{(x-p)^{2}}{2 p(1-p)}
$$

for $1 / 2 \leq p \leq x \leq 1$, and for $0 \leq x \leq p \leq 1 / 2$.
Problem 8.7. Let $B$ be a $\operatorname{Binomial}(n, p)$ random variable. Prove that

$$
\mathbf{P}\{\sqrt{B}-\sqrt{n p} \geq \epsilon \sqrt{n}\}<e^{-2 n \epsilon^{2}},
$$

and

$$
\mathbf{P}\{\sqrt{B}-\sqrt{n p} \leq-\epsilon \sqrt{n}\}<e^{-n \epsilon^{2}}
$$

(Okamoto (1958)). Hint: Use Chernoff's method, and the inequalities

$$
x \log \frac{x}{p}+(1-x) \log \frac{1-x}{1-p} \geq 2(\sqrt{x}-\sqrt{p})^{2} \quad x \in[p, 1]
$$

and

$$
x \log \frac{x}{p}+(1-x) \log \frac{1-x}{1-p} \geq(\sqrt{x}-\sqrt{p})^{2} \quad x \in[0, p],
$$

Problem 8.8. Give a class $\mathcal{C}$ of decision functions of the form $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$ (i.e., the training data do not play any role in the decision) such that for every $\epsilon>0$

$$
\sup _{\phi \in \mathcal{C}} \mathbf{P}\left\{\left|\widehat{L}_{n}(\phi)-L(\phi)\right|>\epsilon\right\} \leq 2 e^{-2 n \epsilon^{2}}
$$

for every distribution, where $\widehat{L}_{n}(\phi)$ is the error-counting estimate of the error probability $L(\phi)=\mathbf{P}\{\phi(X) \neq Y\}$ of decision $\phi$, and at the same time, if $\mathcal{F}_{n}$ is the class of mappings $\phi_{n}^{*}$ minimizing the error count $\widehat{L}_{n}(\phi)$ over the class $\mathcal{C}$, then there exists one distribution such that

$$
\mathbf{P}\left\{\sup _{\phi \in \mathcal{F}_{n}} L(\phi)-\inf _{\phi \in \mathcal{C}} L(\phi)=1\right\}=1
$$

for all $n$.
Problem 8.9. Let $\mathcal{C}$ be a class of classifiers, that is, a class of mappings of the form $\phi_{n}\left(x, D_{n}\right)=\phi_{n}(x)$. Assume that an independent testing sequence $T_{m}$ is given, and that the error count

$$
\widehat{L}_{n, m}\left(\phi_{n}\right)=\frac{1}{m} \sum_{j=1}^{m} I_{\left\{\phi_{n}\left(X_{n+j}\right) \neq Y_{n+j}\right\}}
$$

is used to estimate the error probability $L\left(\phi_{n}\right)=\mathbf{P}\left\{\phi_{n}(X) \neq Y \mid D_{n}\right\}$ of each classifier $\phi_{n} \in \mathcal{C}$. Denote by $\phi_{n, m}^{*}$ the classifier that minimizes the estimated error probability over the class. Prove that for the error probability

$$
L\left(\phi_{n, m}^{*}\right)=\mathbf{P}\left\{\phi_{n, m}^{*}(X) \neq Y \mid D_{n}\right\}
$$

of the selected rule we have

$$
\begin{aligned}
& L\left(\phi_{n, m}^{*}\right)-\inf _{\phi_{n} \in \mathcal{C}} L\left(\phi_{n}\right) \leq 2 \sup _{\phi_{n} \in \mathcal{C}}\left|\widehat{L}_{n, m}\left(\phi_{n}\right)-L\left(\phi_{n}\right)\right|, \\
& \left|\widehat{L}_{n, m}\left(\phi_{n, m}^{*}\right)-L\left(\phi_{n, m}^{*}\right)\right| \leq \sup _{\phi_{n} \in \mathcal{C}}\left|\widehat{L}_{n, m}\left(\phi_{n}\right)-L\left(\phi_{n}\right)\right| .
\end{aligned}
$$

Also, if $\mathcal{C}$ is of finite cardinality with $|\mathcal{C}|=N$, then

$$
\mathbf{P}\left\{\sup _{\phi_{n} \in \mathcal{C}}\left|\widehat{L}_{n, m}\left(\phi_{n}\right)-L\left(\phi_{n}\right)\right|>\epsilon \mid D_{n}\right\} \leq 2 N e^{-2 m \epsilon^{2}} .
$$

Problem 8.10. Show that if a rule $g_{n}$ is consistent, then we can always find an estimate of the error such that $\mathbf{E}\left\{\left|\hat{L}_{n}-L_{n}\right|^{q}\right\} \rightarrow 0$ for all $q>0$. Hint: Split the data sequence $D_{n}$ and use the second half to estimate the error probability of $g_{\lceil n / 2\rceil}$.

Problem 8.11. Open-ended problem. Is there a rule for which no error estimate works for all distributions? More specifically, is there a sequence of classification rules $g_{n}$ such that for all $n$ large enough,

$$
\inf _{\hat{L}_{n} X, Y} \sup _{X} \mathbf{E}\left\{\left(\hat{L}_{n}-L_{n}\right)^{2}\right\} \geq c
$$

for some constant $c>0$, where the infimum is taken over all possible error estimates? Are such rules necessarily inconsistent?

Problem 8.12. Consider the problem of estimating the asymptotic probability of error of the nearest neighbor rule $L_{\mathrm{NN}}=2 \mathbf{E}\{\eta(X)(1-\eta(X))\}$. Show that for every $n$, for any estimate $\widehat{L}_{n}$ of $L_{\mathrm{NN}}$, and for every $\epsilon>0$, there exists a distribution of ( $X, Y$ ), such that

$$
\mathbf{E}\left\{\left|\widehat{L}_{n}-L_{\mathrm{NN}}\right|\right\} \geq \frac{1}{4}-\epsilon
$$

## 9

## The Regular Histogram Rule

In this chapter we study the cubic histogram rule. Recall that this rule partitions $\mathcal{R}^{d}$ into cubes of the same size, and gives the decision according to the number of zeros and ones among the $Y_{i}$ 's such that the corresponding $X_{i}$ falls in the same cube as $X . \mathcal{P}_{n}=\left\{A_{n 1}, A_{n 2}, \ldots\right\}$ denotes a partition of $\mathcal{R}^{d}$ into cubes of size $h_{n}>0$, that is, into sets of the type $\prod_{i=1}^{d}\left[k_{i} h_{n},\left(k_{i}+1\right) h_{n}\right)$, where the $k_{i}$ 's are integers, and the histogram rule is defined by

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} I_{\left\{X_{i} \in A_{n}(x)\right\}} \leq \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} I_{\left\{X_{i} \in A_{n}(x)\right\}} \\ 1 & \text { otherwise }\end{cases}
$$

where for every $x \in \mathcal{R}^{d}, A_{n}(x)=A_{n i}$ if $x \in A_{n i}$. That is, the decision is zero if the number of ones does not exceed the number of zeros in the cell in which $x$ falls. Weak universal consistency of this rule was shown in Chapter 6 under the conditions $h_{n} \rightarrow 0$ and $n h_{n}^{d} \rightarrow \infty$ as $n \rightarrow \infty$. The purpose of this chapter is to introduce some techniques by proving strong universal consistency of this rule. These techniques will prove very useful in handling other problems as well. First we introduce the method of bounded differences.

### 9.1 The Method of Bounded Differences

In this section we present a generalization of Hoeffding's inequality, due to McDiarmid (1989). The result will equip us with a powerful tool to handle complicated functions of independent random variables. This inequality follows by results of Hoeffding (1963) and Azuma (1967) who observed that Theorem 8.1
can be generalized to bounded martingale difference sequences. The inequality has found many applications in combinatorics, as well as in nonparametric statistics (see McDiarmid (1989) and Devroye (1991a) for surveys).

Let us first recall the notion of martingales. Consider a probability space ( $\Omega, \mathcal{F}, \mathbf{P}$ ).
DEFINITION 9.1. A sequence of random variables $Z_{1}, Z_{2}, \ldots$ is called a martingale if

$$
\mathbf{E}\left\{Z_{i+1} \mid Z_{1}, \ldots, Z_{i}\right\}=Z_{i} \text { with probability one }
$$

for each $i>0$.
Let $X_{1}, X_{2}, \ldots$ be an arbitrary sequence of random variables. $Z_{1}, Z_{2}, \ldots$ is called a martingale with respect to the sequence $X_{1}, X_{2}, \ldots$ if for every $i>0, Z_{i}$ is a function of $X_{1}, \ldots, X_{i}$, and

$$
\mathbf{E}\left\{Z_{i+1} \mid X_{1}, \ldots, X_{i}\right\}=Z_{i} \text { with probability one. }
$$

Obviously, if $Z_{1}, Z_{2}, \ldots$ is a martingale with respect to $X_{1}, X_{2}, \ldots$, then $Z_{1}$, $Z_{2}, \ldots$ is a martingale, since

$$
\begin{aligned}
\mathbf{E}\left\{Z_{i+1} \mid Z_{1}, \ldots, Z_{i}\right\} & =\mathbf{E}\left\{\mathbf{E}\left\{Z_{i+1} \mid X_{1}, \ldots, X_{i}\right\} Z_{1}, \ldots, Z_{i}\right\} \\
& =\mathbf{E}\left\{Z_{i} \mid Z_{1}, \ldots, Z_{i}\right\} \\
& =Z_{i} .
\end{aligned}
$$

The most important examples of martingales are sums of independent zeromean random variables. Let $U_{1}, U_{2}, \ldots$ be independent random variables with zero mean. Then the random variables

$$
S_{i}=\sum_{j=1}^{i} U_{j}, \quad i>0
$$

form a martingale (see Problem 9.1). Martingales share many properties of sums of independent variables. Our purpose here is to extend Hoeffding's inequality to martingales. The role of the independent random variables is played here by a so-called martingale difference sequence.

DEFINITION 9.2. A sequence of random variables $V_{1}, V_{2}, \ldots$ is a martingale difference sequence if

$$
\mathbf{E}\left\{V_{i+1} \mid V_{1}, \ldots, V_{i}\right\}=0 \text { with probability one }
$$

for every $i>0$.
A sequence of random variables $V_{1}, V_{2}, \ldots$ is called a martingale difference sequence with respect to a sequence of random variables $X_{1}, X_{2}, \ldots$ iffor every $i>0 V_{i}$ is a function of $X_{1}, \ldots, X_{i}$, and

$$
\mathbf{E}\left\{V_{i+1} \mid X_{1}, \ldots, X_{i}\right\}=0 \text { with probability one. }
$$

Again, it is easily seen that if $V_{1}, V_{2}, \ldots$ is a martingale difference sequence with respect to a sequence $X_{1}, X_{2}, \ldots$ of random variables, then it is a martingale difference sequence. Also, any martingale $Z_{1}, Z_{2}, \ldots$ leads naturally to a martingale difference sequence by defining

$$
V_{i}=Z_{i}-Z_{i-1}
$$

for $i>0$.
The key result in the method of bounded differences is the following inequality that relaxes the independence assumption in Theorem 8.1, allowing martingale difference sequences:

Theorem 9.1. (Hoeffding (1963), Azuma (1967)). Let $X_{1}, X_{2}, \ldots$ be a sequence of random variables, and assume that $V_{1}, V_{2}, \ldots$ is a martingale difference sequence with respect to $X_{1}, X_{2}, \ldots$. Assume furthermore that there exist random variables $Z_{1}, Z_{2}, \ldots$ and nonnegative constants $c_{1}, c_{2}, \ldots$ such thatfor everyi $>0$ $Z_{i}$ is a function of $X_{1}, \ldots, X_{i-1}$, and

$$
Z_{i} \leq V_{i} \leq Z_{i}+c_{i} \text { with probability one. }
$$

Then for any $\epsilon>0$ and $n$

$$
\mathbf{P}\left\{\sum_{i=1}^{n} V_{i} \geq \epsilon\right\} \leq e^{-2 \epsilon^{2} / \sum_{i=1}^{n} c_{i}^{2}}
$$

and

$$
\mathbf{P}\left\{\sum_{i=1}^{n} V_{i} \leq-\epsilon\right\} \leq e^{-2 \epsilon^{2} / \sum_{i=1}^{n} c_{i}^{2}} .
$$

The proof is a rather straightforward extension of that of Hoeffding's inequality. First we need an analog of Lemma 8.1:

Lemma 9.1. Assume that the random variables $V$ and $Z$ satisfy with probability one that $\mathbf{E}\{V \mid Z\}=0$, and for some function $f$ and constant $c \geq 0$

$$
f(Z) \leq V \leq f(Z)+c
$$

Then for every $s>0$

$$
\mathbf{E}\left\{e^{s V} \mid Z\right\} \leq e^{s^{2} c^{2} / 8}
$$

The proof of the lemma is left as an exercise (Problem 9.2).
Proof of Theorem 9.1. As in the proof of Hoeffding's inequality, we proceed by Chernoff's bounding method. Set $S_{k}=\sum_{i=1}^{k} V_{i}$. Then for any $s>0$

$$
\begin{aligned}
\mathbf{P}\left\{S_{n} \geq \epsilon\right\} & \leq e^{-s \epsilon} \mathbf{E}\left\{e^{s S_{n}}\right\} \\
& =e^{-s \epsilon} \mathbf{E}\left\{e^{s S_{n-1}} \mathbf{E}\left\{e^{s V_{n}} \mid X_{1}, \ldots, X_{n-1}\right\}\right\}
\end{aligned}
$$

$$
\begin{aligned}
& \leq e^{-s \epsilon} \mathbf{E}\left\{e^{s S_{n-1}}\right\} e^{s^{2} c_{n}^{2} / 8} \quad \text { (by Lemma 9.1) } \\
& \leq e^{-s \epsilon} e^{s^{2} \sum_{i=1}^{n} c_{i}^{2} / 8} \quad \text { (iterate previous argument) } \\
& \left.=e^{-2 \epsilon^{2} / \sum_{i=1}^{n} c_{i}^{2}} \text { (choose } s=4 \epsilon / \sum_{i=1}^{n} c_{i}^{2}\right) .
\end{aligned}
$$

The second inequality is proved analogously.

Now, we are ready to state the main inequality of this section. It is a large deviation-type inequality for functions of independent random variables such that the function is relatively robust to individual changes in the values of the random variables. The condition of the function requires that by changing the value of its $i$-th variable, the value of the function cannot change by more than a constant $c_{i}$.

Theorem 9.2. (MCDIARMID (1989)). Let $X_{1}, \ldots, X_{n}$ be independent random variables taking values in a set $A$, and assume that $f: A^{n} \rightarrow R$ satisfies

$$
\sup _{\substack{x_{1}, \ldots, x_{n} \\ x_{i}^{\prime} \in A}}\left|f\left(x_{1}, \ldots, x_{n}\right)-f\left(x_{1}, \ldots, x_{i-1}, x_{i}^{\prime}, x_{i+1}, \ldots, x_{n}\right)\right| \leq c_{i}, 1 \leq i \leq n
$$

Then for all $\epsilon>0$

$$
\mathbf{P}\left\{f\left(X_{1}, \ldots, X_{n}\right)-\mathbf{E} f\left(X_{1}, \ldots, X_{n}\right) \geq \epsilon\right\} \leq e^{-2 \epsilon^{2} / \sum_{i=1}^{n} c_{i}^{2}}
$$

and

$$
\mathbf{P}\left\{\mathbf{E} f\left(X_{1}, \ldots, X_{n}\right)-f\left(X_{1}, \ldots, X_{n}\right) \geq \epsilon\right\} \leq e^{-2 \epsilon^{2} / \sum_{i=1}^{n} c_{i}^{2}}
$$

Proof. Define $V=f\left(X_{1}, \ldots, X_{n}\right)-\mathbf{E} f\left(X_{1}, \ldots, X_{n}\right)$. Introduce $V_{1}=\mathbf{E}\left\{V \mid X_{1}\right\}-$ $\mathbf{E} V$, and for $k>1$,

$$
V_{k}=\mathbf{E}\left\{V \mid X_{1} \ldots, X_{k}\right\}-\mathbf{E}\left\{V \mid X_{1}, \ldots, X_{k-1}\right\}
$$

so that $V=\sum_{k=1}^{n} V_{k}$. Clearly, $V_{1}, \ldots, V_{n}$ form a martingale difference sequence with respect to $X_{1}, \ldots, X_{n}$. Define the random variables

$$
H_{k}\left(X_{1}, \ldots, X_{k}\right)=\mathbf{E}\left\{f\left(X_{1}, \ldots, X_{n}\right) \mid X_{1}, \ldots, X_{k}\right\}
$$

and

$$
V_{k}=H_{k}\left(X_{1}, \ldots, X_{k}\right)-\int H_{k}\left(X_{1}, \ldots, X_{k-1}, x\right) F_{k}(d x)
$$

where the integration is with respect to $F_{k}$, the probability measure of $X_{k}$. Introduce the random variables

$$
W_{k}=\sup _{u}\left(H_{k}\left(X_{1}, \ldots, X_{k-1}, u\right)-\int H_{k}\left(X_{1}, \ldots, X_{k-1}, x\right) F_{k}(d x)\right)
$$

and

$$
Z_{k}=\inf _{v}\left(H_{k}\left(X_{1}, \ldots, X_{k-1}, v\right)-\int H_{k}\left(X_{1}, \ldots, X_{k-1}, x\right) F_{k}(d x)\right)
$$

Clearly, $Z_{k} \leq V_{k} \leq W_{k}$ with probability one. Since for every $k Z_{k}$ is a function of $X_{1}, \ldots, X_{k-1}$, we can apply Theorem 9.1 directly to $V=\sum_{k=1}^{n} V_{k}$, if we can show that $W_{k}-Z_{k} \leq c_{k}$. But this follows from

$$
\begin{aligned}
W_{k}-Z_{k} & =\sup _{u} \sup _{v}\left(H_{k}\left(X_{1}, \ldots, X_{k-1}, u\right)-H_{k}\left(X_{1}, \ldots, X_{k-1}, v\right)\right) \\
& \leq c_{k}
\end{aligned}
$$

by the condition of the theorem.

Clearly, if the $X_{i}$ 's are bounded, then the choice $f\left(x_{1}, \ldots, x_{n}\right)=\sum_{i=1}^{n} x_{i}$ yields Hoeffding's inequality. Many times the inequality can be used to handle very complicated functions of independent random variables with great elegance. For examples in nonparametric statistics, see Problems 9.3, 9.6, 10.3.

Similar methods to those used in the proof of Theorem 9.2 may be used to bound the variance $\operatorname{Var}\left\{f\left(X_{1}, \ldots, X_{n}\right)\right\}$. Other inequalities for the variance of general functions of independent random variables were derived by Efron and Stein (1981) and Steele (1986).

Theorem 9.3. (Devroye (1991A)). Assume that the conditions of Theorem 9.2 hold. Then

$$
\operatorname{Var}\left\{f\left(X_{1}, \ldots, X_{n}\right)\right\} \leq \frac{1}{4} \sum_{i=1}^{n} c_{i}^{2}
$$

Proof. Using the notations of the proof of Theorem 9.2, we have to show that

$$
\operatorname{Var}\{V\} \leq \frac{1}{4} \sum_{i=1}^{n} c_{i}^{2}
$$

Observe that

$$
\begin{aligned}
\operatorname{Var}\{V\} & =\mathbf{E} V^{2} \\
& =\mathbf{E}\left\{\left(\sum_{i=1}^{n} V_{i}\right)^{2}\right\} \\
& =\mathbf{E}\left\{\sum_{i=1}^{n} V_{i}^{2}\right\}+2 \sum_{1 \leq i<j \leq n} \mathbf{E}\left\{V_{i} V_{j}\right\} \\
& =\mathbf{E}\left\{\sum_{i=1}^{n} V_{i}^{2}\right\},
\end{aligned}
$$

where in the last step we used the martingale property in the following way: for $i<j$ we have

$$
\begin{aligned}
& \mathbf{E}\left\{V_{i} V_{j} \mid X_{1}, \ldots, X_{j-1}\right\} \\
& =V_{i} \mathbf{E}\left\{V_{j} \mid X_{1}, \ldots, X_{j-1}\right\} \\
& \\
& =0 \text { with probability one. }
\end{aligned}
$$

Thus, the theorem follows if we can show that

$$
\mathbf{E}\left\{V_{i}^{2} \mid X_{1}, \ldots, X_{i-1}\right\} \leq \frac{c_{i}^{2}}{4}
$$

Introducing $W_{i}$ and $Z_{i}$ as in the proof of Theorem 9.2 , we see that with probability one $Z_{i} \leq V_{i} \leq Z_{i}+c_{i}$. Since $Z_{i}$ is a function of $X_{1}, \ldots, X_{i-1}$, therefore, conditioned on $X_{1}, \ldots, X_{i-1}, V_{i}$ is a zero mean random variable taking values in the interval $\left[Z_{i}, Z_{i}+c_{i}\right]$. But an arbitrary random variable $U$ taking values in an interval $[a, b]$ has variance not exceeding

$$
\mathbf{E}\left\{(U-(a+b) / 2)^{2}\right\} \leq \frac{(b-a)^{2}}{4}
$$

so that

$$
\mathbf{E}\left\{V_{i}^{2} \mid X_{1}, \ldots, X_{i-1}\right\} \leq \frac{c_{i}^{2}}{4}
$$

which concludes the proof.

### 9.2 Strong Universal Consistency

The purpose of this section is to prove strong universal consistency of the histogram rule. This is the first such result that we mention. Later we will prove the same property for other rules, too. The theorem, stated here for cubic partitions, is essentially due to Devroye and Györfi (1983). For more general sequences of partitions, see Problem 9.7. An alternative proof of the theorem based on the Vapnik-Chervonenkis inequality will be given later-see the remark following Theorem 17.2.

Theorem 9.4. Assume that the sequence of partitions $\left\{\mathcal{P}_{n}\right\}$ satisfies the following two conditions as $n \rightarrow \infty$ :

$$
h_{n} \rightarrow 0
$$

and

$$
n h_{n}^{d} \rightarrow \infty
$$

For any distribution of $(X, Y)$, and for every $\epsilon>0$ there is an integer $n_{0}$ such that for $n>n_{0}$, for the error probability $L_{n}$ of the histogram rule

$$
\mathbf{P}\left\{L_{n}-L^{*}>\epsilon\right\} \leq 2 e^{-n \epsilon^{2} / 32}
$$

Thus, the cubic histogram rule is strongly universally consistent.

Proof. Define

$$
\eta_{n}^{*}(x)=\frac{\sum_{i=1}^{n} Y_{i} I_{\left\{X_{i} \in A_{n}(x)\right\}}}{n \mu\left(A_{n}(x)\right)}
$$

Clearly, the decision based on $\eta_{n}^{*}$,

$$
g_{n}(x)= \begin{cases}0 & \text { if } \frac{\sum_{i=1}^{n} Y_{i} I_{\left[X_{i} \in A_{n}(x)\right]}^{n \mu\left(A_{n}(x)\right)}}{1}\end{cases}
$$

is just the histogram rule. Therefore, by Theorem 2.3, it suffices to prove that for $n$ large enough,

$$
\mathbf{P}\left\{\int\left|\eta(x)-\eta_{n}^{*}(x)\right| \mu(d x)>\frac{\epsilon}{2}\right\} \leq e^{-n \epsilon^{2} / 32}
$$

Decompose the difference as

$$
\begin{equation*}
\left|\eta(x)-\eta_{n}^{*}(x)\right|=\mathbf{E}\left|\eta(x)-\eta_{n}^{*}(x)\right|+\left(\left|\eta(x)-\eta_{n}^{*}(x)\right|-\mathbf{E}\left|\eta(x)-\eta_{n}^{*}(x)\right|\right) \tag{9.1}
\end{equation*}
$$

The convergence of the first term on the right-hand side implies weak consistency of the histogram rule. The technique we use to bound this term is similar to that which we already saw in the proof of Theorem 6.3. For completeness, we give the details here. However, new ideas have to appear in our handling of the second term.

We begin with the first term. Since the set of continuous functions with bounded support is dense in $L_{1}(\mu)$, it is possible to find a continuous function of bounded support $r(x)$ such that

$$
\int|\eta(x)-r(x)| \mu(d x)<\epsilon / 16
$$

Note that $r(x)$ is uniformly continuous. Introduce the function

$$
r_{n}^{*}(x)=\frac{\mathbf{E}\left\{r(X) I_{\left\{X \in A_{n}(x)\right\}}\right\}}{\mu\left(A_{n}(x)\right)}
$$

Then we can further decompose the first term on the right-hand side of (9.1) as

$$
\begin{align*}
& \mathbf{E}\left|\eta(x)-\eta_{n}^{*}(x)\right| \\
& \quad \leq \quad|\eta(x)-r(x)|+\left|r(x)-r_{n}^{*}(x)\right| \\
& \quad+\left|r_{n}^{*}(x)-\mathbf{E} \eta_{n}^{*}(x)\right|+\mathbf{E}\left|\mathbf{E} \eta_{n}^{*}(x)-\eta_{n}^{*}(x)\right| \tag{9.2}
\end{align*}
$$

We proceed term by term:
FIRST TERM: The integral of $|\eta(x)-r(x)|$ (with respect to $\mu$ ) is smaller than $\epsilon / 16$ by the definition of $r(x)$.

Second term: Using Fubini's theorem we have

$$
\begin{aligned}
& \int\left|r(x)-r_{n}^{*}(x)\right| \mu(d x) \\
& \quad=\sum_{A_{j}: \mu\left(A_{j}\right) \neq 0} \int_{A_{j}}\left|r(x)-\frac{\mathbf{E}\left\{r(X) I_{\left\{X \in A_{j}\right\}}\right\}}{\mu\left(A_{j}\right)}\right| \mu(d x)
\end{aligned}
$$

$$
\begin{aligned}
& =\sum_{A_{j}: \mu\left(A_{j}\right) \neq 0} \frac{1}{\mu\left(A_{j}\right)} \int_{A_{j}}\left|r(x) \mu\left(A_{j}\right)-\mathbf{E}\left\{r(X) I_{\left\{X \in A_{j}\right\}}\right\}\right| \mu(d x) \\
& =\sum_{A_{j}: \mu\left(A_{j}\right) \neq 0} \frac{1}{\mu\left(A_{j}\right)} \int_{A_{j}}\left|r(x) \int_{A_{j}} \mu(d y)-\int_{A_{j}} r(y) \mu(d y)\right| \mu(d x) \\
& \leq \sum_{A_{j}: \mu\left(A_{j}\right) \neq 0} \frac{1}{\mu\left(A_{j}\right)} \int_{A_{j}} \int_{A_{j}}|r(x)-r(y)| \mu(d x) \mu(d y) .
\end{aligned}
$$

As $r(x)$ is uniformly continuous, if $h_{n}$ is small enough, then $|r(x)-r(y)|<\epsilon / 16$ for every $x, y \in A$ for any cell $A \in \mathcal{P}_{n}$. Then the double integral in the above expression can be bounded from above as follows:

$$
\int_{A_{j}} \int_{A_{j}}|r(x)-r(y)| \mu(d x) \mu(d y) \leq \epsilon \mu^{2}\left(A_{j}\right) / 16
$$

Note that we used the condition $h_{n} \rightarrow 0$ here. Summing over the cells we get

$$
\int\left|r(x)-r_{n}^{*}(x)\right| \mu(d x) \leq \epsilon / 16
$$

Third term: We have

$$
\begin{aligned}
\int\left|r_{n}^{*}(x)-\mathbf{E} \eta_{n}^{*}(x)\right| \mu(d x) & =\sum_{A_{j} \in \mathcal{P}_{n}}\left|\mathbf{E}\left\{r(X) I_{\left\{X \in A_{j}\right\}}-Y I_{\left\{X \in A_{j}\right\}}\right\}\right| \\
& =\sum_{A_{j} \in \mathcal{P}_{n}}\left|\int_{A_{j}} r(x) \mu(d x)-\int_{A_{j}} \eta(x) \mu(d x)\right| \\
& \leq \int|r(x)-\eta(x)| \mu(d x)<\epsilon / 16
\end{aligned}
$$

Fourth term: Our aim is to show that for $n$ large enough,

$$
\mathbf{E} \int\left|\mathbf{E} \eta_{n}^{*}(x)-\eta_{n}^{*}(x)\right| \mu(d x)<\epsilon / 16
$$

To this end, let $S$ be an arbitrary large ball centered at the origin. Denote by $m_{n}$ the number of cells of the partition $\mathcal{P}_{n}$ that intersect $S$. Clearly, $m_{n}$ is proportional to $1 / h_{n}^{d}$ as $h_{n} \rightarrow 0$. Using the notation $v_{n}(A)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{Y_{i}=1, X_{i} \in A\right\}}$, it is clear that $v_{n}(A)=\int_{A} \eta_{n}^{*}(x) \mu(d x)$. Now, we can write

$$
\begin{aligned}
& \mathbf{E} \int\left|\mathbf{E} \eta_{n}^{*}(x)-\eta_{n}^{*}(x)\right| \mu(d x) \\
& \quad=\mathbf{E} \sum_{j} \int_{A_{n, j}}\left|\mathbf{E} \eta_{n}^{*}(x)-\eta_{n}^{*}(x)\right| \mu(d x) \\
& \quad=\mathbf{E} \sum_{j}\left|\mathbf{E} v_{n}\left(A_{n, j}\right)-v_{n}\left(A_{n, j}\right)\right|
\end{aligned}
$$

$$
\leq \mathbf{E} \sum_{j: A_{n, j} \cap S \neq \emptyset}\left|\mathbf{E} v_{n}\left(A_{n, j}\right)-v_{n}\left(A_{n, j}\right)\right|+2 \mu\left(S^{c}\right)
$$

(where $S^{c}$ denotes the complement of $S$ )

$$
\leq \sum_{j: A_{n, j} \cap S \neq \emptyset} \sqrt{\mathbf{E}\left|\mathbf{E} v_{n}\left(A_{n, j}\right)-v_{n}\left(A_{n, j}\right)\right|^{2}}+2 \mu\left(S^{c}\right)
$$

(by the Cauchy-Schwarz inequality)

$$
\begin{aligned}
& \leq m_{n} \frac{1}{m_{n}} \sum_{j: A_{n, j} \cap \neq \emptyset \emptyset} \sqrt{\frac{\mu\left(A_{n, j}\right)}{n}}+2 \mu\left(S^{c}\right) \\
& \leq m_{n} \sqrt{\frac{\frac{1}{m_{n}} \sum_{j: A_{n, j} \cap S \neq \emptyset} \mu\left(A_{n, j}\right)}{n}}+2 \mu\left(S^{c}\right)
\end{aligned}
$$

(by Jensen's inequality)

$$
\begin{aligned}
& \leq \sqrt{\frac{m_{n}}{n}}+2 \mu\left(S^{c}\right) \\
& \leq \epsilon / 16
\end{aligned}
$$

if $n$ and the radius of $S$ are large enough, since $m_{n} / n$ converges to zero by the condition $n h_{n}^{d} \rightarrow \infty$, and $\mu\left(S^{c}\right)$ can be made arbitrarily small by choice of $S$.

We have proved for the first term on the right-hand side of (9.1) that for $n$ large enough,

$$
\mathbf{E} \int\left|\eta(x)-\eta_{n}^{*}(x)\right| \mu(d x)<\epsilon / 4
$$

Finally, we handle the second term on the right-hand side of (9.1) by obtaining an exponential bound for

$$
\int\left|\eta(x)-\eta_{n}^{*}(x)\right| \mu(d x)-\mathbf{E} \int\left|\eta(x)-\eta_{n}^{*}(x)\right| \mu(d x)
$$

using Theorem 9.2. Fix the training data $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right) \in \mathcal{R}^{d} \times\{0,1\}$, and replace $\left(x_{i}, y_{i}\right)$ by $\left(\widehat{x}_{i}, \widehat{y}_{i}\right)$ changing the value of $\eta_{n}^{*}(x)$ to $\eta_{n i}^{*}(x)$. Then $\eta_{n}^{*}(x)-\eta_{n i}^{*}(x)$ differs from zero only on $A_{n}\left(x_{i}\right)$ and $A_{n}\left(\widehat{x}_{i}\right)$, and thus

$$
\begin{aligned}
& \int\left|\eta(x)-\eta_{n}^{*}(x)\right| \mu(d x)-\int\left|\eta(x)-\eta_{n i}^{*}(x)\right| \mu(d x) \\
& \quad \leq \int\left|\eta_{n}^{*}(x)-\eta_{n i}^{*}(x)\right| \mu(d x) \\
& \quad \leq\left(\frac{1}{n \mu\left(A_{n}\left(x_{i}\right)\right)} \mu\left(A_{n}\left(x_{i}\right)\right)+\frac{1}{n \mu\left(A_{n}\left(\widehat{x}_{i}\right)\right)} \mu\left(A_{n}\left(\widehat{x_{i}}\right)\right)\right) \leq \frac{2}{n}
\end{aligned}
$$

So by Theorem 9.2, for sufficiently large $n$,

$$
\begin{aligned}
\mathbf{P} & \left\{\int\left|\eta(x)-\eta_{n}^{*}(x)\right| \mu(d x)>\frac{\epsilon}{2}\right\} \\
& \leq \mathbf{P}\left\{\int\left|\eta(x)-\eta_{n}^{*}(x)\right| \mu(d x)-\mathbf{E} \int\left|\eta(x)-\eta_{n}^{*}(x)\right| \mu(d x)>\frac{\epsilon}{4}\right\} \\
& \leq e^{-n \epsilon^{2} / 32} .
\end{aligned}
$$

Remark. Strong universal consistency follows from the exponential bound on the probability $\mathbf{P}\left\{L_{n}-L^{*}>\epsilon\right\}$ via the Borel-Cantelli lemma. The inequality in Theorem 9.4 may seem universal in nature. However, it is distribution-dependent in a surreptitious way because its range of validity, $n \geq n_{0}$, depends heavily on $\epsilon, h_{n}$, and the distribution. We know that distribution-free upper bounds could not exist anyway, in view of Theorem 7.2.

## Problems and Exercises

Problem 9.1. Let $U_{1}, U_{2}, \ldots$ be independent random variables with zero mean. Show that the random variables $S_{i}=\sum_{j=1}^{i} U_{j} \quad i>0$ form a martingale.

Problem 9.2. Prove Lemma 9.1.
Problem 9.3. Let $X_{1}, \ldots, X_{n}$ be real valued i.i.d. random variables with distribution function $F(x)$, and corresponding empirical distribution function $F_{n}(x)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{X_{i} \leq x\right\}}$. Denote the Kolmogorov-Smirnov statistic by

$$
V_{n}=\sup _{x \in \mathcal{R}}\left|F_{n}(x)-F(x)\right| .
$$

Use Theorem 9.2 to show that

$$
\mathbf{P}\left\{\left|V_{n}-\mathbf{E} V_{n}\right| \geq \epsilon\right\} \leq 2 e^{-2 n \epsilon^{2}}
$$

Compare this result with Theorem 12.9. (None of them implies the other.)
Also, consider a class $\mathcal{A}$ of subsets of $\mathcal{R}^{d}$. Let $Z_{1}, \ldots, Z_{n}$ be i.i.d. random variables in $\mathcal{R}^{d}$ with common distribution $\mathbf{P}\left\{Z_{1} \in A\right\}=\nu(A)$, and consider the random variable

$$
W_{n}=\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|,
$$

where $v_{n}(A)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{Z_{i} \in A\right\}}$ denotes the standard empirical measure of $A$. Prove that

$$
\mathbf{P}\left\{\left|W_{n}-\mathbf{E} W_{n}\right| \geq \epsilon\right\} \leq 2 e^{-2 n \epsilon^{2}} .
$$

Compare this result with Theorem 12.5, and note that this result is true even if $s(\mathcal{A}, n)=2^{n}$ for all $n$.

Problem 9.4. The lazy histogram rule. Let $\mathcal{P}_{n}=\left\{A_{n 1}, A_{n 2}, \ldots\right\}$ be a sequence of partitions satisfying the conditions of the convergence Theorem 9.4. Define the lazy histogram rule as follows:

$$
g_{n}(x)=Y_{j}, \quad x \in A_{n i},
$$

where $X_{j}$ is the minimum-index point among $X_{1}, \ldots, X_{n}$ for which $X_{j} \in A_{n i}$. In other words, we ignore all but one point in each set of the partition. If $L_{n}$ is the conditional probability of error for the lazy histogram rule, then show that for any distribution of ( $X, Y$ ),

$$
\limsup _{n \rightarrow \infty} \mathbf{E} L_{n} \leq 2 L^{*}
$$

Problem 9.5. Assume that $\mathcal{P}_{n}=\mathcal{P}=\left\{A_{1}, \ldots, A_{k}\right\}$ is a fixed partition into $k$ sets. Consider the lazy histogram rule defined in Problem 9.4 based on $\mathcal{P}$. Show that for all distributions of $(X, Y), \lim _{n \rightarrow \infty} \mathbf{E} L_{n}$ exists and satisfies

$$
\lim _{n \rightarrow \infty} \mathbf{E} L_{n}=\sum_{i=1}^{k} 2 p_{i}\left(1-p_{i}\right) \mu\left(A_{i}\right)
$$

where $\mu$ is the probability measure for $X$, and $p_{i}=\int_{A_{i}} \eta(X) \mu(d x) / \mu\left(A_{i}\right)$. Show that the limit of the probability of error $\mathbf{E} L_{n}^{\prime}$ for the ordinary histogram rule is

$$
\lim _{n \rightarrow \infty} \mathbf{E} L_{n}^{\prime}=\sum_{i=1}^{k} \min \left(p_{i}, 1-p_{i}\right) \mu\left(A_{i}\right),
$$

and show that

$$
\lim _{n \rightarrow \infty} \mathbf{E} L_{n} \leq 2 \lim _{n \rightarrow \infty} \mathbf{E} L_{n}^{\prime}
$$

Problem 9.6. Histogram density estimation. Let $X_{1}, \ldots, X_{n}$ be i.i.d. random variables in $\mathcal{R}^{d}$ with density $f$. Let $\mathcal{P}$ be a partition of $\mathcal{R}^{d}$, and define the histogram density estimate by

$$
f_{n}(x)=\frac{1}{n \lambda(A(x))} \sum_{i=1}^{n} I_{\left\{X_{i} \in A(x)\right\}},
$$

where $A(x)$ is the set in $\mathcal{P}$ that contains $x$ and $\lambda$ is the Lebesgue measure. Prove for the $L_{1}$-error of the estimate that

$$
\mathbf{P}\left\{\left|\int\right| f_{n}(x)-f(x)\left|d x-\mathbf{E} \int\right| f_{n}(x)-f(x)|d x|>\epsilon\right\} \leq 2 e^{-n \epsilon^{2} / 2}
$$

(Devroye (1991a).) Conclude that weak $L_{1}$-consistency of the estimate implies strong consistency (Abou-Jaoude (1976a; 1976c), see also Problem 6.2).

Problem 9.7. General partitions. Extend the consistency result of Theorem 9.4 for sequences of general, not necessarily cubic partitions. Actually, cells of the partitions need not even be hyperrectangles. Assume that the sequence of partitions $\left\{\mathcal{P}_{n}\right\}$ satisfies the following two conditions. For every ball $S$ centered at the origin

$$
\lim _{n \rightarrow \infty} \max _{i: A_{n i} \cap S \neq \emptyset}\left(\sup _{x, y \in A_{n i}}\|x-y\|\right)=0
$$

and

$$
\lim _{n \rightarrow \infty} \frac{1}{n}\left|\left\{i: A_{n i} \cap S \neq \varnothing\right\}\right|=0
$$

Prove that the corresponding histogram classification rule is strongly universally consistent.

Problem 9.8. Show that for cubic histograms the conditions of Problem 9.7 on the partition are equivalent to the conditions $h_{n} \rightarrow 0$ and $n h_{n}^{d} \rightarrow \infty$, respectively.
Problem 9.9. Linear scaling. Partition $\mathcal{R}^{d}$ into congruent rectangles of the form

$$
\left[k_{1} h_{1},\left(k_{1}+1\right) h_{1}\right) \times \cdots \times\left[k_{d} h_{d},\left(k_{d}+1\right) h_{d}\right),
$$

where $k_{1}, \ldots, k_{d}$ are integers, and $h_{1}, \ldots, h_{d}>0$ denote the size of the edges of the rectangles. Prove that the corresponding histogram rule is strongly universally consistent if $h_{i} \rightarrow 0$ for every $i=1, \ldots, d$, and $n h_{1} h_{2} \ldots h_{d} \rightarrow \infty$ as $n \rightarrow \infty$. Hint: This is a Corollary of Problem 9.7.

Problem 9.10. Nonlinear scaling. Let $F_{1}, \ldots, F_{d}: \mathcal{R} \rightarrow \mathcal{R}$ be invertible, strictly monotone increasing functions. Consider the partition of $\mathcal{R}^{d}$, whose cells are rectangles of the form

$$
\left[F_{1}^{-1}\left(k_{1} h_{1}\right), F_{1}^{-1}\left(\left(k_{1}+1\right) h_{1}\right)\right) \times \ldots \times\left[F_{d}^{-1}\left(k_{d} h_{d}\right), F_{d}^{-1}\left(\left(k_{d}+1\right) h_{d}\right)\right)
$$

(See Problem 9.9.) Prove that the histogram rule corresponding to this partition is strongly universally consistent under the conditions of Problem 9.9. Hint: Use Problem 9.7.

Problem 9.11. Necessary and sufficient conditions for the bias. A sequence of partitions $\left\{\mathcal{P}_{n}\right\}$ is called $\mu$-approximating if for every measurable set $A$, for every $\epsilon>0$, and for all sufficiently large $n$ there is a set $A_{n} \in \sigma\left(\mathcal{P}_{n}\right),(\sigma(\mathcal{P})$ denotes the $\sigma$-algebra generated by cells of the partition $\mathcal{P}$ ) such that $\mu\left(A_{n} \triangle A\right)<\epsilon$. Prove that the bias term $\int\left|\eta(x)-\mathbf{E} \eta_{n}^{*}(x)\right| \mu(d x)$ converges to zero for all distributions of $(X, Y)$ if and only if the sequence of partitions $\left\{\mathcal{P}_{n}\right\}$ is $\mu$-approximating for every probability measure $\mu$ on $\mathcal{R}^{d}$ (Abou-Jaoude (1976a)). Conclude that the first condition of Problem 9.7 implies that $\left\{\mathcal{P}_{n}\right\}$ is $\mu$-approximating for every probability measure $\mu$ (Csiszár (1973)).

Problem 9.12. Necessary and sufficient conditions for the variation. Assume that for every probability measure $\mu$ on $\mathcal{R}^{d}$, every measurable set $A$, every $c>0$, and every $\epsilon>0$, there is an $N(\epsilon, c, A, \mu)$, such that for all $n>N(\epsilon, c, A, \mu)$,

$$
\sum_{j: \mu\left(A_{n, j} \cap A\right) \leq c / n} \mu\left(A_{n, j} \cap A\right)<\epsilon
$$

Prove that the variation term $\int\left|\mathbf{E} \eta_{n}^{*}(x)-\eta_{n}^{*}(x)\right| \mu(d x)$ converges to zero for all distributions of $(X, Y)$ if and only if the sequence of partitions $\left\{\mathcal{P}_{n}\right\}$ satisfies the condition above (AbouJaoude (1976a)).

Problem 9.13. The $\epsilon$-effective cardinality $m(\mathcal{P}, \mu, A, \epsilon)$ of a partition with respect to the probability measure $\mu$, restricted to a set $A$ is the minimum number of sets in $\mathcal{P}$ such that the union of the remaining sets intersected with $A$ has $\mu$-measure less than $\epsilon$. Prove that the sequence of partitions $\left\{\mathcal{P}_{n}\right\}$ satisfies the condition of Problem 9.12 if and only if for every $\epsilon>0$,

$$
\lim _{n \rightarrow \infty} \frac{m\left(\mathcal{P}_{n}, \mu, A, \epsilon\right)}{n}=0
$$

(Barron, Györfi, and van der Meulen (1992)).
Problem 9.14. In $\mathcal{R}^{2}$, partition the plane by taking three fixed points not on a line, $x, y$ and $z$. At each of these points, partition $\mathcal{R}^{2}$ by considering $k$ equal sectors of angle $2 \pi / k$ each. Sets in the histogram partition are obtained as intersections of cones. Is the induced histogram rule strongly universally consistent? If yes, state the conditions on $k$, and if no, provide a counterexample.

Problem 9.15. Partition $\mathcal{R}^{2}$ into shells of size $h$ each. The $i$-th shell contains all points at distance $d \in[(i-1) h, i h)$ from the origin. Let $h \rightarrow 0$ and $n h \rightarrow \infty$ as $n \rightarrow \infty$. Consider the histogram rule. As $n \rightarrow \infty$, to what does $\mathbf{E} L_{n}$ converge?

## 10

## Kernel Rules

Histogram rules have the somewhat undesirable property that the rule is less accurate at borders of cells of the partition than in the middle of cells. Looked at intuitively, this is because points near the border of a cell should have less weight in a decision regarding the cell's center. To remedy this problem, one might introduce the moving window rule, which is smoother than the histogram rule. This classifier simply takes the data points within a certain distance of the point to be classified, and decides according to majority vote. Working formally, let $h$ be a positive number. Then the moving window rule is defined as

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=0, X_{i} \in S_{x, h}\right\}} \geq \sum_{i=1}^{n} I_{\left\{Y_{i}=1, X_{i} \in S_{x, h}\right\}} \\ 1 & \text { otherwise }\end{cases}
$$

where $S_{x, h}$ denotes the closed ball of radius $h$ centered at $x$.
It is possible to make the decision even smoother by giving more weight to closer points than to more distant ones. Let $K: \mathcal{R}^{d} \rightarrow \mathcal{R}$ be a kernel function, which is usually nonnegative and monotone decreasing along rays starting from the origin. The kernel classification rule is given by

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} K\left(\frac{x-X_{i}}{h}\right) \geq \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} K\left(\frac{x-X_{i}}{h}\right) \\ 1 & \text { otherwise }\end{cases}
$$



Figure 10.1. The moving window rule in $\mathcal{R}^{2}$. The decision is 1 in the shaded area.

The number $h$ is called the smoothing factor, or bandwidth. It provides some form of distance weighting.


Figure 10.2. Kernel rule on the real line. The figure shows $\sum_{i=1}^{n}\left(2 Y_{i}-1\right) K\left(\left(x-X_{i}\right) / h\right)$ for $n=20, K(u)=\left(1-u^{2}\right) I_{\{|u| \leq 1\}}$ (the Epanechnikov kernel), and three smoothing factors $h$. One
 definitely undersmooths and one oversmooths. We took $p=1 / 2$, and the class-conditional densities are $f_{0}(x)=2(1-x)$ and $f_{1}(x)=2 x$ on $[0,1]$.

Clearly, the kernel rule is a generalization of the moving window rule, since taking the special kernel $K(x)=I_{\left\{x \in S_{0,1}\right\}}$ yields the moving window rule. This kernel is sometimes called the naïve kernel. Other popular kernels include the Gaussian kernel, $K(x)=e^{-\|x\|^{2}}$; the Cauchy kernel, $K(x)=1 /\left(1+\|x\|^{d+1}\right)$; and the

Epanechnikov kernel $K(x)=\left(1-\|x\|^{2}\right) I_{\{\|x\| \leq 1\}}$, where $\|\cdot\|$ denotes Euclidean distance.


Figure 10.3. Various kernels on $\mathcal{R}$.
Kernel-based rules are derived from the kernel estimate in density estimation originally studied by Parzen (1962), Rosenblatt (1956), Akaike (1954), and Cacoullos (1965) (see Problems 10.2 and 10.3); and in regression estimation, introduced by Nadaraya (1964; 1970), and Watson (1964). For particular choices of $K$, rules of this sort have been proposed by Fix and Hodges (1951; 1952), Sebestyen (1962), Van Ryzin (1966), and Meisel (1969). Statistical analysis of these rules and/or the corresponding regression function estimate can be found in Nadaraya (1964; 1970), Rejtő and Révész (1973), Devroye and Wagner (1976b; 1980a; 1980b), Greblicki (1974; 1978b; 1978a), Krzyżak and Pawlak (1984b), and Devroye and Krzyżak (1989). Usage of Cauchy kernels in discrimination is investigated by Arkadjew and Braverman (1966), Hand (1981), and Coomans and Broeckaert (1986).

### 10.1 Consistency

In this section we demonstrate strong universal consistency of kernel-based rules under general conditions on $h$ and $K$. Let $h>0$ be a smoothing factor depending only on $n$, and let $K$ be a kernel function. If the conditional densities $f_{0}, f_{1}$ exist, then weak and strong consistency follow from Problems 10.2 and 10.3 , respectively, via Problem 2.11. We state the universal consistency theorem for a large class of kernel functions, namely, for all regular kernels.

DEFInITION 10.1. The kernel $K$ is called regular if it is nonnegative, and there is a ball $S_{0, r}$ of radius $r>0$ centered at the origin, and constant $b>0$ such that $K(x) \geq b I_{S_{0, r}}$ and $\int \sup _{y \in x+S_{0, r}} K(y) d x<\infty$.

We provide three informative exercises on regular kernels (Problems 10.18, $10.19,10.20$ ). In all cases, regular kernels are bounded and integrable. The last condition holds whenever $K$ is integrable and uniformly continuous. Introduce the short notation $K_{h}(x)=\frac{1}{h} K\left(\frac{x}{h}\right)$. The next theorem states strong universal consistency of kernel rules. The theorem is essentially due to Devroye and Krzyżak
(1989). Under the assumption that $X$ has a density, it was proven by Devroye and Györfi (1985) and Zhao (1989).

Theorem 10.1. (Devroye and Krzyżak (1989)). Assume that $K$ is a regular kernel. If

$$
h \rightarrow 0 \text { and } n h^{d} \rightarrow \infty \text { as } n \rightarrow \infty
$$

then for any distribution of $(X, Y)$, and for every $\epsilon>0$ there is an integer $n_{0}$ such that for $n>n_{0}$ for the error probability $L_{n}$ of the kernel rule

$$
\mathbf{P}\left\{L_{n}-L^{*}>\epsilon\right\} \leq 2 e^{-n \epsilon^{2} /\left(32 \rho^{2}\right)}
$$

where the constant $\rho$ depends on the kernel $K$ and the dimension only. Thus, the kernel rule is strongly universally consistent.

Clearly, naïve kernels are regular, and moving window rules are thus strongly universally consistent. For the sake of readability, we give the proof for this special case only, and leave the extension to regular kernels to the reader-see Problems $10.14,10.15$, and 10.16 . Before we embark on the proof in the next section, we should warn the reader that Theorem 10.1 is of no help whatsoever regarding the choice of $K$ or $h$. One possible solution is to derive explicit upper bounds for the probability of error as a function of descriptors of the distribution of $(X, Y)$, and of $K, n$ and $h$. Minimizing such bounds with respect to $K$ and $h$ will lead to some expedient choices. Typically, such bounds would be based upon the inequality

$$
\begin{aligned}
& \mathbf{E} L_{n}-L^{*} \\
& \quad \leq \mathbf{E}\left\{\int\left|(1-p) f_{0}(x)-p_{n 0} f_{n 0}(x)\right| d x+\int\left|p f_{1}(x)-p_{n 1} f_{n 1}(x)\right| d x\right\}
\end{aligned}
$$

(see Chapter 6), where $f_{0}, f_{1}$ are the class densities, $f_{n 0}, f_{n 1}$ are their kernel estimates (see Problem 10.2) $(1-p)$ and $p$ are the class probabilities, and $p_{n 0}, p_{n 1}$ are their relative-frequency estimates. Bounds for the expected $L_{1}$-error in density estimation may be found in Devroye (1987) for $d=1$ and Holmström and Klemelä (1992) for $d>1$. Under regularity conditions on the distribution, the choice $h=c n^{-d /(d+4)}$ for some constant is asymptotically optimal in density estimation. However, $c$ depends upon unknown distributional parameters. Rather than following this roundabout process, we ask the reader to be patient and to wait until Chapter 25, where we study automatic kernel rules, i.e., rules in which $h$, and sometimes $K$ as well, is picked by the data without intervention from the statistician.

It is still too early to say meaningful things about the choice of a kernel. The kernel density estimate

$$
f_{n}(x)=\frac{1}{n h^{d}} \sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h}\right)
$$

based upon an i.i.d. sample $X_{1}, \ldots, X_{n}$ drawn from an unknown density $f$ is clearly a density in its own right if $K \geq 0$ and $\int K=1$. Also, there are certain
popular choices of $K$ that are based upon various optimality criteria. In pattern recognition, the story is much more confused, as there is no compelling a priori reason to pick a function $K$ that is nonnegative or integrable. Let us make a few points with the trivial case $n=1$. Taking $h=1$, the kernel rule is given by

$$
g_{1}(x)= \begin{cases}0 & \text { if } Y_{1}=0, K\left(x-X_{1}\right) \geq 0 \text { or if } Y_{1}=1, K\left(x-X_{1}\right) \leq 0 \\ 1 & \text { otherwise }\end{cases}
$$

If $K \geq 0$, then $g_{n}(x)=0$ if $Y_{1}=0$, or if $Y_{1}=1$ and $K\left(x-X_{1}\right)=0$. As we would obviously like $g_{n}(x)=0$ if and only if $Y_{1}=0$, it seems necessary to insist on $K>0$ everywhere. However, this restriction makes the kernel estimate nonlocal in nature.

For $n=1$ and $d=1$, consider next a negative-valued kernel such as the Hermite kernel

$$
K(x)=\left(1-x^{2}\right) e^{-x^{2}}
$$



Figure 10.4. Hermite kernel.

It is easy to verify that $K(x) \geq 0$ if and only if $|x| \leq 1$. Also, $\int K=0$. Nevertheless, we note that it yields a simple rule:

$$
g_{1}(x)= \begin{cases}0 & \text { if } Y_{1}=0,\left|x-X_{1}\right| \leq 1 \text { or if } Y_{1}=1,\left|x-X_{1}\right| \geq 1 \\ 1 & \text { otherwise }\end{cases}
$$

If we have a biatomic distribution for $X$, with equally likely atoms at 0 and 2 , and $\eta(0)=0$ and $\eta(2)=1$ (i.e., $Y=0$ if $X=0$ and $Y=1$ if $X=2$ ), then $L^{*}=0$ and the probability of error for this kernel rule $\left(L_{1}\right)$ is 0 as well. Note also that for all $n, g_{n}=g_{1}$ if we keep the same $K$. Consider now any positive kernel in the same example. If $X_{1}, \ldots, X_{n}$ are all zero, then the decision is $g_{n}(x)=0$ for all $x$. Hence $L_{n} \geq \frac{1}{2} \mathbf{P}\left\{X_{1}=\cdots=X_{n}\right\}=1 / 2^{n+1}>0$. Our negative zero-integral kernel is strictly better for all $n$ than any positive kernel! Such kernels should not be discarded without further thought. In density estimation, negative-valued kernels are used to reduce the bias under some smoothness conditions. Here, as shown above, there is an additional reason-negative weights given to points far away from the $X_{i}$ 's may actually be beneficial.

Staying with the same example, if $K>0$ everywhere, then

$$
\mathbf{E} L_{1}=\mathbf{P}\left\{Y_{1}=0, Y=1\right\}+\mathbf{P}\left\{Y_{1}=1, Y=0\right\}=2 \mathbf{E} \eta(X) \mathbf{E}\{1-\eta(X)\}
$$

which may be $1 / 2($ if $\mathbf{E} \eta(X)=1 / 2)$ even if $L^{*}=0$ (which happens when $\eta \in\{0,1\}$ everywhere). For this particular example, we would have obtained the same result
even if $K \equiv 1$ everywhere. With $K \equiv 1$, we simply ignore the $X_{i}$ 's and take a majority vote among the $Y_{i}$ 's (with $K \equiv-1$, it would be a minority vote!):

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} \geq \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} \\ 1 & \text { otherwise }\end{cases}
$$

Let $N_{n}$ be the number of $Y_{i}$ 's equal to zero. As $N_{n}$ is binomial $(n, 1-p)$ with $p=\mathbf{E} \eta(X)=\mathbf{P}\{Y=1\}$, we see that

$$
\mathbf{E} L_{n}=p \mathbf{P}\left\{N_{n} \geq \frac{n}{2}\right\}+(1-p) \mathbf{P}\left\{N_{n}<\frac{n}{2}\right\} \rightarrow \min (p, 1-p)
$$

simply by invoking the law of large numbers. Thus, $\mathbf{E} L_{n} \rightarrow \min (p, 1-p)$. As in the case with $n=1$, the limit is $1 / 2$ when $p=1 / 2$, even though $L^{*}=0$ when $\eta \in\{0,1\}$ everywhere. It is interesting to note the following though:

$$
\begin{aligned}
\mathbf{E} L_{1} & =2 p(1-p) \\
& =2 \min (p, 1-p)(1-\min (p, 1-p)) \\
& \leq 2 \min (p, 1-p) \\
& =2 \lim _{n \rightarrow \infty} \mathbf{E} L_{n}
\end{aligned}
$$

The expected error with one observation is at most twice as bad as the expected error with an infinite sequence. We have seen various versions of this inequality at work in many instances such as the nearest neighbor rule.

Let us apply the inequality for $\mathbf{E} L_{1}$ to each part in a fixed partition $\mathcal{P}$ of $\mathcal{R}^{d}$. On each of the $k$ sets $A_{1}, \ldots, A_{k}$ of $\mathcal{P}$, we apply a simple majority vote among the $Y_{i}$ 's, as in the histogram rule. If we define the lazy histogram rule as the one in which in each set $A_{i}$, we assign the class according to the $Y_{j}$ for which $X_{j} \in A_{i}$ and $j$ is the lowest such index ("the first point to fall in $A_{i}$ "). It is clear (see Problems 9.4 and 9.5) that

$$
\begin{aligned}
\lim _{n \rightarrow \infty} \mathbf{E} L_{\mathrm{LAZY}, n} & \leq 2 \lim _{n \rightarrow \infty} \mathbf{E} L_{n} \\
& =2 \sum_{i=1}^{k} \frac{1}{\mu\left(A_{i}\right)} \int_{A_{i}} \eta(x) \mu(d x) \int_{A_{i}}(1-\eta(x)) \mu(d x),
\end{aligned}
$$

where $L_{n}$ is the probability of error for the ordinary histogram rule. Again, the vast majority of observations is barely needed to reach a good decision.

Just for fun, let us return to a majority vote rule, now applied to the first three observations only. With $p=\mathbf{P}\{Y=1\}$, we see that

$$
\mathbf{E} L_{3}=p\left((1-p)^{3}+3(1-p)^{2} p\right)+(1-p)\left(3(1-p) p^{2}+p^{3}\right)
$$

by just writing down binomial probabilities. Observe that

$$
\begin{aligned}
\mathbf{E} L_{3} & =p(1-p)(1+4 p(1-p)) \\
& \leq \min (p, 1-p)(1+4 \min (p, 1-p)) \\
& =\lim _{n \rightarrow \infty} \mathbf{E} L_{n}\left(1+4 \lim _{n \rightarrow \infty} 4 \mathbf{E} L_{n}\right)
\end{aligned}
$$

If $\lim _{n \rightarrow \infty} \mathbf{E} L_{n}$ is small to start with, e.g., $\lim _{n \rightarrow \infty} \mathbf{E} L_{n}=0.01$, then $\mathbf{E} L_{3} \leq$ $0.01 \times 1.04=0.0104$. In such cases, it just does not pay to take more than three observations.

Kernels with fixed smoothing factors have no local sensitivity and, except in some circumstances, have probabilities of error that do not converge to $L^{*}$. The universal consistency theorem makes a strong case for decreasing smoothing factorsthere is no hope in general of approaching $L^{*}$ unless decisions are asymptotically local.

The consistency theorem describes kernel rules with $h \rightarrow 0$ : these rules become more and more local in nature as $n \rightarrow \infty$. The necessity of local rules is not apparent from the previous biatomic example. However, it is clear that if we consider a distribution in which given $Y=0, X$ is uniform on $\{\delta, 3 \delta, \ldots,(2 k+1) \delta\}$, and given $Y=1, X$ is uniform on $\{2 \delta, 4 \delta, \ldots, 2 k \delta\}$, that is, with the two classes intimately interwoven, a kernel rule with $K \geq 0$ of compact support [ $-1,1$ ], and $h<\delta<1$ will have

$$
\left.L_{n} \leq I_{\left\{U_{i=1}^{2 k}\left\{N_{i}=0\right\}\right.}\right\}
$$

where $N_{i}$ is the number of $X_{j}$ 's at the $i$-th atom. Hence $\mathbf{E} L_{n}$ goes to zero exponentially fast. If in the above example we assign $X$ by a geometric distribution on $\delta, \delta^{3}, \delta^{5}, \ldots$ when $Y=0$, and by a geometric distribution on $\delta^{2}, \delta^{4}, \delta^{6}, \ldots$ when $Y=1$, then to obtain $\mathbf{E} L_{n} \rightarrow 0$, it is necessary that $h \rightarrow 0$ (see Problem 10.1).
Remark. It is worthwhile to investigate what happens for negative-valued kernels $K$ when $h \rightarrow 0, n h^{d} \rightarrow \infty$ and $K$ has compact support. Every decision becomes an average over many local decisions. If $\mu$ has a density $f$, then at almost all points $x, f$ may be approximated very nicely by $\int_{S_{x, \delta}} f / \lambda\left(S_{x, \delta}\right)$ for small $\delta>0$, where $S_{x, \delta}$ is the closed ball of radius $\delta$ about $x$. This implies, roughly speaking, that the number of weighted votes from class 0 observations in a neighborhood of $x$ is about $(1-\eta(x)) f(x) n h^{d} \int K$, while for class 1 the weight is about $\eta(x) f(x) n h^{d} \int K$. The correct decision is nearly always made for $n h^{d}$ large enough provided that $\int K>0$. See Problem 10.4 on why kernels with $\int K<0$ should be avoided.

### 10.2 Proof of the Consistency Theorem

In the proof we can proceed as for the histogram. The crucial difference is captured in the following covering lemmas. Let $\beta_{d}$ denote the minimum number of balls of radius $1 / 2$ that cover the ball $S_{0,1}$. If $K$ is the naïve kernel, then $\rho=\beta_{d}$ in Theorem 10.1.

Lemma 10.1. (Covering Lemma). If $K(x)=I_{\left\{x \in S_{0,\}}\right\}}$, then for any $y \in \mathcal{R}^{d}, h>0$, and probability measure $\mu$,

$$
\int \frac{K_{h}(x-y)}{\int K_{h}(x-z) \mu(d z)} \mu(d x) \leq \beta_{d} .
$$

Proof. Cover the ball $S_{y, h}$ by $\beta_{d}$ balls of radius $h / 2$. Denote their centers by $x_{1}, \ldots, x_{\beta_{d}}$. Then $x \in S_{x_{i}, h / 2}$ implies $S_{x_{i}, h / 2} \subset S_{x, h}$ and thus

$$
\mu\left(S_{x, h}\right) \geq \mu\left(S_{x_{i}, h / 2}\right)
$$

We may write

$$
\begin{aligned}
\int \frac{K_{h}(x-y)}{\int K_{h}(x-z) \mu(d z)} \mu(d x) & =\int \frac{I_{\left\{x \in S_{y, h}\right\}}}{\mu\left(S_{x, h}\right)} \mu(d x) \\
& \leq \sum_{i=1}^{\beta_{d}} \int \frac{I_{\left\{x \in S_{x_{i}, h / 2}\right\}}}{\mu\left(S_{x, h}\right)} \mu(d x) \\
& \leq \sum_{i=1}^{\beta_{d}} \int \frac{I_{\left\{x \in S_{x_{i}, h / 2}\right\}}}{\mu\left(S_{x_{i}, h / 2}\right)} \mu(d x) \\
& =\beta_{d}
\end{aligned}
$$

Lemma 10.2. Let $0<h \leq R<\infty$, and let $S \subset \mathcal{R}^{d}$ be a ball of radius $R$. Then for every probability measure $\mu$,

$$
\int_{S} \frac{1}{\sqrt{\mu\left(S_{x, h}\right)}} \mu(d x) \leq\left(1+\frac{R}{h}\right)^{d / 2} c_{d}
$$

where $c_{d}$ depends upon the dimension $d$ only.
Proof. Cover $S$ with balls of radius $h / 2$, centered at center points of a regular grid of dimension $h /(2 \sqrt{d}) \times \ldots \times h /(2 \sqrt{d})$. Denote these centers by $x_{1}, \ldots, x_{m}$, where $m$ is the number of balls that cover $S$. Clearly,

$$
\begin{aligned}
m & \leq \frac{\operatorname{volume}\left(S_{0, R+h}\right)}{\operatorname{volume}(\text { grid cell })} \\
& =\frac{V_{d}(R+h)^{d}}{(h /(2 \sqrt{d}))^{d}}\left(V_{d} \text { is the volume of the unit ball in } \mathcal{R}^{d}\right) \\
& \leq\left(1+\frac{R}{h}\right)^{d} c_{d}^{\prime}
\end{aligned}
$$

where the constant $c_{d}^{\prime}$ depends upon the dimension only. Every $x$ gets covered at most $k_{1}$ times where $k_{1}$ depends upon $d$ only. Then we have

$$
\begin{aligned}
\int_{S} \frac{1}{\sqrt{\mu\left(S_{x, h}\right)}} \mu(d x) & \leq \sum_{i=1}^{m} \int_{S_{x_{i}, h / 2}} \frac{1}{\sqrt{\mu\left(S_{x, h}\right)}} \mu(d x) \\
& =\sum_{i=1}^{m} \int \frac{I_{\left\{x \in S_{x_{i}, h / 2}\right\}}}{\sqrt{\mu\left(S_{x, h}\right)}} \mu(d x)
\end{aligned}
$$

$$
\leq \sum_{i=1}^{m} \int \frac{I_{\left\{x \in S_{x_{i}, h / 2}\right\}}}{\sqrt{\mu\left(S_{x_{i}, h / 2}\right)}} \mu(d x)
$$

(by the same argument as in Lemma 10.1)

$$
\begin{aligned}
& \leq \sum_{i=1}^{m} \sqrt{\mu\left(S_{x_{i}, h / 2}\right)} \\
& \leq \sqrt{m \sum_{i=1}^{m} \mu\left(S_{x_{i}, h / 2}\right)}
\end{aligned}
$$

(by the Cauchy-Schwarz inequality)
$\leq \sqrt{k_{1} m}$

$$
\leq \sqrt{\left(1+\frac{R}{h}\right)^{d}} c_{d}
$$

where $c_{d}$ depends upon the dimension only.

Proof of Theorem 10.1. Define

$$
\eta_{n}(x)=\frac{\sum_{j=1}^{n} Y_{j} K_{h}\left(x-X_{j}\right)}{n \mathbf{E} K_{h}(x-X)}
$$

Since the decision rule can be written as

$$
g_{n}(x)= \begin{cases}0 & \text { if } \frac{\sum_{j=1}^{n} Y_{j} K_{h}\left(x-X_{j}\right)}{n \mathbf{E} K_{h}(x-X)} \leq \frac{\sum_{j=1}^{n}\left(1-Y_{j}\right) K_{h}\left(x-X_{j}\right)}{n \mathbf{E} K_{h}(x-X)} \\ 1 & \text { otherwise },\end{cases}
$$

by Theorem 2.3, what we have to prove is that for $n$ large enough

$$
\mathbf{P}\left\{\int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)>\frac{\epsilon}{2}\right\} \leq e^{-n \epsilon^{2} /\left(32 \rho^{2}\right)}
$$

We use a decomposition as in the proof of strong consistency of the histogram rule:

$$
\begin{align*}
& \left|\eta(x)-\eta_{n}(x)\right| \\
& \quad=\mathbf{E}\left|\eta(x)-\eta_{n}(x)\right|+\left(\left|\eta(x)-\eta_{n}(x)\right|-\mathbf{E}\left|\eta(x)-\eta_{n}(x)\right|\right) . \tag{10.1}
\end{align*}
$$

To handle the first term on the right-hand side, fix $\epsilon^{\prime}>0$, and let $r: \mathcal{R}^{d} \rightarrow \mathcal{R}$ be a continuous function of bounded support satisfying

$$
\int|\eta(x)-r(x)| \mu(d x)<\epsilon^{\prime}
$$

Obviously, we can choose the function $r$ such that $0 \leq r(x) \leq 1$ for all $x \in \mathcal{R}^{d}$. Then we have the following simple upper bound:

$$
\begin{align*}
& \mathbf{E}\left|\eta(x)-\eta_{n}(x)\right| \\
& \leq \quad|\eta(x)-r(x)|+\left|r(x)-\frac{\mathbf{E}\left\{r(X) K_{h}(x-X)\right\}}{\mathbf{E} K_{h}(x-X)}\right| \\
& \quad+\left|\frac{\mathbf{E}\left\{r(X) K_{h}(x-X)\right\}}{\mathbf{E} K_{h}(x-X)}-\mathbf{E} \eta_{n}(x)\right|+\mathbf{E}\left|\mathbf{E} \eta_{n}(x)-\eta_{n}(x)\right| . \tag{10.2}
\end{align*}
$$

Next we bound the integral of each term on the right-hand side of the inequality above.

FIRST TERM: By the definition of $r$,

$$
\int|\eta(x)-r(x)| \mu(d x)<\epsilon^{\prime}
$$

SECOND TERM: Since $r(x)$ is continuous and zero outside of a bounded set, it is also uniformly continuous, that is, there exists a $\delta>0$ such that $\|x-y\|<\delta$ implies $|r(x)-r(y)|<\epsilon^{\prime}$. Also, $r(x)$ is bounded. Thus,

$$
\begin{aligned}
& \int\left|r(x)-\frac{\mathbf{E}\left\{r(X) K_{h}(x-X)\right\}}{\mathbf{E} K_{h}(x-X)}\right| \mu(d x) \\
&= \int\left|r(x)-\frac{\int r(y) K_{h}(x-y) \mu(d y)}{\mathbf{E} K_{h}(x-X)}\right| \mu(d x) \\
& \leq \iint_{S_{x, \delta}} \frac{K_{h}(x-y)}{\mathbf{E} K_{h}(x-X)}|r(x)-r(y)| \mu(d y) \mu(d x) \\
&+\iint_{S_{x, \delta}^{c}} \frac{K_{h}(x-y)}{\mathbf{E} K_{h}(x-X)}|r(x)-r(y)| \mu(d y) \mu(d x) \\
& \leq \int\left(\int_{S_{x, \delta}} \frac{K_{h}(x-y)}{\mathbf{E} K_{h}(x-X)} \mu(d y) \sup _{z \in S_{x, \delta}}|r(x)-r(z)|\right) \mu(d x) \\
&+\int\left(\int_{S_{x, \delta}^{c}} \frac{K_{h}(x-y)}{\mathbf{E} K_{h}(x-X)} \mu(d y)\right) \mu(d x) .
\end{aligned}
$$

In the last step we used the fact that $\sup _{x, y}|r(x)-r(y)| \leq 1$. Clearly, we have $\int_{S_{x, \delta}} \frac{K_{h}(x-y)}{\mathbf{E} K_{h}(x-X)} \mu(d y) \leq 1$, and by the uniform continuity of $r(x), \sup _{z \in S_{x, \delta}} \mid r(x)-$ $r(z) \mid<\epsilon^{\prime}$. Thus, the first term at the end of the chain of inequalities above is bounded by $\epsilon^{\prime}$. The second term converges to zero since $h<\delta$ for all $n$ large enough, which in turn implies $\int_{S_{x, \delta}^{c}} \frac{K_{h}(x-y)}{\mathbf{E} K_{h}(x-X)} \mu(d y)=0$. (This is obvious for the naïve kernel. For regular kernels, convergence to zero follows from Problem 10.15.) The convergence of the integral (with respect to $\mu(d x)$ ) follows from the dominated
convergence theorem. In summary, we have shown that

$$
\limsup _{n \rightarrow \infty} \int\left|r(x)-\frac{\mathbf{E}\left\{r(X) K_{h}(x-X)\right\}}{\mathbf{E} K_{h}(x-X)}\right| \mu(d x) \leq \epsilon^{\prime}
$$

Third term:

$$
\begin{aligned}
\int & \left|\frac{\mathbf{E}\left\{r(X) K_{h}(x-X)\right\}}{\mathbf{E} K_{h}(x-X)}-\mathbf{E} \eta_{n}(x)\right| \mu(d x) \\
& =\int\left|\frac{\int(r(y)-\eta(y)) K_{h}(x-y) \mu(d y)}{\mathbf{E} K_{h}(x-X)}\right| \mu(d x) \\
& \leq \iint|r(y)-\eta(y)| \frac{K_{h}(x-y)}{\int K_{h}(x-z) \mu(d z)} \mu(d y) \mu(d x) \\
& =\int\left(\int \frac{K_{h}(x-y)}{\int K_{h}(x-z) \mu(d z)} \mu(d x)\right)|r(y)-\eta(y)| \mu(d y)
\end{aligned}
$$

(by Fubini's theorem)
$\leq \int \rho|r(y)-\eta(y)| \mu(d y) \leq \rho \epsilon^{\prime}$,
where in the last two steps we used the covering lemma (see Lemma 10.1 for the naïve kernel, and Problem 10.14 for general kernels), and the definition of $r(x)$. $\rho$ is the constant appearing in the covering lemma:

FOURTH TERM: We show that

$$
\mathbf{E}\left\{\int\left|\mathbf{E} \eta_{n}(x)-\eta_{n}(x)\right| \mu(d x)\right\} \rightarrow 0
$$

For the naïve kernel, we have

$$
\begin{aligned}
\mathbf{E} & \left\{\left|\mathbf{E} \eta_{n}(x)-\eta_{n}(x)\right|\right\} \\
& \leq \sqrt{\mathbf{E}\left\{\left|\mathbf{E} \eta_{n}(x)-\eta_{n}(x)\right|^{2}\right\}} \\
& =\sqrt{\frac{\mathbf{E}\left\{\left(\sum_{j=1}^{n}\left(Y_{j} K_{h}\left(x-X_{j}\right)-\mathbf{E}\left\{Y K_{h}(x-X)\right\}\right)\right)^{2}\right\}}{n^{2}\left(\mathbf{E} K_{h}(x-X)\right)^{2}}} \\
& =\sqrt{\frac{\mathbf{E}\left\{\left(Y K_{h}(x-X)-\mathbf{E}\left\{Y K_{h}(x-X)\right\}\right)^{2}\right\}}{n\left(\mathbf{E} K_{h}(x-X)\right)^{2}}} \\
& \leq \sqrt{\frac{\mathbf{E}\left\{\left(Y K_{h}(x-X)\right)^{2}\right\}}{n\left(\mathbf{E} K_{h}(x-X)\right)^{2}}} \leq \sqrt{\frac{\mathbf{E}\left\{\left(K_{h}(x-X)\right)^{2}\right\}}{n\left(\mathbf{E} K_{h}(x-X)\right)^{2}}} \\
& \leq \sqrt{\frac{\mathbf{E} K((x-X) / h)}{n(\mathbf{E} K((x-X) / h))^{2}}} \leq \sqrt{\frac{1}{n \mu\left(S_{x, h}\right)}},
\end{aligned}
$$

where we used the Cauchy-Schwarz inequality, and properties of the naïve kernel. Extension to regular kernels is straightforward.

Next we use the inequality above to show that the integral converges to zero. Divide the integral over $\mathcal{R}^{d}$ into two terms, namely an integral over a large ball $S$ centered at the origin, of radius $R>0$, and an integral over $S^{c}$. For the integral outside of the ball we have

$$
\int_{S^{c}} \mathbf{E}\left\{\left|\mathbf{E} \eta_{n}(x)-\eta_{n}(x)\right|\right\} \mu(d x) \leq 2 \int_{S^{c}} \mathbf{E} \eta_{n}(x) \mu(d x) \rightarrow 2 \int_{S^{c}} \eta(x) \mu(d x)
$$

with probability one as $n \rightarrow \infty$, which can be shown in the same way we proved

$$
\int_{\mathcal{R}^{d}} \mathbf{E} \eta_{n}(x) \mu(d x) \rightarrow \int_{\mathcal{R}^{d}} \eta(x) \mu(d x)
$$

(see the first, second, and third terms of (10.2)). Clearly, the radius $R$ of the ball $S$ can be chosen such that $2 \int_{S^{c}} \eta(x) \mu(d x)<\epsilon / 4$. To bound the integral over $S$ we employ Lemma 10.2:

$$
\begin{aligned}
\int_{S} \mathbf{E}\left\{\left|\mathbf{E} \eta_{n}(x)-\eta_{n}(x)\right|\right\} \mu(d x) & \leq \\
& \quad \frac{1}{\sqrt{n}} \int_{S} \frac{1}{\sqrt{\mu\left(S_{x, h}\right)}} \mu(d x) \\
\leq & \frac{1}{\sqrt{n}}\left(1+\frac{R}{h}\right)^{d / 2} c_{d} \\
\rightarrow & \left.0 \text { (since by assumption } n h^{d} \rightarrow \infty\right) .
\end{aligned}
$$

Therefore, if $n$ is sufficiently large, then for the first term on the right-hand side of (10.1) we have

$$
\mathbf{E}\left\{\int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)\right\}<\epsilon^{\prime}(\rho+3)=\epsilon / 4
$$

if we take $\epsilon^{\prime}=\epsilon /(4 \rho+12)$.
It remains to show that the second term on the right-hand side of (10.1) is small with large probability. To do this, we use McDiarmid's inequality (Theorem 9.2) for

$$
\int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)-\mathbf{E}\left\{\int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)\right\}
$$

Fix the training data at $\left(\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)\right)$ and replace the $i$-th pair $\left(x_{i}, y_{i}\right)$ by $\left(\widehat{x}_{i}, \widehat{y}_{i}\right)$, changing the value of $\eta_{n}(x)$ to $\eta_{n i}^{*}(x)$. Clearly, by the covering lemma (Lemma 10.1),

$$
\begin{aligned}
& \int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)-\int\left|\eta(x)-\eta_{n i}^{*}(x)\right| \mu(d x) \\
& \quad \leq \int\left|\eta_{n}(x)-\eta_{n i}^{*}(x)\right| \mu(d x) \leq \sup _{y \in \mathcal{R}^{d}} \int \frac{2 K_{h}(x-y)}{n \mathbf{E} K_{h}(x-X)} \mu(d x) \\
& \quad \leq \frac{2 \rho}{n}
\end{aligned}
$$

So by Theorem 9.2,

$$
\begin{aligned}
& \mathbf{P}\left\{\int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)>\frac{\epsilon}{2}\right\} \\
& \quad \leq \mathbf{P}\left\{\int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)-\mathbf{E}\left\{\int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)\right\}>\frac{\epsilon}{4}\right\} \\
& \quad \leq e^{-n \epsilon^{2} /\left(32 \rho^{2}\right)}
\end{aligned}
$$

The proof is now completed.

### 10.3 Potential Function Rules

Kernel classification rules may be formulated in terms of the so-called potential function rules. These rules were originally introduced and studied by Bashkirov, Braverman and Muchnik (1964), Aizerman, Braverman and Rozonoer (1964c; 1964b; 1964a; 1970), Braverman (1965), and Braverman and Pyatniskii (1966). The original idea was the following: put a unit of positive electrical charge at every data point $X_{i}$, where $Y_{i}=1$, and a unit of negative charge, at data points $X_{i}$ where $Y_{i}=0$. The resulting potential field defines an intuitively appealing rule: the decision at a point $x$ is one if the potential at that point is positive, and zero if it is negative. This idea leads to a rule that can be generalized to obtain rules of the form

$$
g_{n}(x)= \begin{cases}0 & \text { if } f_{n}(x) \leq 0 \\ 1 & \text { otherwise }\end{cases}
$$

where

$$
f_{n}(x)=\sum_{i=1}^{n} r_{n, i}\left(D_{n}\right) K_{n, i}\left(x, X_{i}\right)
$$

where the $K_{n, i}$ 's describe the potential field around $X_{i}$, and the $r_{n, i}$ 's are their weights. Rules that can be put into this form are often called potential function rules. Here we give a brief survey of these rules.

Kernel rules. Clearly, kernel rules studied in the previous section are potential function rules with

$$
K_{n, i}(x, y)=K\left(\frac{x-y}{h_{n}}\right), \quad \text { and } \quad r_{n, i}\left(D_{n}\right)=2 Y_{i}-1
$$

Here $K$ is a fixed kernel function, and $h_{1}, h_{2}, \ldots$ is a sequence of positive numbers.
Histogram rules. Similarly, histogram rules (see Chapters 6 and 9) can be put in this form, by choosing

$$
K_{n, i}(x, y)=I_{\left\{y \in A_{n}(x)\right\}}, \quad \text { and } \quad r_{n, i}\left(D_{n}\right)=2 Y_{i}-1
$$

Recall that $A_{n}(x)$ denotes the cell of the partition in which $x$ falls.
POLYNOMIAL DISCRIMINANT FUNCTIONS. Specht (1967) suggested applying a polynomial expansion to the kernel $K\left(\frac{x-y}{h}\right)$. This led to the choice

$$
K_{n, i}(x, y)=\sum_{j=1}^{k} \psi_{j}(x) \psi_{j}(y), \quad \text { and } \quad r_{n, i}\left(D_{n}\right)=2 Y_{i}-1
$$

where $\psi_{1}, \ldots, \psi_{k}$ are fixed real-valued functions on $\mathcal{R}^{d}$. When these functions are polynomials, the corresponding classifier $g_{n}$ is called a polynomial discriminant function. The potential function rule obtained this way is a generalized linear rule (see Chapter 17) with

$$
f_{n}(x)=\sum_{j=1}^{k} a_{n, j} \psi_{j}(x)
$$

where the coefficients $a_{n, j}$ depend on the data $D_{n}$ only, through

$$
a_{n, j}=\sum_{i=1}^{n}\left(2 Y_{i}-1\right) \psi_{j}\left(X_{i}\right)
$$

This choice of the coefficients does not necessarily lead to a consistent rule, unless the functions $\psi_{1}, \ldots, \psi_{k}$ are allowed to change with $n$, or $k$ is allowed to vary with $n$. Nevertheless, the rule has some computational advantages over kernel rules. In many practical situations there is enough time to preprocess the data $D_{n}$, but once the observation $X$ becomes known, the decision has to be made very quickly. Clearly, the coefficients $a_{n, 1}, \ldots, a_{n, n}$ can be computed by knowing the training data $D_{n}$ only, and if the values $\psi_{1}(X), \ldots, \psi_{k}(X)$ are easily computable, then $f_{n}(X)$ can be computed much more quickly than in a kernel-based decision, where all $n$ terms of the sum have to be computed in real time, if no preprocessing is done. However, using preprocessing of the data may also help with kernel rules, especially when $d=1$. For a survey of computational speed-up with kernel methods, see Devroye and Machell (1985).

Recursive kernel rules. Consider the choice

$$
\begin{equation*}
K_{n, i}(x, y)=K\left(\frac{x-y}{h_{i}}\right), \quad \text { and } \quad r_{n, i}\left(D_{n}\right)=2 Y_{i}-1 \tag{10.3}
\end{equation*}
$$

Observe that the only difference between this and the ordinary kernel rule is that in the expression of $K_{n, i}$, the smoothing parameter $h_{n}$ is replaced with $h_{i}$. With this change, we can compute the rule recursively by observing that

$$
f_{n+1}(x)=f_{n}(x)+\left(2 Y_{n+1}-1\right) K\left(\frac{x-X_{n+1}}{h_{n+1}}\right)
$$

The computational advantage of this rule is that if one collects additional data, then the rule does not have to be entirely recomputed. It can be adjusted using the formula above. Consistency properties of this rule were studied by Devroye and Wagner (1980b), Krzyżak and Pawlak (1984a), Krzyżak (1986), and Greblicki and Pawlak (1987). Several similar recursive kernel rules have been studied in the literature. Wolverton and Wagner (1969b), Greblicki (1974), and Krzyżak and Pawlak (1983), studied the situation when

$$
\begin{equation*}
K_{n, i}(x, y)=\frac{1}{h_{i}^{d}} K\left(\frac{x-y}{h_{i}}\right) \quad \text { and } \quad r_{n, i}\left(D_{n}\right)=2 Y_{i}-1 . \tag{10.4}
\end{equation*}
$$

The corresponding rule can be computed recursively by

$$
f_{n+1}(x)=f_{n}(x)+\left(2 Y_{n+1}-1\right) \frac{1}{h_{n+1}^{d}} K\left(\frac{x-X_{n+1}}{h_{n+1}}\right)
$$

Motivated by stochastic approximation methods (see Chapter 17), Révész (1973) suggested and studied the rule obtained from

$$
f_{n+1}(x)=f_{n}(x)+\frac{1}{n+1}\left(2 Y_{n+1}-1-f_{n}(x)\right) \frac{1}{h_{n+1}^{d}} K\left(\frac{x-X_{n+1}}{h_{n+1}}\right)
$$

A similar rule was studied by Györfi (1981):

$$
f_{n+1}(x)=f_{n}(x)+\frac{1}{n+1}\left(2 Y_{n+1}-1-f_{n}(x)\right) K\left(\frac{x-X_{n+1}}{h_{n+1}}\right)
$$

## Problems and Exercises

Problem 10.1. Let $K$ be a nonnegative kernel with compact support on $[-1,1]$. Show that for some distribution, $h \rightarrow 0$ is necessary for consistency of the kernel rule. To this end, consider the following example. Given $Y=0, X$ has a geometric distribution on $\delta, \delta^{3}, \delta^{5}, \ldots$, and given $Y=1, X$ has a geometric distribution on $\delta^{2}, \delta^{4}, \delta^{6}, \ldots$ Then show that to obtain $\mathbf{E} L_{n} \rightarrow L^{*}=0$, it is necessary that $h \rightarrow 0$.

Problem 10.2. Kernel density estimation. Let $X_{1}, \ldots, X_{n}$ be i.i.d. random variables in $\mathcal{R}^{d}$ with density $f$. Let $K$ be a kernel function integrating to one, and let $h_{n}>0$ be a smoothing factor. The kernel density estimate is defined by

$$
f_{n}(x)=\frac{1}{n h_{n}^{d}} \sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h_{n}}\right)
$$

(Rosenblatt (1956), Parzen (1962)). Prove that the estimate is weakly universally consistent in $L_{1}$ if $h_{n} \rightarrow 0$ and $n h_{n}^{d} \rightarrow \infty$ as $n \rightarrow \infty$. Hint: Proceed as in Problem 6.2.

Problem 10.3. Strong consistency of kernel density estimation. Let $X_{1}, \ldots, X_{n}$ be i.i.d. random variables in $\mathcal{R}^{d}$ with density $f$. Let $K$ be a nonnegative function integrating to one (a kernel) and $h>0$ a smoothing factor. As in the previous exercise, the kernel density estimate is defined by

$$
f_{n}(x)=\frac{1}{n h^{d}} \sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h}\right) .
$$

Prove for the $L_{1}$-error of the estimate that

$$
\mathbf{P}\left\{\left|\int\right| f_{n}(x)-f(x)\left|d x-\mathbf{E} \int\right| f_{n}(x)-f(x)|d x|>\epsilon\right\} \leq 2 e^{-n \epsilon^{2} / 2}
$$

(Devroye (1991a)). Conclude that weak $L_{1}$-consistency of the estimate implies strong consistency (see Problem 10.2). This is a way to show that weak and strong $L_{1}$-consistencies of the kernel density estimate are equivalent (Devroye (1983).) Also, for $d=1$, since if $K$ is nonnegative, then $\mathbf{E} \int\left|f_{n}(x)-f(x)\right| d x$ cannot converge to zero faster than $n^{-2 / 5}$ for any density (see Devroye and Györf (1985)), therefore, the inequality above implies that for any density

$$
\lim _{n \rightarrow \infty} \frac{\int\left|f_{n}(x)-f(x)\right| d x}{\mathbf{E} \int\left|f_{n}(x)-f(x)\right| d x}=0
$$

with probability one (Devroye (1988d)). This property is called the relative stability of the $L_{1}$ error. It means that the asymptotic behavior of the $L_{1}$-error is the same as that of its expected value. Hint: Use McDiarmid's inequality.

Problem 10.4. If $\int K<0$, show that under the assumption that $\mu$ has a density $f$, and that $h \rightarrow 0, n h^{d} \rightarrow \infty$, the kernel rule has

$$
\lim _{n \rightarrow \infty} \mathbf{E} L_{n}=\mathbf{E}\{\max (\eta(X), 1-\eta(X))\}=1-L^{*}
$$

Thus, the rule makes the wrong decisions, and such kernels should be avoided. Hint: You may use the fact that for the kernel density estimate with kernel $L$ satisfying $\int L=1$,

$$
\int\left|f_{n}(x)-f(x)\right| d x \rightarrow 0
$$

with probability one, if $h \rightarrow 0$ and $n h^{d} \rightarrow \infty$ (see Problems 10.2 and 10.3).
Problem 10.5. Consider a devilish kernel that attaches counterproductive weight to the origin:

$$
K(x)= \begin{cases}-1 & \text { if }\|x\| \leq 1 / 3 \\ 1 & \text { if } 1 / 3<\|x\| \leq 1 \\ 0 & \text { if }\|x\|>1\end{cases}
$$

Assume that $h \rightarrow 0$, yet $n h^{d} \rightarrow \infty$. Assume that $L^{*}=0$. Show that $L_{n} \rightarrow 0$ with probability one. CONCESSION: if you find that you can't handle the universality, try first proving the statement for strictly separable distributions.

Problem 10.6. Show that for the distribution depicted in Figure 10.2, the kernel rule with kernel $K(u)=\left(1-u^{2}\right) I_{\{|u| \leq 1\}}$ is consistent whenever $h$, the smoothing factor, remains fixed and $0<h \leq 1 / 2$.

Problem 10.7. The limit for fixed $h$. Consider a kernel rule with fixed $h \equiv 1$, and fixed kernel $K$. Find a simple argument that proves

$$
\lim _{n \rightarrow \infty} \mathbf{E} L_{n}=L_{\infty} ;
$$

where $L_{\infty}$ is the probability of error for the decision $g_{\infty}$ defined by

$$
g_{\infty}(x)= \begin{cases}0 & \text { if } \mathbf{E}\{K(x-X)(2 \eta(X)-1)\} \leq 0 \\ 1 & \text { if } \mathbf{E}\{K(x-X)(2 \eta(X)-1)\}>0\end{cases}
$$

Find a distribution such that for the window kernel, $L_{\infty}=1 / 2$, yet $L^{*}=0$. Is there such a distribution for any kernel? Hint: Try proving a convergence result at each $x$ by invoking the law of large numbers, and then replace $x$ by $X$.

Problem 10.8. Show that the conditions $h_{n} \rightarrow 0$ and $n h^{d} \rightarrow \infty$ of Theorem 10.1 are not necessary for consistency, that is, exhibit a distribution such that the kernel rule is consistent with $h_{n}=1$, and exhibit another distribution for which the kernel rule is consistent with $h_{n} \sim 1 / n^{1 / d}$.

Problem 10.9. Prove that the conditions $h_{n} \rightarrow 0$ and $n h^{d} \rightarrow \infty$ of Theorem 10.1 are necessary for universal consistency, that is, show that if one of these conditions are violated then there is a distribution for which the kernel rule is not consistent (Krzyżak (1991)).

Problem 10.10. This exercise provides an argument in favor of monotonicity of the kernel $K$. In $\mathcal{R}^{2}$, find a nonatomic distribution for ( $X, Y$ ), and a positive kernel with $\int K>0$, $K$ vanishing off $S_{0, \delta}$ for some $\delta>0$, such that for all $h>0$, and all $n$, the kernel rule has $\mathbf{E} L_{n}=1 / 2$, while $L^{*}=0$. This result says that the condition $K(x) \geq b I_{\left\{S_{0, \delta}\right\}}$ for some $b>0$ in the universal consistency theorem cannot be abolished altogether.

Problem 10.11. With $K$ as in the previous problem, and taking $h=1$, show that

$$
\lim _{n \rightarrow \infty} \mathbf{E} L_{n}=L^{*}
$$

under the following conditions:
(1) $K$ has compact support vanishing off $S_{0, \delta}, K \geq 0$, and $K \geq b I_{\left\{S_{0, \epsilon}\right\}}$ for some $\epsilon>0$.
(2) We say that we have agreement on $S_{x, \delta}$ when for all $z \in S_{x, \delta}$, either $\eta(z) \leq 1 / 2$, or $\eta(z) \geq 1 / 2$. We ask that $\mathbf{P}\left\{\right.$ Agreement on $\left.S_{X, \delta}\right\}=1$.

Problem 10.12. The previous exercise shows that at points where there is agreement, we make asymptotically the correct decision with kernels with fixed smoothing factor. Let $D$ be the set $\{x: \eta(x)=1 / 2\}$, and let the $\delta$-neighborhood of $D$ be defined by $D_{\delta}=\{y$ : $\|y-x\| \leq \delta$ for some $x \in D\}$. Let $\mu$ be the probability measure for $X$. Take $K, h$ as in the previous exercise. Noting that $x \notin D_{\delta}$ means that we have agreement on $S_{x, \delta}$, show that for all distributions of $(X, Y)$,

$$
\limsup _{n \rightarrow \infty} \mathbf{E} L_{n} \leq L^{*}+\mu\left(D_{\delta}\right) .
$$



Figure 10.5. $\delta$-neighborhood of a set $D$.

Problem 10.13. Continuation. Clearly, $\mu\left(D_{\delta}\right) \rightarrow 0$ as $\delta \rightarrow 0$ when $\mu(D)=0$. Convince yourself that $\mu(D)=0$ for most problems. If you knew how fast $\mu\left(D_{\delta}\right)$ tended to zero, then the previous exercise would enable you to pick $h$ as a function of $n$ such that $h \rightarrow 0$ and such that the upper bound for $\mathbf{E} L_{n}$ obtained by analogy from the previous exercise is approximately minimal. If in $\mathcal{R}^{d}, D$ is the surface of the unit ball, $X$ has a bounded density $f$, and $\eta$ is Lipschitz, determine a bound for $\mu\left(D_{\delta}\right)$. By considering the proof of the universal consistency theorem, show how to choose $h$ such that

$$
\mathbf{E} L_{n}-L^{*}=O\left(\frac{1}{n^{1 /(2+d)}}\right) .
$$

Problem 10.14. Extension of the covering lemma (Lemma 10.1) to regular kernels. Let $K$ be a regular kernel, and let $\mu$ be an arbitrary probability measure. Prove that there exists a finite constant $\rho=\rho(K)$ only depending upon $K$ such that for any $y$ and $h$

$$
\int \frac{K_{h}(x-y)}{\int K_{h}(x-z) \mu(d z)} \mu(d x) \leq \rho
$$

(Devroye and Krzyżak (1989)). Hint: Prove this by checking the following:
(1) First take a bounded overlap cover of $\mathcal{R}^{d}$ with translates of $S_{0, r / 2}$, where $r>0$ is the constant appearing in the definition of a regular kernel. This cover has an infinite number of member balls, but every $x$ gets covered at most $k_{1}$ times where $k_{1}$ depends upon $d$ only.
(2) The centers of the balls are called $x_{i}, i=1,2, \ldots$ The integral condition on $K$ implies that

$$
\sum_{i=1}^{\infty} \sup _{x \in x_{i}+S_{0, r / 2}} K(x) \leq \frac{k_{1}}{\int_{S_{0, r / 2}} d x} \int \sup _{y \in x+S_{0, r / 2}} K(y) d x \leq k_{2}
$$

for another finite constant $k_{2}$.
(3) Show that

$$
K_{h}(x-y) \leq \sum_{i=1}^{\infty} \sup _{x \in y+h x_{i}+S_{0, r h / 2}} K_{h}(x-y) I_{\left[x \in y+h x_{i}+S_{0, r h / 2}\right]},
$$

and

$$
\int K_{h}(x-z) \mu(d z) \geq b \mu\left(y+h x_{i}+S_{0, r h / 2}\right)\left(x \in y+h x_{i}+S_{0, r h / 2}\right) .
$$

From (c) conclude

$$
\begin{aligned}
\int & \frac{K_{h}(x-y)}{\int K_{h}(x-z) \mu(d z)} \mu(d x) \\
& \leq \sum_{i=1}^{\infty} \int_{x \in y+h x_{i}+S_{0, r h / 2}} \frac{\sup _{z \in h x_{i}+S_{0, r h / 2}} K_{h}(z)}{b \mu\left(y+h x_{i}+S_{0, r h / 2}\right)} \mu(d x) \\
& =\sum_{i=1}^{\infty} \frac{\mu\left(y+h x_{i}+S_{0, r h / 2}\right) \sup _{z \in h x_{i}+S_{0, r h / 2}} K_{h}(z)}{b \mu\left(y+h x_{i}+S_{0, r h / 2}\right)} \\
& =\frac{1}{b} \sum_{i=1}^{\infty} \sup _{z \in h x_{i}+S_{0, r h / 2}} K_{h}(z) \leq \frac{k_{2}}{b},
\end{aligned}
$$

where $k_{2}$ depends on $K$ and $d$ only.
Problem 10.15 . Let $K$ be a regular kernel, and let $\mu$ be an arbitrary probability measure. Prove that for any $\delta>0$

$$
\lim _{h \rightarrow 0} \sup _{y} \int \frac{K_{h}(x-y) I_{\{||x-y| \geq \delta\}}}{\int K_{h}(x-z) \mu(d z)} \mu(d x)=0 .
$$

Hint: Substitute $K_{h}(z)$ in the proof of Problem 10.15 by $K_{h}(z) I(\|z\| \geq \delta)$ and notice that

$$
\sup _{y} \int \frac{K_{h}(x-y) I_{\{\|x-y\| \geq \delta\}}}{\int K_{h}(x-z) \mu(d z)} \mu(d x) \leq \sum_{i=1}^{\infty} \sup _{z \in h x_{i}+\delta_{0, r h / 2}} K_{h}(z) I(\|z\| \geq \delta) \rightarrow 0
$$

as $h \rightarrow 0$.
Problem 10.16. Use Problems 10.14 and 10.15 to extend the proof of Theorem 10.1 for arbitrary regular kernels.

Problem 10.17. Show that the constant $\beta_{d}$ in Lemma 10.1 is never more than $4^{d}$.
Problem 10.18. Show that if $L \geq 0$ is a bounded function that is monotonically decreasing on $[0, \infty)$ with the property that $\int u^{d-1} L(u) d u<\infty$, and if $K: \mathcal{R}^{d} \rightarrow[0, \infty)$ is a function with $K(x) \leq L(\|x\|)$, then $K$ is regular.

Problem 10.19. Find a kernel $K \geq 0$ that is monotonically decreasing along rays (i.e., $K(r x) \leq K(x)$ for all $x \in \mathcal{R}^{d}$ and all $r \geq 1$ ) such that $K$ is not regular. (This exercise is intended to convince you that it is very difficult to find well-behaved kernels that are not regular.)

Problem 10.20. Let $K(x)=L(\|x\|)$ for some bounded function $L \geq 0$. Show that $K$ is regular if $L$ is decreasing on $[0, \infty)$ and $\int K(x) d x<\infty$. Conclude that the Gaussian and Cauchy kernels are regular.

Problem 10.21. Regularity of the kernel is not necessary for universal consistency. Investigate universal consistency with a nonintegrable kernel-that is, for which $\int K(x) d x=\infty$ such as $K(x)=1 /(1+|x|)$. Greblicki, Krzyżak, and Pawlak (1984) proved consistency of the kernel rule with smoothing factor $h_{n}$ satisfying $h_{n} \rightarrow 0$ and $n h_{n}^{d} \rightarrow \infty$ if the kernel $K$ satisfies the following conditions: $K(x) \geq c I_{\{\|x\| \leq 1\}}$ for some $c>0$ and for some $c_{1}, c_{2}>0, c_{1} H(\|x\|) \leq K(x) \leq c_{2} H(\|x\|)$, where $H$ is a nonincreasing function on $[0, \infty)$ with $u^{d} H(u) \rightarrow 0$ as $u \rightarrow \infty$.

Problem 10.22. Consider the kernel rule with kernel $K(x)=1 /\|x\|^{r}, r>0$. Such kernels are useless for atomic distributions unless we take limits and define $g_{n}$ as usual when $x \notin S$, the collection of points $z$ with $X_{i}=X_{j}=z$ for some pair $(i \neq j)$. For $x \in S$, we take a majority vote over the $Y_{i}$ 's for which $X_{i}=x$. Discuss the weak universal consistency of this rule, which has the curious property that $g_{n}$ is invariant to the smoothing factor $h$-so, we might as well set $h=1$ without loss of generality. Note also that for $r \geq d$, $\int_{S_{0,1}} K(x) d x=\infty$, and for $r \leq d, \int_{S_{0,1}^{c}} K(x) d x=\infty$, where $S_{0,1}$ is the unit ball of $\mathcal{R}^{d}$ centered at the origin. In particular, if $r \leq d$, by considering $X$ uniform on $S_{0,1}$ if $Y=1$ and $X$ uniform on the surface of $S_{0.1}$ if $Y=0$, show that even though $L^{*}=0$, the probability of error of the rule may tend to a nonzero limit for certain values of $\mathbf{P}\{Y=1\}$. Hence, the rule is not universally consistent. For $r \geq d$, prove or disprove the weak universal consistency, noting that the rules' decisions are by-and-large based on the few nearest neighbors. Prove the rule is weakly consistent for all $r \leq d$ whenever $X$ has a density.

Problem 10.23. Assume that the class densities coincide, that is, $f_{0}(x)=f_{1}(x)$ for every $x \in \mathcal{R}$, and assume $p=\mathbf{P}\{Y=1\}>1 / 2$. Show that the expected probability of error of the kernel rule with $K \equiv 1$ is smaller than that with any unimodal regular kernel for every $n$ and $h$ small enough. Exhibit a distribution such that the kernel rule with a symmetric kernel such that $K(|x|)$ is monotone increasing has smaller expected error probability than that with any unimodal regular kernel.

Problem 10.24. Scaling. Assume that the kernel $K$ can be written into the following product form of one-dimensional kernels:

$$
K(x)=K\left(x^{(1)}, \ldots, x^{(d)}\right)=\prod_{i=1}^{d} K_{i}\left(x^{(i)}\right) .
$$

Assume also that $K$ is regular. One can use different smoothing factors along the different coordinate axes to define a kernel rule by

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n}\left(2 Y_{i}-1\right) \prod_{j=1}^{d} K_{j}\left(\frac{x^{(j)}-X_{i}^{(j)}}{h_{j n}}\right) \leq 0 \\ 1 & \text { otherwise, }\end{cases}
$$

where $X_{i}^{(j)}$ denotes the $j$-th component of $X_{i}$. Prove that $g_{n}$ is strongly universally consistent if $h_{i n} \rightarrow 0$ for all $i=1, \ldots, d$, and $n h_{1 n} h_{2 n} \ldots h_{d n} \rightarrow \infty$.

Problem 10.25. Let $K:[0, \infty) \rightarrow[0, \infty)$ be a function, and $\Sigma$ a symmetric positive definite $d \times d$ matrix. For $x \in \mathcal{R}^{d}$ define $K^{\prime}(x)=K\left(x^{T} \Sigma x\right)$. Find conditions on $K$ such that the kernel rule with kernel $K^{\prime}$ is universally consistent.

Problem 10.26. Prove that the recursive kernel rule defined by (10.4) is strongly universally consistent if $K$ is a regular kernel, $h_{n} \rightarrow 0$, and $n h_{n}^{d} \rightarrow \infty$ as $n \rightarrow \infty$ (Krzyżak and Pawlak (1983)).

Problem 10.27. Show that the recursive kernel rule of (10.3) is strongly universally consistent whenever $K$ is a regular kernel, $\lim _{n \rightarrow \infty} h_{n}=0$, and $\sum_{n=1}^{\infty} h_{n}^{d}=\infty$ (Greblicki and Pawlak (1987)). Note: Greblicki and Pawlak (1987) showed convergence under significantly weaker assumptions on the kernel. They assume that $K(x) \geq c I_{\{\|x\| \leq 1\}}$ for some $c>0$ and that for some $c_{1}, c_{2}>0, c_{1} H(\|x\|) \leq K(x) \leq c_{2} H(\|x\|)$, where $H$ is a nonincreasing function on $[0, \infty)$ with $u^{d} H(u) \rightarrow 0$ as $u \rightarrow \infty$. They also showed that under the additional
assumption $\int K(x) d x<\infty$ the following conditions on $h_{n}$ are necessary and sufficient for universal consistency:

$$
\frac{\sum_{i=1}^{n} h_{i}^{d} I_{\left\{h_{i}>a\right\}}}{\sum_{i=1}^{n} h_{i}^{d}} \rightarrow 0 \text { for all } a>0 \text { and } \sum_{i=1}^{\infty} h_{i}^{d}=\infty .
$$

Problem 10.28. Open-ended problem. Let $\mathbf{P}\{Y=1\}=1 / 2$. Given $Y=1$, let $X$ be uniformly distributed on $[0,1]$. Given $Y=0$, let $X$ be atomic on the rationals with the following distribution: let $X=V / W$, where $V$ and $W$ are independent identically distributed, and $\mathbf{P}\{V=i\}=1 / 2^{i}, i \geq 1$. Consider the kernel rule with the window kernel. What is the behavior of the smoothing factor $h_{n}^{*}$ that minimizes the expected probability of error $\mathbf{E} L_{n}$ ?

## 11

## Consistency of the $k$-Nearest Neighbor Rule

In Chapter 5 we discuss results about the asymptotic behavior of $k$-nearest neighbor classification rules, where the value of $k$-the number of neighbors taken into account at the decision-is kept at a fixed number as the size of the training data $n$ increases. This choice leads to asymptotic error probabilities smaller than $2 L^{*}$, but no universal consistency. In Chapter 6 we showed that if we let $k$ grow to infinity as $n \rightarrow \infty$ such that $k / n \rightarrow 0$, then the resulting rule is weakly consistent. The main purpose of this chapter is to demonstrate strong consistency, and to discuss various versions of the rule.

We are not concerned here with the data-based choice of $k$-that subject deserves a chapter of its own (Chapter 26). We are also not tackling the problem of the selection of a suitable-even data-based-metric. At the end of this chapter and in the exercises, we draw the attention to 1 -nearest neighbor relabeling rules, which combine the computational comfort of the 1 -nearest neighbor rule with the asymptotic performance of variable- $k$ nearest neighbor rules.

Consistency of $k$-nearest neighbor classification, and corresponding regression and density estimation has been studied by many researchers. See Fix and Hodges (1951; 1952), Cover (1968a), Stone (1977), Beck (1979), Györfi and Györfi (1975), Devroye (1981a; 1982b), Collomb (1979; 1980; 1981), Bickel and Breiman (1983), Mack (1981), Stute (1984), Devroye and Györfi (1985), Bhattacharya and Mack (1987), Zhao (1987), and Devroye, Györfi, Krzyżak, and Lugosi (1994).

Recall the definition of the $k$-nearest neighbor rule: first reorder the data

$$
\left(X_{(1)}(x), Y_{(1)}(x)\right), \ldots,\left(X_{(n)}(x), Y_{(n)}(x)\right)
$$

according to increasing Euclidean distances of the $X_{j}$ 's to $x$. In other words, $X_{(i)}(x)$ is the $i$-th nearest neighbor of $x$ among the points $X_{1}, \ldots, X_{n}$. If distance ties occur,
a tie-breaking strategy must be defined. If $\mu$ is absolutely continuous with respect to the Lebesgue measure, that is, it has a density, then no ties occur with probability one, so formally we break ties by comparing indices. However, for general $\mu$, the problem of distance ties turns out to be important, and its solution is messy. The issue of tie breaking becomes important when one is concerned with convergence of $L_{n}$ with probability one. For weak universal consistency, it suffices to break ties by comparing indices.

The $k$-NN classification rule is defined as

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{k} I_{\left\{Y_{i(i)}(x)=1\right\}} \leq \sum_{i=1}^{k} I_{\left\{Y_{(i)}(x)=0\right\}} \\ 1 & \text { otherwise }\end{cases}
$$

In other words, $g_{n}(x)$ is a majority vote among the labels of the $k$ nearest neighbors of $x$.

### 11.1 Strong Consistency

In this section we prove Theorem 11.1. We assume the existence of a density for $\mu$, so that we can avoid messy technicalities necessary to handle distance ties. We discuss this issue briefly in the next section.

The following result implies strong consistency whenever $X$ has an absolutely continuous distribution. The result was proved by Devroye and Györfi (1985), and Zhao (1987). The proof presented here basically appears in Devroye, Györfi, Krzyżak, and Lugosi (1994), where strong universal consistency is proved under an appropriate tie-breaking strategy (see discussion later). Some of the main ideas appeared in the proof of the strong universal consistency of the regular histogram rule (Theorem 9.4).

Theorem 11.1. (Devroye and Györfi (1985), Zhao (1987)). Assume that $\mu$ has $a$ density. If $k \rightarrow \infty$ and $k / n \rightarrow 0$ then for every $\epsilon>0$ there is an $n_{0}$ such that for $n>n_{0}$

$$
\mathbf{P}\left\{L_{n}-L^{*}>\epsilon\right\} \leq 2 e^{-n \epsilon^{2} /\left(72 \gamma_{d}^{2}\right)}
$$

where $\gamma_{d}$ is the minimal number of cones centered at the origin of angle $\pi / 6$ that cover $\mathcal{R}^{d}$. (For the definition of a cone, see Chapter 5.) Thus, the $k$-NN rule is strongly consistent.

REmARK. At first glance the upper bound in the theorem does not seem to depend on $k$. It is $n_{0}$ that depends on the sequence of $k$ 's. What we really prove is the following: for every $\epsilon>0$ there exists a $\beta_{0} \in(0,1)$ such that for any $\beta<\beta_{0}$ there is an $n_{0}$ such that if $n>n_{0}, k>1 / \beta$, and $k / n<\beta$, the exponential inequality holds.

For the proof we need a generalization of Lemma 5.3. The role of this covering lemma is analogous to that of Lemma 10.1 in the proof of consistency of kernel rules.

Lemma 11.1. (Devroye and Györfi (1985)). Let

$$
B_{a}\left(x^{\prime}\right)=\left\{x: \mu\left(S_{x,\left\|x-x^{\prime}\right\|}\right) \leq a\right\} .
$$

Then for all $x^{\prime} \in \mathcal{R}^{d}$

$$
\mu\left(B_{a}\left(x^{\prime}\right)\right) \leq \gamma_{d} a .
$$

Proof. For $x \in \mathcal{R}^{d}$ let $C(x, s) \subset \mathcal{R}^{d}$ be a cone of angle $\pi / 6$ centered at $x$. The cone consists of all $y$ with the property that either $y=x$ or angle $(y-x, s) \leq \pi / 6$, where $s$ is a fixed direction. If $y, y^{\prime} \in C(x, s)$, and $\|x-y\|<\left\|x-y^{\prime}\right\|$, then $\left\|y-y^{\prime}\right\|<\left\|x-y^{\prime}\right\|$. This follows from a simple geometric argument in the vector space spanned by $x, y$ and $y^{\prime}$ (see the proof of Lemma 5.3).

Now, let $C_{1}, \ldots, C_{\gamma_{d}}$ be a collection of cones centered at $x$ with different central direction covering $\mathcal{R}^{d}$. Then

$$
\mu\left(B_{a}\left(x^{\prime}\right)\right) \leq \sum_{i=1}^{\gamma_{d}} \mu\left(C_{i} \cap B_{a}\left(x^{\prime}\right)\right) .
$$

Let $x^{*} \in C_{i} \cap B_{a}\left(x^{\prime}\right)$. Then by the property of the cones mentioned above we have

$$
\mu\left(C_{i} \cap S_{x^{\prime},\left\|x^{\prime}-x^{*}\right\|} \cap B_{a}\left(x^{\prime}\right)\right) \leq \mu\left(S_{x^{*},\left\|x^{\prime}-x^{*}\right\|}\right) \leq a,
$$

where we use the fact that $x^{*} \in B_{a}\left(x^{\prime}\right)$. Since $x^{*}$ is arbitrary,

$$
\mu\left(C_{i} \cap B_{a}\left(x^{\prime}\right)\right) \leq a,
$$

which completes the proof of the lemma.
An immediate consequence of the lemma is that the number of points among $X_{1}, \ldots, X_{n}$ such that $X$ is one of their $k$ nearest neighbors is not more than a constant times $k$.
Corollary 11.1.

Proof. Apply Lemma 11.1 with $a=k / n$ and let $\mu$ be the empirical measure $\mu_{n}$ of $X_{1}, \ldots, X_{n}$, that is, for each Borel set $A \subseteq \mathcal{R}^{d}, \mu_{n}(A)=(1 / n) \sum_{i=1}^{n} I_{\left\{X_{i} \in A\right\}}$.

Proof of Theorem 11.1. Since the decision rule $g_{n}$ may be rewritten as

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

where $\eta_{n}$ is the corresponding regression function estimate

$$
\eta_{n}(x)=\frac{1}{k} \sum_{i=1}^{k} Y_{(i)}(x),
$$

the statement follows from Theorem 2.2 if we show that for sufficiently large $n$

$$
\mathbf{P}\left\{\int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)>\frac{\epsilon}{2}\right\} \leq 2 e^{-n \epsilon^{2} /\left(72 \gamma_{d}^{2}\right)}
$$

Define $\rho_{n}(x)$ as the solution of the equation

$$
\frac{k}{n}=\mu\left(S_{x, \rho_{n}(x)}\right)
$$

Note that the absolute continuity of $\mu$ implies that the solution always exists. (This is the only point in the proof where we use this assumption.) Also define

$$
\eta_{n}^{*}(x)=\frac{1}{k} \sum_{j=1}^{n} Y_{j} I_{\left\{\left\|X_{j}-x\right\|<\rho_{n}(x)\right\}}
$$

The basis of the proof is the following decomposition:

$$
\left|\eta(x)-\eta_{n}(x)\right| \leq\left|\eta_{n}(x)-\eta_{n}^{*}(x)\right|+\left|\eta_{n}^{*}(x)-\eta(x)\right| .
$$

For the first term on the right-hand side, observe that denoting $R_{n}(x)=\| X_{(k)}(x)-$ $x \|$,

$$
\begin{aligned}
\left|\eta_{n}^{*}(x)-\eta_{n}(x)\right| & =\frac{1}{k}\left|\sum_{j=1}^{n} Y_{j} I_{\left\{X_{j} \in S_{x, \rho_{n}(x)}\right\}}-\sum_{j=1}^{n} Y_{j} I_{\left\{X_{j} \in S_{x, R_{n}(x)}\right\}}\right| \\
& \leq \frac{1}{k} \sum_{j=1}^{n}\left|I_{\left\{X_{j} \in S_{\left.x, \rho_{n}(x)\right\}}\right.}-I_{\left\{X_{j} \in S_{x, R_{n}(x)}\right\}}\right| \\
& =\left|\frac{1}{k} \sum_{j=1}^{n} I_{\left\{X_{j} \in S_{\left.x, \rho_{n}(x)\right\}}\right.}-1\right|=\left|\widehat{\eta}_{n}^{*}(x)-\widehat{\eta}(x)\right|
\end{aligned}
$$

where $\widehat{\eta}_{n}^{*}$ is defined as $\eta_{n}^{*}$ with $Y$ replaced by the constant random variable $\widehat{Y}=1$, and $\widehat{\eta} \equiv 1$ is the corresponding regression function. Thus,

$$
\begin{equation*}
\left|\eta(x)-\eta_{n}(x)\right| \leq\left|\widehat{\eta}_{n}^{*}(x)-\widehat{\eta}(x)\right|+\left|\eta_{n}^{*}(x)-\eta(x)\right| . \tag{11.1}
\end{equation*}
$$

First we show that the expected values of the integrals of both terms on the righthand side converge to zero. Then we use McDiarmid's inequality to prove that both terms are very close to their expected values with large probability.

For the expected value of the first term on the right-hand side of (11.1), using the Cauchy-Schwarz inequality, we have

$$
\begin{aligned}
\mathbf{E} \int\left|\eta_{n}^{*}(x)-\eta_{n}(x)\right| \mu(d x) & \leq \mathbf{E} \int\left|\widehat{\eta}_{n}^{*}(x)-\widehat{\eta}(x)\right| \mu(d x) \\
& \leq \int \sqrt{\mathbf{E}\left\{\left|\widehat{\eta}_{n}^{*}(x)-\widehat{\eta}(x)\right|^{2}\right\}} \mu(d x)
\end{aligned}
$$

$$
\begin{aligned}
& =\int \sqrt{\frac{1}{k^{2}} n \operatorname{Var}\left\{I_{\left\{X \in S_{x, \rho_{n}(x)}\right\}}\right\}} \mu(d x) \\
& \leq \int \sqrt{\frac{1}{k^{2}} n \mu\left(S_{x, \rho_{n}(x)}\right)} \mu(d x) \\
& =\int \sqrt{\frac{n}{k^{2}} \frac{k}{n}} \mu(d x) \\
& =\frac{1}{\sqrt{k}}
\end{aligned}
$$

which converges to zero.
For the expected value of the second term on the right-hand side of (11.1), note that in the proof of Theorems 6.3 and 6.4 we already showed that

$$
\lim _{n \rightarrow \infty} \mathbf{E} \int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)=0
$$

Therefore,

$$
\begin{aligned}
& \mathbf{E} \int\left|\eta_{n}^{*}(x)-\eta(x)\right| \mu(d x) \\
& \quad \leq \mathbf{E} \int\left|\eta_{n}^{*}(x)-\eta_{n}(x)\right| \mu(d x)+\mathbf{E} \int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x) \rightarrow 0
\end{aligned}
$$

Assume now that $n$ is so large that

$$
\mathbf{E} \int\left|\widehat{\eta}_{n}^{*}(x)-\widehat{\eta}(x)\right| \mu(d x)+\mathbf{E} \int\left|\eta_{n}^{*}(x)-\eta(x)\right| \mu(d x)<\frac{\epsilon}{6}
$$

Then, by (11.1), we have

$$
\begin{align*}
& \mathbf{P}\left\{\int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)>\frac{\epsilon}{2}\right\}  \tag{11.2}\\
& \leq \mathbf{P}\left\{\int\left|\eta_{n}^{*}(x)-\eta(x)\right| \mu(d x)-\mathbf{E} \int\left|\eta_{n}^{*}(x)-\eta(x)\right| \mu(d x)>\frac{\epsilon}{6}\right\} \\
& \quad+\mathbf{P}\left\{\int\left|\widehat{\eta}_{n}^{*}(x)-\widehat{\eta}(x)\right| \mu(d x)-\mathbf{E} \int\left|\widehat{\eta}_{n}^{*}(x)-\widehat{\eta}(x)\right| \mu(d x)>\frac{\epsilon}{6}\right\}
\end{align*}
$$

Next we get an exponential bound for the first probability on the right-hand side of (11.2) by McDiarmid's inequality (Theorem 9.2). Fix an arbitrary realization of the data $D_{n}=\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)$, and replace $\left(x_{i}, y_{i}\right)$ by $\left(\widehat{x_{i}}, \widehat{y_{i}}\right)$, changing the value of $\eta_{n}^{*}(x)$ to $\eta_{n i}^{*}(x)$. Then
$\left|\int\right| \eta_{n}^{*}(x)-\eta(x)\left|\mu(d x)-\int\right| \eta_{n i}^{*}(x)-\eta(x)|\mu(d x)| \leq \int\left|\eta_{n}^{*}(x)-\eta_{n i}^{*}(x)\right| \mu(d x)$.

But $\left|\eta_{n}^{*}(x)-\eta_{n i}^{*}(x)\right|$ is bounded by $2 / k$ and can differ from zero only if $\| x-$ $x_{i} \|<\rho_{n}(x)$ or $\left\|x-\widehat{x}_{i}\right\|<\rho_{n}(x)$. Observe that $\left\|x-x_{i}\right\|<\rho_{n}(x)$ if and only if $\mu\left(S_{x,\left\|x-x_{i}\right\|}\right)<k / n$. But the measure of such $x$ 's is bounded by $\gamma_{d} k / n$ by Lemma 11.1. Therefore

$$
\sup _{x_{1}, y_{1}, \ldots, x_{n}, y_{n}, \widehat{x}_{i}, \hat{y}_{i}} \int\left|\eta_{n}^{*}(x)-\eta_{n i}^{*}(x)\right| \mu(d x) \leq \frac{2}{k} \frac{\gamma_{d} k}{n}=\frac{2 \gamma_{d}}{n}
$$

and by Theorem 9.2

$$
\mathbf{P}\left\{\left|\int\right| \eta(x)-\eta_{n}^{*}(x)\left|\mu(d x)-\mathbf{E} \int\right| \eta(x)-\eta_{n}^{*}(x)|\mu(d x)|>\frac{\epsilon}{6}\right\} \leq 2 e^{-n \epsilon^{2} /\left(72 \gamma_{d}^{2}\right)}
$$

Finally, we need a bound for the second term on the right-hand side of (11.2). This probability may be bounded by McDiarmid's inequality exactly the same way as for the first term, obtaining

$$
\mathbf{P}\left\{\left|\int\right| \widehat{\eta}_{n}^{*}(x)-\widehat{\eta}(x)\left|\mu(d x)-\mathbf{E} \int\right| \widehat{\eta}_{n}^{*}(x)-\widehat{\eta}(x)|\mu(d x)|>\frac{\epsilon}{6}\right\} \leq 2 e^{-n \epsilon^{2} /\left(72 \gamma_{d}^{2}\right)}
$$

and the proof is completed.

REmark. The conditions $k \rightarrow \infty$ and $k / n \rightarrow 0$ are optimal in the sense that they are also necessary for consistency for some distributions with a density. However, for some distributions they are not necessary for consistency, and in fact, keeping $k=1$ for all $n$ may be a better choice. This latter property, dealt with in Problem 11.1, shows that the 1 -nearest neighbor rule is admissible.

### 11.2 Breaking Distance Ties

Theorem 11.1 provides strong consistency under the assumption that $X$ has a density. This assumption was needed to avoid problems caused by equal distances. Turning to the general case, we see that if $\mu$ does not have a density then distance ties can occur with nonzero probability, so we have to deal with the problem of breaking them. To see that the density assumption cannot be relaxed to the condition that $\mu$ is merely nonatomic without facing frequent distance ties, consider the following distribution on $\mathcal{R}^{d} \times \mathcal{R}^{d^{\prime}}$ with $d, d^{\prime} \geq 2$ :

$$
\mu=\frac{1}{2}\left(\tau_{d} \times \sigma_{d^{\prime}}\right)+\frac{1}{2}\left(\sigma_{d} \times \tau_{d^{\prime}}\right),
$$

where $\tau_{d}$ denotes the uniform distribution on the surface of the unit sphere of $\mathcal{R}^{d}$ and $\sigma_{d}$ denotes the unit point mass at the origin of $\mathcal{R}^{d}$. Observe that if $X$ has distribution $\tau_{d} \times \sigma_{d^{\prime}}$ and $X^{\prime}$ has distribution $\sigma_{d} \times \tau_{d^{\prime}}$, then $\left\|X-X^{\prime}\right\|=\sqrt{2}$.

Hence, if $X_{1}, X_{2}, X_{3}, X_{4}$ are independent with distribution $\mu$, then $\mathbf{P}\left\{\left\|X_{1}-X_{2}\right\|=\right.$ $\left.\left\|X_{3}-X_{3}\right\|\right\}=1 / 4$.

Next we list some methods of breaking distance ties.

- Tie-breaking by indices: If $X_{i}$ and $X_{j}$ are equidistant from $x$, then $X_{i}$ is declared closer if $i<j$. This method has some undesirable properties. For example, if $X$ is monoatomic, with $\eta<1 / 2$, then $X_{1}$ is the nearest neighbor of all $X_{j}$ 's, $j>1$, but $X_{j}$ is only the $j-1$-st nearest neighbor of $X_{1}$. The influence of $X_{1}$ in such a situation is too large, making the estimate very unstable and thus undesirable. In fact, in this monoatomic case, if $L^{*}+\epsilon<1 / 2$,

$$
\mathbf{P}\left\{L_{n}-L^{*}>\epsilon\right\} \geq e^{-c k}
$$

for some $c>0$ (see Problem 11.2). Thus, we cannot expect a distributionfree version of Theorem 11.1 with this tie-breaking method.

- Stone's tie-breaking: Stone (1977) introduced a version of the nearest neighbor rule, where the labels of the points having the same distance from $x$ as the $k$-th nearest neighbor are averaged. If we denote the distance of the $k$-th nearest neighbor to $x$ by $R_{n}(x)$, then Stone's rule is the following:

$$
g_{n}(x)=\left\{\begin{aligned}
0 & \text { if } \sum_{i:\left\|x-x_{i}\right\|<R_{n}(x)} I_{\left\{Y_{i}=0\right\}}+\frac{k-\left|\left\{i:\left\|x-x_{i}\right\|<R_{n}(x)\right\}\right|}{\left|\left\{i:\left\|x-x_{i}\right\|=R_{n}(x)\right\}\right|} \\
& \times \sum_{i:\left\|x-x_{i}\right\|=R_{n}(x)} I_{\left\{Y_{i}=0\right\}} \geq \frac{k}{2} \\
1 & \text { otherwise. }
\end{aligned}\right.
$$

This is not a $k$-nearest neighbor rule in a strict sense, since this estimate, in general, uses more than $k$ neighbors. Stone (1977) proved weak universal consistency of this rule.

- Adding a random component: To circumvent the aforementioned difficulties, we may artificially increase the dimension of the feature vector by one. Define the the $d+1$-dimensional random vectors

$$
X^{\prime}=(X, U), X_{1}^{\prime}=\left(X_{1}, U_{1}\right), \ldots, X_{n}^{\prime}=\left(X_{n}, U_{n}\right)
$$

where the randomizing variables $U, U_{1}, \ldots, U_{n}$ are real-valued i.i.d. random variables independent of $X, Y$, and $D_{n}$, and their common distribution has a density. Clearly, because of the independence of $U$, the Bayes error corresponding to the pair $\left(X^{\prime}, Y\right)$ is the same as that of $(X, Y)$. The algorithm performs the $k$-nearest neighbor rule on the modified data set

$$
D_{n}^{\prime}=\left(\left(X_{1}^{\prime}, Y_{1}\right), \ldots,\left(X_{n}^{\prime}, Y_{n}\right)\right)
$$

It finds the $k$ nearest neighbors of $X^{\prime}$, and uses a majority vote among these labels to guess $Y$. Since $U$ has a density and is independent of $X$, distance
ties occur with zero probability. Strong universal consistency of this rule can be seen by observing that the proof of Theorem 11.1 used the existence of the density in the definition of $\rho_{n}(x)$ only. With our randomization, $\rho_{n}(x)$ is well-defined, and the same proof yields strong universal consistency. Interestingly, this rule is consistent whenever $U$ has a density and is independent of $(X, Y)$. If, for example, the magnitude of $U$ is much larger than that of $\|X\|$, then the rule defined this way will significantly differ from the $k$ nearest neighbor rule, though it still preserves universal consistency. One should expect however a dramatic decrease in the performance. Of course, if $U$ is very small, then the rule remains intuitively appealing.

- Tie-breaking by randomization: There is another, perhaps more natural way of breaking ties via randomization. We assume that $(X, U)$ is a random vector independent of the data, where $U$ is independent of $X$ and uniformly distributed on $[0,1]$. We also artificially enlarge the data by introducing $U_{1}, U_{2}, \ldots, U_{n}$, where the $U_{i}$ 's are i.i.d. uniform [0,1] as well. Thus, each ( $X_{i}, U_{i}$ ) is distributed as $(X, U)$. Let

$$
\left(X_{(1)}(x, u), Y_{(1)}(x, u)\right), \ldots,\left(X_{(n)}(x, u), Y_{(n)}(x, u)\right)
$$

be a reordering of the data according to increasing values of $\left\|x-X_{i}\right\|$. In case of distance ties, we declare $\left(X_{i}, U_{i}\right)$ closer to $(x, u)$ than $\left(X_{j}, U_{j}\right)$ provided that

$$
\left|U_{i}-u\right| \leq\left|U_{j}-u\right|
$$

Define the $k$-NN classification rule as

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{k} I_{\left\{Y_{(i)}(x, u)=1\right\}} \leq \sum_{i=1}^{k} I_{\left\{Y_{(i)}(x, u)=0\right\}}  \tag{11.3}\\ 1 & \text { otherwise },\end{cases}
$$

and denote the error probability of $g_{n}$ by

$$
L_{n}=\mathbf{P}\left\{g_{n}(X, U) \neq Y \mid X_{1}, U_{1}, Y_{1}, \ldots, X_{n}, U_{n}, Y_{n}\right\}
$$

Devroye, Györfi, Krzyżak, and Lugosi (1994) proved that $L_{n} \rightarrow L^{*}$ with probability one for all distributions if $k \rightarrow \infty$ and $k / n \rightarrow 0$. The basic argument in (1994) is the same as that of Theorem 11.1, except that the covering lemma (Lemma 11.1) has to be appropriately modified.

It should be stressed again that if $\mu$ has a density, or just has an absolutely continuous component, then tie-breaking is needed with zero probability, and becomes therefore irrelevant.

### 11.3 Recursive Methods

To find the nearest neighbor of a point $x$ among $X_{1}, \ldots, X_{n}$, we may preprocess the data in $O(n \log n)$ time, such that each query may be answered in $O(\log n)$
worst-case time-see, for example, Preparata and Shamos (1985). Other recent developments in computational geometry have made the nearest neighbor rules computationally feasible even when $n$ is formidable. Without preprocessing however, one must resort to slow methods. If we need to find a decision at $x$ and want to process the data file once when doing so, a simple rule was proposed by Devroye and Wise (1980). It is a fully recursive rule that may be updated as more observations become available.

Split the data sequence $D_{n}$ into disjoint blocks of length $l_{1}, \ldots, l_{N}$, where $l_{1}, \ldots, l_{N}$ are positive integers satisfying $\sum_{i=1}^{N} l_{i}=n$. In each block find the nearest neighbor of $x$, and denote the nearest neighbor of $x$ from the $i$-th block by $X_{i}^{*}(x)$. Let $Y_{i}^{*}(x)$ be the corresponding label. Ties are broken by comparing indices. The classification rule is defined as a majority vote among the nearest neighbors from each block:

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{N} I_{\left\{Y_{i}^{*}(x)=1\right\}} \leq \sum_{i=1}^{N} I_{\left\{Y_{i}^{*}(x)=0\right\}} \\ 1 & \text { otherwise }\end{cases}
$$

Note that we have only defined the rule $g_{n}$ for $n$ satisfying $\sum_{i=1}^{N} l_{i}=n$ for some $N$. A possible extension for all $n$ 's is given by $g_{n}(x)=g_{m}(x)$, where $m$ is the largest integer not exceeding $n$ that can be written as $\sum_{i=1}^{N} l_{i}$ for some $N$. The rule is weakly universally consistent if

$$
\lim _{N \rightarrow \infty} l_{N}=\infty
$$

(Devroye and Wise (1980), see Problem 11.3).

### 11.4 Scale-Invariant Rules

A scale-invariant rule is a rule that is invariant under rescalings of the components. It is motivated by the lack of a universal yardstick when components of a vector represent physically different quantities, such as temperature, blood pressure, alcohol, and the number of lost teeth. More formally, let $x^{(1)}, \ldots, x^{(d)}$ be the $d$ components of a vector $x$. If $\psi_{1}, \ldots, \psi_{d}$ are strictly monotone mappings: $\mathcal{R} \rightarrow \mathcal{R}$, and if we define

$$
\psi(x)=\left(\psi_{1}\left(x^{(1)}\right), \ldots, \psi_{d}\left(x^{(d)}\right)\right)
$$

then $g_{n}$ is scale-invariant if

$$
g_{n}\left(x, D_{n}\right)=g_{n}\left(\psi(x), D_{n}^{\prime}\right),
$$

where $D_{n}^{\prime}=\left(\left(\psi\left(X_{1}\right), Y_{1}\right), \ldots,\left(\psi\left(X_{n}\right), Y_{n}\right)\right)$. In other words, if all the $X_{i}$ 's and $x$ are transformed in the same manner, the decision does not change.

Some rules based on statistically equivalent blocks (discussed in Chapters 21 and 22) are scale-invariant, while the $k$-nearest neighbor rule clearly is not. Here we describe a scale-invariant modification of the $k$-nearest neighbor rule, suggested by Olshen (1977) and Devroye (1978).

The scale-invariant $k$-nearest neighbor rule is based upon empirical distances that are defined in terms of the order statistics along the $d$ coordinate axes. First order the points $x, X_{1}, \ldots, X_{n}$ according to increasing values of their first components $x^{(1)}, X_{1}^{(1)}, \ldots, X_{n}^{(1)}$, breaking ties via randomization. Denote the rank of $X_{i}^{(1)}$ by $r_{i}^{(1)}$, and the rank of $x^{(1)}$ by $r^{(1)}$. Repeating the same procedure for the other coordinates, we obtain the ranks

$$
r_{i}^{(j)}, r^{(j)}, \quad j=1, \ldots, d, i=1, \ldots, n
$$

Define the empirical distance between $x$ and $X_{i}$ by

$$
\rho\left(x, X_{i}\right)=\max _{1 \leq j \leq d}\left|r_{i}^{(j)}-r^{(j)}\right|
$$

A $k$-NN rule can be defined based on these distances, by a majority vote among the $Y_{i}$ 's with the corresponding $X_{i}$ 's whose empirical distance from $x$ are among the $k$ smallest. Since these distances are integer-valued, ties frequently occur. These ties should be broken by randomization. Devroye (1978) proved that this rule (with randomized tie-breaking) is weakly universally consistent when $k \rightarrow \infty$ and $k / n \rightarrow 0$ (see Problem 11.5). For another consistent scale-invariant nearest neighbor rule we refer to Problem 11.6.


Figure 11.1. Scale-invariant distances of 15 points from a fixed point are shown here.

### 11.5 Weighted Nearest Neighbor Rules

In the $k-\mathrm{NN}$ rule, each of the $k$ nearest neighbors of a point $x$ plays an equally important role in the decision. However, intuitively speaking, nearer neighbors should
provide more information than more distant ones. Royall (1966) first suggested using rules in which the labels $Y_{i}$ are given unequal voting powers in the decision according to the distances of the $X_{i}$ 's from $x$ : the $i$-th nearest neighbor receives weight $w_{n i}$, where usually $w_{n 1} \geq w_{n 2} \geq \ldots \geq w_{n n} \geq 0$ and $\sum_{i=1}^{n} w_{n i}=1$. The rule is defined as

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} w_{n i} I_{\left\{Y_{(i)}(x)=1\right\}} \leq \sum_{i=1}^{n} w_{n i} I_{\left\{Y_{(i)}(x)=0\right\}} \\ 1 & \text { otherwise } .\end{cases}
$$

We get the ordinary $k$-nearest neighbor rule back by the choice

$$
w_{n i}= \begin{cases}1 / k & \text { if } i \leq k \\ 0 & \text { otherwise. }\end{cases}
$$

The following conditions for consistency were established by Stone (1977):

$$
\lim _{n \rightarrow \infty} \max _{1 \leq i \leq n} w_{n i}=0,
$$

and

$$
\lim _{n \rightarrow \infty} \sum_{k \leq i \leq n} w_{n i}=0
$$

for some $k$ with $k / n \rightarrow 0$ (see Problem 11.7). Weighted versions of the recursive and scale-invariant methods described above can also be defined similarly.

### 11.6 Rotation-Invariant Rules

Assume that an affine transformation $T$ is applied to $x$ and $X_{1}, \ldots, X_{n}$ (i.e., any number of combinations of rotations, translations, and linear rescalings), and that for any such linear transformation $T$,

$$
g_{n}\left(x, D_{n}\right)=g_{n}\left(T(x), D_{n}^{\prime}\right),
$$

where $D_{n}^{\prime}=\left(\left(T\left(X_{1}\right), Y_{1}\right), \ldots,\left(T\left(X_{n}\right), Y_{n}\right)\right)$. Then we call $g_{n}$ rotation-invariant. Rotation-invariance is indeed a very strong property. In $\mathcal{R}^{d}$, in the context of $k$-NN estimates, it suffices to be able to define a rotation-invariant distance measure. These are necessarily data-dependent. An example of this goes as follows. Any collection of $d$ points in general position defines a polyhedron in a hyperplane of $\mathcal{R}^{d}$. For points ( $X_{i_{1}}, \ldots, X_{i_{d}}$ ), we denote this polyhedron by $P\left(i_{1}, \ldots, i_{d}\right)$. Then we define the distance

$$
\rho\left(X_{i}, x\right)=\sum_{\substack{i_{1}, \ldots, i_{d}, i \notin\left\{i_{1}, \ldots, i_{d}\right\}}} I_{\left\{\text {segment }\left(X_{i}, x\right) \text { intersects } P\left(i_{1}, \ldots, i_{d}\right)\right\}} .
$$

Near points have few intersections. Using $\rho(\cdot, \cdot)$ in a $k$-nN rule with $k \rightarrow \infty$ and $k / n \rightarrow 0$, we expect weak universal consistency under an appropriate scheme of tie-breaking. The answer to this is left as an open problem for the scholars.


Figure 11.2. Rotation-invariant distances from $x$.

### 11.7 Relabeling Rules

The 1-NN rule appeals to the masses who crave simplicity and attracts the programmers who want to write short understandable code. Nearest neighbors may be found efficiently if the data are preprocessed (see Chapter 5 for references). Can we make the $1-\mathrm{NN}$ rule universally consistent as well? In this section we introduce a tool called relabeling, which works as follows. Assume that we have a classification rule $\left\{g_{n}\left(x, D_{n}\right), x \in \mathcal{R}^{d}, n \geq 1\right\}$, where $D_{n}$ is the data $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$. This rule will be called the ancestral rule. Define the labels

$$
Z_{i}=g_{n}\left(X_{i}, D_{n}\right), \quad 1 \leq i \leq n .
$$

These are the decisions for the $X_{i}$ 's themselves obtained by mere resubstitution. In the relabeling method, we apply the $1-\mathrm{NN}$ rule to the new data ( $X_{1}, Z_{1}$ ), $\ldots$, ( $X_{n}, Z_{n}$ ). If all goes well, when the ancestral rule $g_{n}$ is universally consistent, so should the relabeling rule. We will show this by example, starting from a consistent $k$-NN rule as ancestral rule (with $k \rightarrow \infty, k / n \rightarrow 0$ ).

Unfortunately, relabeling rules do not always inherit consistency from their ancestral rules, so that a more general theorem is more difficult to obtain, unless one adds in a lot of regularity conditions-this does not seem to be the right time for that sort of effort. To see that universal consistency of the ancestral rule does not imply consistency of the relabeling rule, consider the following rule $h_{n}$ :

$$
h_{n}(x)= \begin{cases}1-Y_{i} & \text { if } x=X_{i} \text { and } x \neq X_{j}, \text { all } j \neq i \\ g_{n}\left(x, D_{n}\right) & \text { otherwise, }\end{cases}
$$

where $g_{n}$ is a weakly universally consistent rule. It is easy to show (see Problem 11.15) that $h_{n}$ is universally consistent as well. Changing a rule on a set of measure zero indeed does not affect $L_{n}$. Also, if $x=X_{i}$ is at an atom of the distribution of $X$, we only change $g_{n}$ to $1-Y_{i}$ if $X_{i}$ is the sole occurrence of that atom in the data. This has asymptotically no impact on $L_{n}$. However, if $h_{n}$ is used as an ancestral rule, and $X$ is nonatomic, then $h_{n}\left(X_{i}, D_{n}\right)=1-Y_{i}$ for all $i$, and therefore, the relabeling rule is a $1-\mathrm{NN}$ rule based on the data $\left(X_{1}, 1-Y_{1}\right), \ldots,\left(X_{n}, 1-Y_{n}\right)$. If $L^{*}=0$ for the distribution of $(X, Y)$, then the relabeling rule has probability of error converging to one!

For most nonpathological ancestral rules, relabeling does indeed preserve universal consistency. We offer a prototype proof for the $k$-NN rule.

Theorem 11.2. Let $g_{n}$ be the $k$-NN rule in which tie-breaking is done by randomization as in (11.3). Assume that $k \rightarrow \infty$ and $k / n \rightarrow 0$ (so that $g_{n}$ is weakly universally consistent). Then the relabeling rule based upon $g_{n}$ is weakly universally consistent as well.

Proof. We verify the conditions of Stone's weak convergence theorem (see Theorem 6.3). To keep things simple, we assume that the distribution of $X$ has a density so that distance ties happen with probability zero.

In our case, the weight $W_{n i}(X)$ of Theorem 6.3 equals $1 / k$ iff $X_{i}$ is among the $k$ nearest neighbors of $X_{(1)}(X)$, where $X_{(1)}(X)$ is the nearest neighbor of $X$. It is zero otherwise.

Condition 6.3 (iii) is trivially satisfied since $k \rightarrow \infty$. For condition 6.3 (ii), we note that if $X_{(i)}(x)$ denotes the $i$-th nearest neighbor of $X$ among $X_{1}, \ldots, X_{n}$, then

$$
\left\|X_{(1)}(x)-X_{(k)}\left(X_{(1)}(x)\right)\right\| \leq 2\left\|x-X_{(k)}\right\| .
$$

Just note that $S_{X_{(1)}(x), 2\left\|x-X_{(k)}\right\|} \supseteq S_{x,\left\|x-X_{(k)}\right\|}$ and that the latter sphere contains $k$ data points. But we already know from the proof of weak consistency of the $k$-NN rule that if $k / n \rightarrow 0$, then for all $\epsilon>0$,

$$
\mathbf{P}\left\{\left\|X_{(k)}(X)-X\right\|>\epsilon\right\} \rightarrow 0
$$

Finally, we consider condition 6.3 (i). Here we have, arguing partially as in Stone (1977),

$$
\begin{aligned}
& \mathbf{E}\left\{\sum_{i=1}^{n} \frac{1}{k} I_{\left\{X_{i} \text { is among the } k \text { NN's of } X_{(1)}(X)\right\}} f\left(X_{i}\right)\right\} \\
& \quad=\mathbf{E}\left\{\frac{1}{k} \sum_{i=1}^{n} I_{\left\{X \text { is among the } k \text { NN's of } X_{(1)}\left(X_{i}\right) \text { in }\left\{X_{1}, \ldots, X_{n}, X\right\}-\left\{X_{i}\right\}\right\}} f(X)\right\}
\end{aligned}
$$

(reverse the roles of $X_{i}$ and $X$ )
$=\mathbf{E}\left\{\frac{1}{k} \sum_{i=1}^{n} \sum_{j=1}^{n} I_{\left\{X_{j}\right.}\right.$ is the NN of $X_{i}$ in $\left.\left\{X_{1}, \ldots, X_{n}, X\right\}-\left\{X_{i}\right\}\right\}$
$\left.\times I_{\left\{X \text { is among the } k \text { NN's of } X_{j} \text { in }\left\{X_{1}, \ldots, X_{n}, X\right\}-\left\{X_{i}\right\}\right\}} f(X)\right\}$.
However, by Lemma 11.1,

$$
\left.\sum_{i=1}^{n} I_{\left\{X_{j}\right.} \text { is the NN of } X_{i} \text { in }\left\{X_{1}, \ldots, X_{n}, X\right\}-\left\{X_{i}\right\}\right\} \leq \gamma_{d} .
$$

Also,

Therefore, by a double application of Lemma 11.1,

$$
\mathbf{E}\left\{\sum_{i=1}^{n} \frac{1}{k} I_{\left\{X_{i} \text { is among the } k \text { NN's of } X_{(1)}(X)\right\}} f\left(X_{i}\right)\right\} \leq \gamma_{d}^{2} \mathbf{E} f(X)
$$

and condition 6.3 (i) is verified.

## Problems and Exercises

Problem 11.1. Show that the conditions $k \rightarrow \infty$ and $k / n \rightarrow 0$ are necessary for universal consistency of the $k$-nearest neighbor rule. That is, exhibit a distribution such that if $k$ remains bounded, $\lim \inf _{n \rightarrow \infty} \mathbf{E} L_{n}>L^{*}$. Exhibit a second distribution such that if $k / n \geq \epsilon>0$ for all $n$, and some $\epsilon$, then $\liminf _{n \rightarrow \infty} \mathbf{E} L_{n}>L^{*}$.

Problem 11.2. Let $X$ be monoatomic, with $\eta<1 / 2$. Show that for $\epsilon<1 / 2-\eta$,

$$
\mathbf{P}\left\{L_{n}-L^{*}>\epsilon\right\} \geq e^{-c k}
$$

for some $c>0$, where $L_{n}$ is the error probability of the $k$-NN rule with tie-breaking by indices.

Problem 11.3. Prove that the recursive nearest neighbor rule is universally consistent provided that $\lim _{N \rightarrow \infty} l_{N}=\infty$ (Devroye and Wise (1980)). Hint: Check the conditions of Theorem 6.3.

Problem 11.4. Prove that the nearest neighbor rule defined by any $L_{p}$-distance measure $0<p \leq \infty$ is universally consistent under the usual conditions on $k$. The $L_{p}$-distance between $x, y \in \mathcal{R}^{d}$ is defined by $\left(\sum_{i=1}^{d}\left|x^{(i)}-y^{(i)}\right|^{p}\right)^{1 / p}$ for $0<p<\infty$, and by $\sup _{i} \mid x^{(i)}-$ $y^{(i)} \mid$ for $p=\infty$, where $x=\left(x^{(1)}, \ldots, x^{(d)}\right)$. Hint: Check the conditions of Theorem 6.3.

Problem 11.5. Let $\sigma\left(x, z, X_{i}, Z_{i}\right)=\rho\left(x, X_{i}\right)+\left|Z_{i}-z\right|$ be a generalized distance between $(x, z)$ and $\left(X_{i}, Z_{i}\right)$, where $x, X_{1}, \ldots, X_{n}$ are as in the description of the scale-invariant $k$-NN rule, $\rho\left(x, X_{i}\right)$ is the empirical distance defined there, and $z, Z_{i} \in[0,1]$ are real numbers added to break ties at random. The sequence $Z_{1}, \ldots, Z_{n}$ is i.i.d. uniform [0,1] and is independent of the data $D_{n}$. With the $k$-NN rule based on the artificial distances $\sigma$, show that the rule is universally consistent by verifying the conditions of Theorem 6.3, when $k \rightarrow \infty$ and $k / n \rightarrow 0$. In particular, show first that if $Z$ is uniform [0, 1] and independent of $X, Y, D_{n}$ and $Z_{1}, \ldots, Z_{n}$, and if $W_{n i}(X, Z)$ is the weight of $\left(X_{i}, Z_{i}\right)$ in this $k$-nN rule (i.e., it is $1 / k$ iff $\left(X_{i}, Z_{i}\right)$ is among the $k$ nearest neighbors of ( $X, Z$ ) according to $\sigma$ ), then

$$
\begin{equation*}
\mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X, Z) f\left(X_{i}\right)\right\} \leq 2^{d} \mathbf{E} f(X) \tag{1}
\end{equation*}
$$

for all nonnegative measurable $f$ with $\mathbf{E} f(X)<\infty$.
(2) If $k / n \rightarrow 0$, then

$$
\mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X, Z) I_{\left\{\left\|X_{i}-X\right\|>a\right\}}\right\}
$$

for all $a>0$.
Hint: Check the conditions of Theorem 6.3 (Devroye (1978)).
Problem 11.6. The layered nearest neighbor rule partitions the space at $x$ into $2^{d}$ quadrants. In each quadrant, the outer-layer points are marked, that is, those $X_{i}$ for which the hyperrectangle defined by $x$ and $X_{i}$ contains no other data point. Then it takes a majority vote over the $Y_{i}$ 's for the marked points. Observe that this rule is scale-invariant. Show that whenever $X$ has nonatomic marginals (to avoid ties), $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ in probability.


Hint: It suffices to show that the number of marked points increases unboundedly in probability, and that its proportion to unmarked points tends to zero in probability.

Problem 11.7. Prove weak universal consistency of the weighted nearest neighbor rule if the weights satisfy

$$
\lim _{n \rightarrow \infty} \max _{1 \leq i \leq n} w_{n i}=0
$$

and

$$
\lim _{n \rightarrow \infty} \sum_{k \leq i \leq n} w_{n i}=0
$$

for some $k$ with $k / n \rightarrow 0$ (Stone (1977)). Hint: Check the conditions of Theorem 6.3.
Problem 11.8. If $\left(w_{n 1}, \ldots, w_{n n}\right)$ is a probability vector, then $\lim _{n \rightarrow \infty} \sum_{i>n \delta} w_{n i}=0$ for all $\delta>0$ if and only if there exists a sequence of integers $k=k_{n}$ such that $k=o(n)$, $k \rightarrow \infty$, and $\sum_{i>k} w_{n i}=o(1)$. Show this. Conclude that the conditions of Problem 11.7 are equivalent to
(i) $\lim _{n \rightarrow \infty} \max _{1 \leq i \leq n} w_{n i}=0$,

$$
\begin{equation*}
\lim _{n \rightarrow \infty} \sum_{i>n \delta} w_{n i}=0 \text { for all } \delta>0 \tag{ii}
\end{equation*}
$$

Problem 11.9. Verify the conditions of Problems 11.7 and 11.8 for weight vectors of the form $w_{n i}=c_{n} / i^{\alpha}$, where $\alpha>0$ is a constant and $c_{n}$ is a normalizing constant. In particular, check that they do not hold for $\alpha>1$ but that they do hold for $0<\alpha \leq 1$.

Problem 11.10. Consider

$$
w_{n i}=\frac{\rho_{n}}{\left(1+\rho_{n}\right)^{i}} \times \frac{1}{1-\left(1+\rho_{n}\right)^{-n}}, \quad 1 \leq i \leq n,
$$

as weight vector. Show that the conditions of Problem 11.8 hold if $\rho_{n} \rightarrow 0$ and $n \rho_{n} \rightarrow \infty$.

Problem 11.11. Let $w_{n i}=\mathbf{P}\{Z=i\} / \mathbf{P}\{Z \leq n\}, 1 \leq i \leq n$, where $Z$ is a Poisson $\left(\lambda_{n}\right)$ random variable. Show that there is no choice of $\lambda_{n}$ such that $\left\{w_{n i}, 1 \leq i \leq n\right\}$ is a consistent weight sequence following the conditions of Problem 11.8.
$\operatorname{Problem} 11.12$. Let $w_{n i}=\mathbf{P}\left\{\operatorname{Binomial}\left(n, p_{n}\right)=i\right\}=\binom{n}{i} p_{n}^{i}\left(1-p_{n}\right)^{n-i}$. Derive conditions on $p_{n}$ for this choice of weight sequence to be consistent in the sense of Problem 11.8.

Problem 11.13. $k$-nn density estimation. We recall from Problem 2.11 that if the conditional densities $f_{0}, f_{1}$ exist, then $L_{1}$-consistent density estimation leads to a consistent classification rule. Consider now the $k$-nearest neighbor density estimate introduced by Loftsgaarden and Quesenberry (1965). Let $X_{1}, \ldots, X_{n}$ be independent, identically distributed random variables in $\mathcal{R}^{d}$, with common density $f$. The $k$-NN estimate of $f$ is defined by

$$
f_{n}(x)=\frac{k}{\lambda\left(S_{x,\left\|x-X_{(k)}(x)\right\|}\right)}
$$

where $X_{(k)}(x)$ is the $k$-th nearest neighbor of $x$ among $X_{1}, \ldots, X_{n}$. Show that for every $n, \int\left|f(x)-f_{n}(x)\right| d x=\infty$, so that the density estimate is never consistent in $L_{1}$. On the other hand, according to Problem 11.14, the corresponding classification rule is consistent, so read on.

Problem 11.14. Assume that the conditional densities $f_{0}$ and $f_{1}$ exist. Then we can use a rule suggested by Patrick and Fischer (1970): let $N_{0}=\sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}}$ and $N_{1}=n-N_{0}$ be the number of zeros and ones in the training data. Denote by $X_{0}^{(k)}(x)$ the $k$-th nearest neighbor of $x$ among the $X_{i}$ 's with $Y_{i}=0$. Define $X_{1}^{(k)}(x)$ similarly. If $\lambda(A)$ denotes the volume of a set $A \subset \mathcal{R}^{d}$, then the rule is defined as

$$
g_{n}(x)= \begin{cases}0 & \text { if } N_{0} / \lambda\left(S_{x,\left\|x-x_{0}^{(k)}(x)\right\|}\right) \geq N_{1} / \lambda\left(S_{x,\left\|x-X_{1}^{(k)}(x)\right\|}\right) \\ 1 & \text { otherwise. }\end{cases}
$$

This estimate is based on the $k$-nearest neighbor density estimate introduced by Loftsgaarden and Quesenberry (1965). Their estimate of $f_{0}$ is

$$
\widehat{f_{0, n}}(x)=\frac{k}{\lambda\left(S_{x,\left\|x-X_{0}^{(k)}(x)\right\|}\right)} .
$$

Then the rule $g_{n}$ can be re-written as

$$
g_{n}(x)= \begin{cases}0 & \text { if } p_{0, n} \widehat{f}_{0, n} \geq p_{1, n} \widehat{f}_{1, n} \\ 1 & \text { otherwise }\end{cases}
$$

where $p_{0, n}=N_{0} / n$ and $p_{1, n}=N_{1} / n$ are the obvious estimates of the class probabilities. Show that $g_{n}$ is weakly consistent if $k \rightarrow \infty$ and $k / n \rightarrow 0$, whenever the conditional densities exist.

Problem 11.15. Consider a weakly universally consistent rule $g_{n}$, and define the rule

$$
h_{n}(x)= \begin{cases}1-Y_{i} & \text { if } x=X_{i} \text { and } x \neq X_{j}, \text { all } j \neq i \\ g_{n}\left(x, D_{n}\right) & \text { otherwise }\end{cases}
$$

Show that $h_{n}$ too is weakly universally consistent. Note: it is the atomic (or partially atomic) distributions of $X$ that make this exercise interesting. Hint: The next exercise may help.

Problem 11.16. Let $X$ have an atomic distribution which puts probability $p_{i}$ at atom $i$. Let $X_{1}, \ldots, X_{n}$ be an i.i.d. sample drawn from this distribution. Then show the following.
(1) $\mathbf{P}\{|N-\mathbf{E} N|>\epsilon\} \leq 2 e^{-2 \epsilon^{2} / n}$, where $\epsilon>0$, and $N$ is the number of "occupied" atoms (the number of different values in the data sequence).
(2) $\mathbf{E} N / n \rightarrow 0$.
(3) $N / n \rightarrow 0$ almost surely.
(4) $\sum_{i: X_{j} \neq i}$ for all $j \leq n=0$ almost surely.

Problem 11.17. Royall's rule. Royall (1966) proposes the regression function estimate

$$
\eta_{n}(x)=\frac{1}{n h_{n}} \sum_{i=1}^{\left\lfloor n h_{n}\right\rfloor} J\left(\frac{i}{n h_{n}}\right) Y_{(i)}(x),
$$

where $J(u)$ is a smooth kernel-like function on [0,1] with $\int_{0}^{1} J(u) d u=1$ and $h_{n}>0$ is a smoothing factor. Suggestions included
(i). $\quad J(u) \equiv 1 ;$
(ii) $\quad J(u)=\frac{(d+2)^{2}}{4}\left(1-\frac{d+4}{d+2} u^{2 / d}\right)$
(note that this function becomes negative).
Define

$$
g_{n}(x)= \begin{cases}1 & \text { if } \eta_{n}(x)>1 / 2 \\ 0 & \text { otherwise }\end{cases}
$$

Assume that $h_{n} \rightarrow 0, n h_{n} \rightarrow \infty$. Derive sufficient conditions on $J$ that guarantee the weak universal consistency of Royall's rule. In particular, insure that choice (ii) is weakly universally consistent. Hint: Try adding an appropriate smoothness condition to $J$.

Problem 11.18. Let $K$ be a kernel and let $R_{n}(x)$ denote the distance between $x$ and its $k$-th nearest neighbor, $X_{(k)}(x)$, among $X_{1}, \ldots, X_{n}$. The discrimination rule that corresponds to a kernel-type nearest neighbor regression function estimate of Mack (1981) is

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n}\left(2 Y_{i}-1\right) K\left(\frac{x-X_{i}}{R_{n}(x)}\right) \leq 0 \\ 1 & \text { otherwise }\end{cases}
$$

(The idea of replacing the smoothing factor in the kernel estimate by a local rank-based value such as $R_{n}(x)$ is due to Breiman, Meisel, and Purcell (1977).) For the kernel $K=I_{S_{0,1}}$, this rule coincides with the $k$-nN rule. For regular kernels (see Chapter 10), show that the rule remains weakly universally consistent whenever $k \rightarrow \infty$ and $k / n \rightarrow 0$ by verifying Stone's conditions of Theorem 6.3.

Problem 11.19. Let $\left\{g_{n}\right\}$ be a weakly universally consistent sequence of classifiers. Split the data sequence $D_{n}$ into two parts: $D_{m}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{m}, Y_{m}\right)\right)$ and $T_{n-m}=\left(\left(X_{m+1}\right.\right.$, $\left.\left.Y_{m+1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$. Use $g_{n-m}$ and the second part $T_{n-m}$ to relabel the first part, i.e., define $Y_{i}^{\prime}=g_{n-m}\left(X_{i}, T_{n-m}\right)$ for $i=1, \ldots, m$. Prove that the $1-\mathrm{NN}$ rule based on the data $\left(X_{1}, Y_{1}^{\prime}\right)$, $\ldots,\left(X_{m}, Y_{m}^{\prime}\right)$ is weakly universally consistent whenever $m \rightarrow \infty$ and $n-m \rightarrow \infty$. HinT: Use Problem 5.40.

Problem 11.20. Consider the $k$-nN rule with a fixed $k$ as the ancestral rule, and apply the $1-\mathrm{NN}$ rule using the relabeled data. Investigate the convergence of $\mathbf{E} L_{n}$. Is the limit $L_{k \mathrm{NN}}$ or something else?

## 12

## Vapnik-Chervonenkis Theory

### 12.1 Empirical Error Minimization

In this chapter we select a decision rule from a class of rules with the help of training data. Working formally, let $\mathcal{C}$ be a class of functions $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$. One wishes to select a function from $\mathcal{C}$ with small error probability. Assume that the training data $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$ are given to pick one of the functions from $\mathcal{C}$ to be used as a classifier. Perhaps the most natural way of selecting a function is to minimize the empirical error probability

$$
\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi\left(X_{i}\right) \neq Y_{i}\right\}}
$$

over the class $\mathcal{C}$. Denote the empirically optimal rule by $\phi_{n}^{*}$ :

$$
\phi_{n}^{*}=\underset{\phi \in \mathcal{C}}{\arg \min } \widehat{L}_{n}(\phi)
$$

Thus, $\phi_{n}^{*}$ is the classifier that, according to the data $D_{n}$, "looks best" among the classifiers in $\mathcal{C}$. This idea of minimizing the empirical risk in the construction of a rule was developed to great extent by Vapnik and Chervonenkis (1971; 1974c; 1974a; 1974b).

Intuitively, the selected classifier $\phi_{n}^{*}$ should be good in the sense that its true error probability $L\left(\phi_{n}^{*}\right)=\mathbf{P}\left\{\phi_{n}^{*}(X) \neq Y \mid D_{n}\right\}$ is expected to be close to the optimal error probability within the class. Their difference is the quantity that primarily interests us in this chapter:

$$
L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)
$$

The latter difference may be bounded in a distribution-free manner, and a rate of convergence results that only depends on the structure of $\mathcal{C}$. While this is very exciting, we must add that $L\left(\phi_{n}^{*}\right)$ may be far away from the Bayes error $L^{*}$. Note that

$$
L\left(\phi_{n}^{*}\right)-L^{*}=\left(L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)\right)+\left(\inf _{\phi \in \mathcal{C}} L(\phi)-L^{*}\right)
$$

The size of $\mathcal{C}$ is a compromise: when $\mathcal{C}$ is large, $\inf _{\phi \in \mathcal{C}} L(\phi)$ may be close to $L^{*}$, but the former error, the estimation error, is probably large as well. If $\mathcal{C}$ is too small, there is no hope to make the approximation error $\inf _{\phi \in \mathcal{C}} L(\phi)-L^{*}$ small. For example, if $\mathcal{C}$ is the class of all (measurable) decision functions, then we can always find a classifier in $\mathcal{C}$ with zero empirical error, but it may have arbitrary values outside of the points $X_{1}, \ldots, X_{n}$. For example, an empirically optimal classifier is

$$
\phi_{n}^{*}(x)= \begin{cases}Y_{i} & \text { if } x=X_{i}, i=1, \ldots, n \\ 0 & \text { otherwise }\end{cases}
$$

This is clearly not what we are looking for. This phenomenon is called overfitting, as the overly large class $\mathcal{C}$ overfits the data. We will give precise conditions on $\mathcal{C}$ that allow us to avoid this anomaly. The choice of $\mathcal{C}$ such that $\inf _{\phi \in \mathcal{C}} L(\phi)$ is close to $L^{*}$ has been the subject of various chapters on consistency-just assume that $\mathcal{C}$ is allowed to grow with $n$ in some manner. Here we take the point of view that $\mathcal{C}$ is fixed, and that we have to live with the functions in $\mathcal{C}$. The best we may then hope for is to minimize $L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)$. A typical situation is shown in Figure 12.1.


Figure 12.1. Various errors in empirical classifier selection.

Consider first a finite collection $\mathcal{C}$, and assume that one of the classifiers in $\mathcal{C}$ has zero error probability, that is, $\min _{\phi \in \mathcal{C}} L(\phi)=0$. Then clearly, $\widehat{L}_{n}\left(\phi_{n}^{*}\right)=0$ with probability one. We then have the following performance bound:

Theorem 12.1. (Vapnik and Chervonenkis (1974c)). Assume $|\mathcal{C}|<\infty$ and $\min _{\phi \in \mathcal{C}} L(\phi)=0$. Then for every $n$ and $\epsilon>0$,

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)>\epsilon\right\} \leq|\mathcal{C}| e^{-n \epsilon},
$$

and

$$
\mathbf{E}\left\{L\left(\phi_{n}^{*}\right)\right\} \leq \frac{1+\log |\mathcal{C}|}{n}
$$

Proof. Clearly,

$$
\begin{aligned}
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)>\epsilon\right\} & \leq \mathbf{P}\left\{\max _{\phi \in \mathcal{C}: \widehat{L}_{n}(\phi)=0} L(\phi)>\epsilon\right\} \\
& =\mathbf{E}\left\{I_{\left\{\max _{\phi \in C: \overparen{L}_{n}(\phi)=0} L(\phi)>\epsilon\right\}}\right\} \\
& =\mathbf{E}\left\{\max _{\phi \in \mathcal{C}} I_{\left\{\widehat{L}_{n}(\phi)=0\right\}} I_{\{L(\phi)>\epsilon\}}\right\} \\
& \leq \sum_{\phi \in \mathcal{C}: L(\phi)>\epsilon} \mathbf{P}\left\{\widehat{L}_{n}(\phi)=0\right\} \\
& \leq|\mathcal{C}|(1-\epsilon)^{n},
\end{aligned}
$$

since the probability that no $\left(X_{i}, Y_{i}\right)$ pair falls in the set $\{(x, y): \phi(x) \neq y\}$ is less than $(1-\epsilon)^{n}$ if the probability of the set is larger than $\epsilon$. The probability inequality of the theorem follows from the simple inequality $1-x \leq e^{-x}$.

To bound the expected error probability, note that for any $u>0$,

$$
\begin{aligned}
\mathbf{E}\left\{L\left(\phi_{n}^{*}\right)\right\} & =\int_{0}^{\infty} \mathbf{P}\left\{L\left(\phi_{n}^{*}\right)>t\right\} d t \\
& \leq u+\int_{u}^{\infty} \mathbf{P}\left\{L\left(\phi_{n}^{*}\right)>t\right\} d t \\
& \leq u+|\mathcal{C}| \int_{u}^{\infty} e^{-n t} d t \\
& =u+\frac{|\mathcal{C}|}{n} e^{-n u}
\end{aligned}
$$

Since $u$ was arbitrary, we may choose it to minimize the obtained upper bound. The optimal choice is $u=\log |\mathcal{C}| / n$, which yields the desired inequality.

Theorem 12.1 shows that empirical selection works very well if the sample size $n$ is much larger than the logarithm of the size of the family $\mathcal{C}$. Unfortunately, the
assumption on the distribution of $(X, Y)$, that is, that $\min _{\phi \in \mathcal{C}} L(\phi)=0$, is very restrictive. In the sequel we drop this assumption, and deal with the distributionfree problem.

One of our main tools is taken from Lemma 8.2:

$$
L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi) \leq 2 \sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|
$$

This leads to the study of uniform deviations of relative frequencies from their probabilities by the following simple observation: let $v$ be a probability measure of $(X, Y)$ on $\mathcal{R}^{d} \times\{0,1\}$, and let $v_{n}$ be the empirical measure based upon $D_{n}$. That is, for any fixed measurable set $A \subset \mathcal{R}^{d} \times\{0,1\}, v(A)=\mathbf{P}\{(X, Y) \in A\}$, and $v_{n}(A)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\left(X_{i}, Y_{i}\right) \in A\right\}}$. Then

$$
L(\phi)=v(\{(x, y): \phi(x) \neq y\})
$$

is just the $v$-measure of the set of pairs $(x, y) \in \mathcal{R}^{d} \times\{0,1\}$, where $\phi(x) \neq y$. Formally, $L(\phi)$ is the $v$-measure of the set

$$
\{\{x: \phi(x)=1\} \times\{0\}\} \bigcup\{\{x: \phi(x)=0\} \times\{1\}\}
$$

Similarly, $\widehat{L_{n}}(\phi)=v_{n}(\{(x, y): \phi(x) \neq y\})$. Thus,

$$
\sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|=\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|,
$$

where $\mathcal{A}$ is the collection of all sets

$$
\{\{x: \phi(x)=1\} \times\{0\}\} \bigcup\{\{x: \phi(x)=0\} \times\{1\}\}, \phi \in \mathcal{C} .
$$

For a fixed set $A$, for any probability measure $v$, by the law of large numbers $v_{n}(A)-v(A) \rightarrow 0$ almost surely as $n \rightarrow \infty$. Moreover, by Hoeffding's inequality (Theorem 8.1),

$$
\mathbf{P}\left\{\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 2 e^{-2 n \epsilon^{2}}
$$

However, it is a much harder problem to obtain such results for $\sup _{A \in \mathcal{A}} \mid v_{n}(A)-$ $v(A) \mid$. If the class of sets $\mathcal{A}$ (or, analogously, in the pattern recognition context, $\mathcal{C}$ ) is of finite cardinality, then the union bound trivially gives

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 2|\mathcal{A}| e^{-2 n \epsilon^{2}}
$$

However, if $\mathcal{A}$ contains infinitely many sets (as in many of the interesting cases) then the problem becomes nontrivial, spawning a vast literature. The most powerful weapons to attack these problems are distribution-free large deviation-type inequalities first proved by Vapnik and Chervonenkis (1971) in their pioneering work. However, in some situations, we can handle the problem in a much simpler way. We have already seen such an example in Section 4.5.

### 12.2 Fingering

Recall that in Section 4.5 we studied a specific rule that selects a linear classifier by minimizing the empirical error. The performance bounds provided by Theorems 4.5 and 4.6 show that the selected rule performs very closely to the best possible linear rule. These bounds apply only to the specific algorithm used to find the empirical minima-we have not showed that any classifier minimizing the empirical error performs well. This matter will be dealt with in later sections. In this section we extend Theorems 4.5 and 4.6 to classes other than linear classifiers.

Let $\mathcal{C}$ be the class of classifiers assigning 1 to those $x$ contained in a closed hyperrectangle, and 0 to all other points. Then a classifier minimizing the empirical error $\widehat{L}_{n}(\phi)$ over all $\phi \in \mathcal{C}$ may be obtained by the following algorithm: to each $2 d$-tuple ( $X_{i_{1}}, \ldots, X_{i_{2 d}}$ ) of points from $X_{1}, \ldots, X_{n}$, assign the smallest hyperrectangle containing these points. If we assume that $X$ has a density, then the points $X_{1}, \ldots, X_{n}$ are in general position with probability one. This way we obtain at most $\binom{n}{2 d}$ sets. Let $\phi_{i}$ be the classifier corresponding to the $i$-th such hyperrectangle, that is, the one assigning 1 to those $x$ contained in the hyperrectangle, and 0 to other points. Clearly, for each $\phi \in \mathcal{C}$, there exists a $\phi_{i}, i=1, \ldots,\binom{n}{2 d}$, such that

$$
\phi\left(X_{j}\right)=\phi_{i}\left(X_{j}\right)
$$

for all $X_{j}$, except possibly for those on the boundary of the hyperrectangle. Since the points are in general position, there are at most $2 d$ such exceptional points. Therefore, if we select a classifier $\widehat{\phi}$ among $\phi_{1}, \ldots, \phi_{\binom{n}{d}}$ to minimize the empirical error, then it approximately minimizes the empirical error over the whole class $\mathcal{C}$ as well. A quick scan through the proof of Theorem 4.5 reveals that by similar arguments we may obtain the performance bound

$$
\mathbf{P}\left\{L(\widehat{\phi})-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq e^{4 d \epsilon}\left(\binom{n}{2 d}+1\right) e^{-n \epsilon^{2} / 2}
$$

for $n \geq 2 d$ and $\epsilon \geq 4 d / n$.
The idea may be generalized. It always works, if for some $k, k$-tuples of points determine classifiers from $\mathcal{C}$ such that no matter where the other data points fall, the minimal empirical error over these sets coincides with the overall minimum. Then we may fix these sets-"put our finger on them"-and look for the empirical minima over this finite collection. The next theorems, whose proofs are left as an exercise (Problem 12.2), show that if $\mathcal{C}$ has this property, then "fingering" works extremely well whenever $n \gg k$.

Theorem 12.2. Assume that the class $\mathcal{C}$ of classifiers has the following property: for some integer $k$ there exists a function $\Psi:\left(\mathcal{R}^{d}\right)^{k} \rightarrow \mathcal{C}$ such that for all $x_{1}, \ldots, x_{n} \in \mathcal{R}^{d}$ and all $\phi \in \mathcal{C}$, there exists a $k$-tuple $i_{1}, \ldots, i_{k} \in\{1, \ldots, n\}$ of different indices such that

$$
\Psi\left(x_{i_{1}}, \ldots, x_{i_{k}}\right)\left(x_{j}\right)=\phi\left(x_{j}\right) \quad \text { for all } j=1, \ldots, n \text { with } j \neq i_{l}, l=1, \ldots, k
$$

with probability one. Let $\widehat{\phi}$ be found by fingering, that is, by empirical error minimization over the collection of $n!/(n-k)$ ! classifiers of the form

$$
\Psi\left(X_{i_{1}}, \ldots, X_{i_{k}}\right), \quad i_{1}, \ldots, i_{k} \in\{1, \ldots, n\}, \quad \text { different } .
$$

Then for $n \geq k$ and $2 k / n \leq \epsilon \leq 1$,

$$
\mathbf{P}\left\{L(\widehat{\phi})-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq e^{2 k \epsilon}\left(n^{k}+1\right) e^{-n \epsilon^{2} / 2}
$$

Moreover, if $n \geq k$, then

$$
\mathbf{E}\left\{L(\widehat{\phi})-\inf _{\phi \in \mathcal{C}} L(\phi)\right\} \leq \sqrt{\frac{2(k+1) \log n+(2 k+2)}{n}}
$$

The smallest $k$ for which $\mathcal{C}$ has the property described in the theorem may be called the fingering dimension of $\mathcal{C}$. In most interesting cases, it is independent of $n$. Problem 12.3 offers a few such classes. We will see later in this chapter that the fingering dimension is closely related the so-called vc dimension of $\mathcal{C}$ (see also Problem 12.4).

Again, we get much smaller errors if $\inf _{\phi \in \mathcal{C}} L(\phi)=0$. The next inequality generalizes Theorem 4.6.

Theorem 12.3. Assume that $\mathcal{C}$ has the property described in Theorem 12.2 with fingering dimension $k$. Assume, in addition, that $\inf _{\phi \in \mathcal{C}} L(\phi)=0$. Then for all $n$ and $\epsilon$,

$$
\mathbf{P}\{L(\widehat{\phi})>\epsilon\} \leq n^{k} e^{-(n-k) \epsilon}
$$

and

$$
\mathbf{E}\{L(\widehat{\phi})\} \leq \frac{k \log n+2}{n-k}
$$

Remark. Even though the results in the next few sections based on the VapnikChervonenkis theory supersede those of this section (by requiring less from the class $\mathcal{C}$ and being able to bound the error of any classifier minimizing the empirical risk), we must remark that the exponents in the above probability inequalities are the best possible, and bounds of the same type for the general case can only be obtained with significantly more effort.

### 12.3 The Glivenko-Cantelli Theorem

In the next two sections, we prove the Vapnik-Chervonenkis inequality, a powerful generalization of the classical Glivenko-Cantelli theorem. It provides upper bounds on random variables of the type

$$
\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right| .
$$

As we noted in Section 12.1, such bounds yield performance bounds for any classifier selected by minimizing the empirical error. To make the material more digestible, we first present the main ideas in a simple one-dimensional setting, and then prove the general theorem in the next section.

We drop the pattern recognition setting momentarily, and return to probability theory. The following theorem is sometimes referred to as the fundamental theorem of mathematical statistics, stating uniform almost sure convergence of the empirical distribution function to the true one:

Theorem 12.4. (Glivenko-Cantelli theorem). Let $Z_{1}, \ldots, Z_{n}$ be i.i.d. realvalued random variables with distribution function $F(z)=\mathbf{P}\left\{Z_{1} \leq z\right\}$. Denote the standard empirical distribution function by

$$
F_{n}(z)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{Z_{i} \leq z\right\}} .
$$

Then

$$
\mathbf{P}\left\{\sup _{z \in \mathcal{R}}\left|F(z)-F_{n}(z)\right|>\epsilon\right\} \leq 8(n+1) e^{-n \epsilon^{2} / 32}
$$

and, in particular, by the Borel-Cantelli lemma,

$$
\lim _{n \rightarrow \infty} \sup _{z \in \mathcal{R}}\left|F(z)-F_{n}(z)\right|=0 \quad \text { with probability one } .
$$

Proof. The proof presented here is not the simplest possible, but it contains the main ideas leading to a powerful generalization. Introduce the notation $\nu(A)=$ $\mathbf{P}\left\{Z_{1} \in A\right\}$ and $\nu_{n}(A)=(1 / n) \sum_{j=1}^{n} I_{\left\{Z_{j} \in A\right\}}$ for all measurable sets $A \subset \mathcal{R}$. Let $\mathcal{A}$ denote the class of sets of form $(-\infty, z]$ for $z \in \mathcal{R}$. With these notations,

$$
\sup _{z \in \mathcal{R}}\left|F(z)-F_{n}(z)\right|=\sup _{A \in \mathcal{A}}\left|\nu_{n}(A)-v(A)\right| .
$$

We prove the theorem in several steps, following symmetrization ideas of Dudley (1978), and Pollard (1984). We assume that $n \epsilon^{2} \geq 2$, since otherwise the bound is trivial. In the first step we introduce a symmetrization.

Step 1. First symmetrization by a ghost sample. Define the random variables $Z_{1}^{\prime}, \ldots, Z_{n}^{\prime} \in \mathcal{R}$ such that $Z_{1}, \ldots, Z_{n}, Z_{1}^{\prime}, \ldots, Z_{n}^{\prime}$ are all independent and identically distributed. Denote by $v_{n}^{\prime}$ the empirical measure corresponding to the new sample:

$$
v_{n}^{\prime}(A)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{Z_{i}^{\prime} \in A\right\}}
$$

Then for $n \epsilon^{2} \geq 2$ we have

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 2 \mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v_{n}^{\prime}(A)\right|>\frac{\epsilon}{2}\right\} .
$$

To see this, let $A^{*} \in \mathcal{A}$ be a set for which $\left|v_{n}\left(A^{*}\right)-v\left(A^{*}\right)\right|>\epsilon$ if such a set exists, and let $A^{*}$ be a fixed set in $\mathcal{A}$ otherwise. Then

$$
\begin{aligned}
\mathbf{P} & \left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v_{n}^{\prime}(A)\right|>\epsilon / 2\right\} \\
& \geq \mathbf{P}\left\{\left|v_{n}\left(A^{*}\right)-v_{n}^{\prime}\left(A^{*}\right)\right|>\epsilon / 2\right\} \\
& \geq \mathbf{P}\left\{\left|v_{n}\left(A^{*}\right)-v\left(A^{*}\right)\right|>\epsilon,\left|v_{n}^{\prime}\left(A^{*}\right)-v\left(A^{*}\right)\right|<\frac{\epsilon}{2}\right\} \\
& =\mathbf{E}\left\{I_{\left\{\left|v_{n}\left(A^{*}\right)-v\left(A^{*}\right)\right|>\epsilon\right\}} \mathbf{P}\left\{\left.\left|v_{n}^{\prime}\left(A^{*}\right)-v\left(A^{*}\right)\right|<\frac{\epsilon}{2} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\}\right\} .
\end{aligned}
$$

The conditional probability inside may be bounded by Chebyshev's inequality as follows:

$$
\begin{aligned}
\mathbf{P}\left\{\left.\left|v_{n}^{\prime}\left(A^{*}\right)-v\left(A^{*}\right)\right|<\frac{\epsilon}{2} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} & \geq 1-\frac{v\left(A^{*}\right)\left(1-v\left(A^{*}\right)\right)}{n \epsilon^{2} / 4} \\
& \geq 1-\frac{1}{n \epsilon^{2}} \geq \frac{1}{2}
\end{aligned}
$$

whenever $n \epsilon^{2} \geq 2$. In summary,

$$
\begin{aligned}
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v_{n}^{\prime}(A)\right|>\epsilon / 2\right\} & \geq \frac{1}{2} \mathbf{P}\left\{\left|v_{n}\left(A^{*}\right)-v\left(A^{*}\right)\right|>\epsilon\right\} \\
& \geq \frac{1}{2} \mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\}
\end{aligned}
$$

STEP 2. SECOND SYMMETRIZATION BY RANDOM SIGNS. Let $\sigma_{1}, \ldots, \sigma_{n}$ be i.i.d. sign variables, independent of $Z_{1}, \ldots, Z_{n}$ and $Z_{1}^{\prime}, \ldots, Z_{n}^{\prime}$, with $\mathbf{P}\left\{\sigma_{i}=-1\right\}=\mathbf{P}\left\{\sigma_{i}=\right.$ $1\}=1 / 2$. Clearly, because $Z_{1}, Z_{1}^{\prime}, \ldots, Z_{n}, Z_{n}^{\prime}$ are all independent and identically distributed, the distribution of

$$
\sup _{A \in \mathcal{A}}\left|\sum_{i=1}^{n}\left(I_{A}\left(Z_{i}\right)-I_{A}\left(Z_{i}^{\prime}\right)\right)\right|
$$

is the same as the distribution of

$$
\sup _{A \in \mathcal{A}}\left|\sum_{i=1}^{n} \sigma_{i}\left(I_{A}\left(Z_{i}\right)-I_{A}\left(Z_{i}^{\prime}\right)\right)\right| .
$$

Thus, by Step 1,

$$
\begin{aligned}
& \mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \\
& \quad \leq 2 \mathbf{P}\left\{\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n}\left(I_{A}\left(Z_{i}\right)-I_{A}\left(Z_{i}^{\prime}\right)\right)\right|>\frac{\epsilon}{2}\right\} \\
& \\
& \quad=2 \mathbf{P}\left\{\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i}\left(I_{A}\left(Z_{i}\right)-I_{A}\left(Z_{i}^{\prime}\right)\right)\right|>\frac{\epsilon}{2}\right\} .
\end{aligned}
$$

Simply applying the union bound, we can remove the auxiliary random variables $Z_{1}^{\prime}, \ldots, Z_{n}^{\prime}$ :

$$
\begin{aligned}
& \mathbf{P}\left\{\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i}\left(I_{A}\left(Z_{i}\right)-I_{A}\left(Z_{i}^{\prime}\right)\right)\right|>\frac{\epsilon}{2}\right\} \\
& \\
& \leq \mathbf{P}\left\{\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4}\right\}+\mathbf{P}\left\{\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}^{\prime}\right)\right|>\frac{\epsilon}{4}\right\} \\
& \\
& =2 \mathbf{P}\left\{\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4}\right\} .
\end{aligned}
$$

STEP 3. COnditioning. To bound the probability

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4}\right\}=\mathbf{P}\left\{\sup _{z \in \mathcal{R}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{\left\{Z_{i} \leq z\right\}}\right|>\frac{\epsilon}{4}\right\},
$$

we condition on $Z_{1}, \ldots, Z_{n}$. Fix $z_{1}, \ldots, z_{n} \in \mathcal{R}^{d}$, and note that as $z$ ranges over $\mathcal{R}$, the number of different vectors $\left(I_{\left\{z_{1} \leq z\right\}}, \ldots, I_{\left\{z_{n} \leq z\right\}}\right)$ is at most $n+1$. Thus, conditional on $Z_{1}, \ldots, Z_{n}$, the supremum in the probability above is just a maximum taken over at most $n+1$ random variables. Thus, applying the union bound gives

$$
\begin{aligned}
& \mathbf{P}\left\{\left.\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} \\
& \quad \leq(n+1) \sup _{A \in \mathcal{A}} \mathbf{P}\left\{\left.\frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} .
\end{aligned}
$$

With the supremum now outside the probability, it suffices to find an exponential bound on the conditional probability

$$
\mathbf{P}\left\{\left.\frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} .
$$

Step 4. Hoeffding's inequality. With $z_{1}, \ldots, z_{n}$ fixed, $\sum_{i=1}^{n} \sigma_{i} I_{A}\left(z_{i}\right)$ is the sum of $n$ independent zero mean random variables bounded between -1 and 1 . Therefore, Theorem 8.1 applies in a straightforward manner:

$$
\mathbf{P}\left\{\left.\frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} \leq 2 e^{-n \epsilon^{2} / 32}
$$

Thus,

$$
\mathbf{P}\left\{\left.\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} \leq 2(n+1) e^{-n \epsilon^{2} / 32} .
$$

Taking the expected value on both sides we have

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4}\right\} \leq 2(n+1) e^{-n \epsilon^{2} / 32}
$$

In summary,

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 8(n+1) e^{-n \epsilon^{2} / 32}
$$

### 12.4 Uniform Deviations of Relative Frequencies from Probabilities

In this section we prove the Vapnik-Chervonenkis inequality, a mighty generalization of Theorem 12.4. In the proof we need only a slight adjustment of the proof above. In the general setting, let the independent identically distributed random variables $Z_{1}, \ldots, Z_{n}$ take their values from $\mathcal{R}^{d}$. Again, we use the notation $\nu(A)=\mathbf{P}\left\{Z_{1} \in A\right\}$ and $v_{n}(A)=(1 / n) \sum_{j=1}^{n} I_{\left\{Z_{j} \in A\right\}}$ for all measurable sets $A \subset \mathcal{R}^{d}$. The Vapnik-Chervonenkis theory begins with the concepts of shatter coefficient and Vapnik-Chervonenkis (or VC) dimension:

Definition 12.1. Let $\mathcal{A}$ be a collection of measurable sets. For $\left(z_{1}, \ldots, z_{n}\right) \in$ $\left\{\mathcal{R}^{d}\right\}^{n}$, let $N_{\mathcal{A}}\left(z_{1}, \ldots, z_{n}\right)$ be the number of different sets in

$$
\left\{\left\{z_{1}, \ldots, z_{n}\right\} \cap A ; A \in \mathcal{A}\right\}
$$

The n-th shatter coefficient of $\mathcal{A}$ is

$$
s(\mathcal{A}, n)=\max _{\left(z_{1}, \ldots, z_{n}\right) \in\left\{\mathcal{R}^{d}\right\}^{n}} N_{\mathcal{A}}\left(z_{1}, \ldots, z_{n}\right) .
$$

That is, the shatter coefficient is the maximal number of different subsets of $n$ points that can be picked out by the class of sets $\mathcal{A}$.

The shatter coefficients measure the richness of the class $\mathcal{A}$. Clearly, $s(\mathcal{A}, n) \leq$ $2^{n}$, as there are $2^{n}$ subsets of a set with $n$ elements. If $N_{\mathcal{A}}\left(z_{1}, \ldots, z_{n}\right)=2^{n}$ for some $\left(z_{1}, \ldots, z_{n}\right)$, then we say that $\mathcal{A}$ shatters $\left\{z_{1}, \ldots, z_{n}\right\}$. If $s(\mathcal{A}, n)<2^{n}$, then any set of $n$ points has a subset such that there is no set in $\mathcal{A}$ that contains exactly that subset of the $n$ points. Clearly, if $s(\mathcal{A}, k)<2^{k}$ for some integer $k$, then $s(\mathcal{A}, n)<2^{n}$ for all $n>k$. The first time when this happens is important:

Definition 12.2. Let $\mathcal{A}$ be a collection of sets with $|\mathcal{A}| \geq 2$. The largest integer $k \geq 1$ for which $s(\mathcal{A}, k)=2^{k}$ is denoted by $V_{\mathcal{A}}$, and it is called the VapnikChervonenkis dimension (or vc dimension) of the class $\mathcal{A}$. If $s(\mathcal{A}, n)=2^{n}$ for all $n$, then by definition, $V_{\mathcal{A}}=\infty$.

For example, if $\mathcal{A}$ contains all halflines of form $(-\infty, x], x \in \mathcal{R}$, then $s(\mathcal{A}, 2)=$ $3<2^{2}$, and $V_{\mathcal{A}}=1$. This is easily seen by observing that for any two different points $z_{1}<z_{2}$ there is no set of the form $(-\infty, x]$ that contains $z_{2}$, but not $z_{1}$. A class of sets $\mathcal{A}$ for which $V_{\mathcal{A}}<\infty$ is called a Vapnik-Chervonenkis (or vc) class. In a sense, $V_{\mathcal{A}}$ may be considered as the complexity, or size, of $\mathcal{A}$. Several properties of the shatter coefficients and the vc dimension will be shown in Chapter 13. The main purpose of this section is to prove the following important result by Vapnik and Chervonenkis (1971):

Theorem 12.5. (VAPNIK And Chervonenkis (1971)). For any probability measure $\nu$ and class of sets $\mathcal{A}$, and for any $n$ and $\epsilon>0$,

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 8 s(\mathcal{A}, n) e^{-n \epsilon^{2} / 32} .
$$

Proof. The proof parallels that of Theorem 12.4. We may again assume that $n \epsilon^{2} \geq 2$. In the first two steps we prove that

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 4 \mathbf{P}\left\{\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4}\right\} .
$$

This may be done exactly the same way as in Theorem 12.4; we do not repeat the argument. The only difference appears in Step 3:
Step 3. Conditioning. To bound the probability

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4}\right\},
$$

again we condition on $Z_{1}, \ldots, Z_{n}$. Fix $z_{1}, \ldots, z_{n} \in \mathcal{R}^{d}$, and observe that as $A$ ranges over $\mathcal{A}$, the number of different vectors $\left(I_{A}\left(z_{1}\right), \ldots, I_{A}\left(z_{n}\right)\right)$ is just the number of different subsets of $\left\{z_{1}, \ldots, z_{n}\right\}$ produced by intersecting it with sets in $\mathcal{A}$, which, by definition, cannot exceed $s(\mathcal{A}, n)$. Therefore, with $Z_{1}, \ldots, Z_{n}$ fixed, the supremum in the above probability is a maximum of at most $N_{\mathcal{A}}\left(Z_{1}, \ldots, Z_{n}\right)$ random variables. This number, by definition, is bounded from above by $s(\mathcal{A}, n)$. By the union bound we get

$$
\begin{aligned}
& \mathbf{P}\left\{\left.\sup _{A \in \mathcal{A}} \frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} \\
& \quad \leq s(\mathcal{A}, n) \sup _{A \in \mathcal{A}} \mathbf{P}\left\{\left.\frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} .
\end{aligned}
$$

Therefore, as before, it suffices to bound the conditional probability

$$
\mathbf{P}\left\{\left.\frac{1}{n}\left|\sum_{i=1}^{n} \sigma_{i} I_{A}\left(Z_{i}\right)\right|>\frac{\epsilon}{4} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} .
$$

This may be done by Hoeffding's inequality exactly as in Step 4 of the proof of Theorem 12.4. Finally, we obtain

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 8 s(\mathcal{A}, n) e^{-n \epsilon^{2} / 32}
$$

The bound of Theorem 12.5 is useful when the shatter coefficients do not increase too quickly with $n$. For example, if $\mathcal{A}$ contains all Borel sets of $\mathcal{R}^{d}$, then we can shatter any collection of $n$ different points at will, and obtain $s(\mathcal{A}, n)=2^{n}$. This would be useless, of course. The smaller $\mathcal{A}$, the smaller the shatter coefficient is. To apply the vc bound, it suffices to compute shatter coefficients for certain families of sets. Examples may be found in Cover (1965), Vapnik and Chervonenkis (1971), Devroye and Wagner (1979a), Feinholz (1979), Devroye (1982a), Massart (1983), Dudley (1984), Simon (1991), and Stengle and Yukich (1989). This list of references is far from exhaustive. More information about shatter coefficients is given in Chapter 13.
Remark. Measurability. The supremum in Theorem 12.5 is not always measurable. Measurability must be verified for every family $\mathcal{A}$. For all our examples, the quantities are indeed measurable. For more on the measurability question, see Dudley (1978; 1984), Massart (1983), and Gaenssler (1983). Giné and Zinn (1984) and Yukich (1985) provide further work on suprema of the type shown in Theorem 12.5.

Remark. Optimal exponent. For the sake of readability we followed the line of Pollard's proof (1984) instead of the original by Vapnik and Chervonenkis. In particular, the exponent $-n \epsilon^{2} / 32$ in Theorem 12.5 is worse than the $-n \epsilon^{2} / 8$ established in the original paper. The best known exponents together with some other related results are mentioned in Section 12.8. The basic ideas of the original proof by Vapnik and Chervonenkis appear in the proof of Theorem 12.7 below.

Remark. Necessary and Sufficient Conditions. It is clear from the proof of the theorem that it can be strengthened to

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 8 \mathbf{E}\left\{N_{\mathcal{A}}\left(Z_{1}, \ldots, Z_{n}\right)\right\} e^{-n \epsilon^{2} / 32}
$$

where $Z_{1}, \ldots, Z_{n}$ are i.i.d. random variables with probability measure $v$. Although this upper bound is tighter than that in the stated inequality, it is usually more difficult to handle, since the coefficient in front of the exponential term depends on the distribution of $Z_{1}$, while $s(\mathcal{A}, n)$ is purely combinatorial in nature. However, this form is important in a different setting: we say that the uniform law of large numbers holds if

$$
\sup \left|v_{n}(A)-v(A)\right| \rightarrow 0 \quad \text { in probability }
$$

It follows from this form of Theorem 12.5 that the uniform law of large numbers holds if

$$
\frac{\mathbf{E}\left\{\log \left(N_{\mathcal{A}}\left(Z_{1}, \ldots, Z_{n}\right)\right)\right\}}{n} \rightarrow 0
$$

Vapnik and Chervonenkis showed $(1971 ; 1981)$ that this condition is also necessary for the uniform law of large numbers. Another characterization of the uniform law of large numbers is given by Talagrand (1987), who showed that the uniform law of large numbers holds if and only if there does not exist a set $A \subset \mathcal{R}^{d}$ with $v(A)>0$ such that, with probability one, the set $\left\{Z_{1}, \ldots, Z_{n}\right\} \cap A$ is shattered by $\mathcal{A}$.

### 12.5 Classifier Selection

The following theorem relates the results of the previous sections to empirical classifier selection, that is, when the empirical error probability $\widehat{L}_{n}(\phi)$ is minimized over a class of classifiers $\mathcal{C}$. We emphasize that unlike in Section 12.2, here we allow any classifier with the property that has minimal empirical error $\widehat{L}_{n}(\phi)$ in $\mathcal{C}$. First introduce the shatter coefficients and vc dimension of $\mathcal{C}$ :

DEfinition 12.3. Let $\mathcal{C}$ be a class of decision functions of the form $\phi: \mathcal{R}^{d} \rightarrow$ $\{0,1\}$. Define $\mathcal{A}$ as the collection of all sets

$$
\{\{x: \phi(x)=1\} \times\{0\}\} \bigcup\{\{x: \phi(x)=0\} \times\{1\}\}, \quad \phi \in \mathcal{C} .
$$

Define the $n$-th shatter coefficient $\mathcal{S}(\mathcal{C}, n)$ of the class of classifiers $\mathcal{C}$ as

$$
\mathcal{S}(\mathcal{C}, n)=s(\mathcal{A}, n)
$$

Furthermore, define the vc dimension $V_{\mathcal{C}}$ of $\mathcal{C}$ as

$$
V_{\mathcal{C}}=V_{\mathcal{A}}
$$

For the performance of the empirically selected decision $\phi_{n}^{*}$, we have the following:

Theorem 12.6. Let $\mathcal{C}$ be a class of decision functions of the form $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$. Then using the notation $\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi\left(X_{i}\right) \neq Y_{i}\right\}}$ and $L(\phi)=\mathbf{P}\{\phi(X) \neq Y\}$, we have

$$
\mathbf{P}\left\{\sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|>\epsilon\right\} \leq 8 \mathcal{S}(\mathcal{C}, n) e^{-n \epsilon^{2} / 32}
$$

and therefore

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq 8 \mathcal{S}(\mathcal{C}, n) e^{-n \epsilon^{2} / 128}
$$

where $\phi_{n}^{*}$ denotes the classifier minimizing $\widehat{L}_{n}(\phi)$ over the class $\mathcal{C}$.

Proof. The statements are immediate consequences of Theorem 12.5 and Lemma 8.2.

The next corollary, an easy application of Theorem 12.6 (see Problem 12.1), makes things a little more transparent.

Corollary 12.1. In the notation of Theorem 12.6,

$$
\mathbf{E}\left\{L\left(\phi_{n}^{*}\right)\right\}-\inf _{\phi \in \mathcal{C}} L(\phi) \leq 16 \sqrt{\frac{\log (8 e \mathcal{S}(\mathcal{C}, n))}{2 n}}
$$

If $\mathcal{S}(\mathcal{C}, n)$ increases polynomially with $n$, then the average error probability of the selected classifier is within $O(\sqrt{\log n / n})$ of the error of the best rule in the class. We point out that this result is completely distribution-free. Furthermore, note the nonasymptotic nature of these inequalities: they hold for every $n$. From here on the problem is purely combinatorial-one has to estimate the shatter coefficients. Many properties are given in Chapter 13. In particular, if $V_{\mathcal{C}}>2$, then $\mathcal{S}(\mathcal{C}, n) \leq$ $n^{V_{C}}$, that is, if the class $\mathcal{C}$ has finite vc dimension, then $\mathcal{S}(\mathcal{C}, n)$ increases at a polynomial rate, and

$$
\mathbf{E}\left\{L\left(\phi_{n}^{*}\right)\right\}-\inf _{\phi \in \mathcal{C}} L(\phi) \leq 16 \sqrt{\frac{V_{\mathcal{C}} \log n+4}{2 n}}
$$

REMARK. Theorem 12.6 provides a bound for the behavior of the empirically optimal classifier. In practice, finding an empirically optimal classifier is often computationally very expensive. In such cases, the designer is often forced to put up with algorithms yielding suboptimal classifiers. Assume for example, that we have an algorithm which selects a classifier $g_{n}$ such that its empirical error is not too far from the optimum with large probability:

$$
\mathbf{P}\left\{\widehat{L}_{n}\left(g_{n}\right) \leq \inf _{\phi \in \mathcal{C}} \widehat{L}_{n}(\phi)+\epsilon_{n}\right\} \geq 1-\delta_{n}
$$

where $\left\{\epsilon_{n}\right\}$ and $\left\{\delta_{n}\right\}$ are sequences of positive numbers converging to zero. Then it is easy to see (see Problem 12.6) that

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq \delta_{n}+\mathbf{P}\left\{2 \sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|>\epsilon-\epsilon_{n}\right\}
$$

and Theorem 12.6 may be used to obtain bounds for the error probability of $g_{n}$. $\square$

Interestingly, the empirical error probability of the empirically optimal classifier is always close to its expected value, as may be seen from the following example:

Corollary 12.2. Let $\mathcal{C}$ be an arbitrary class of classification rules (that is, functions of the form $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$ ). Let $\phi_{n}^{*} \in \mathcal{C}$ be the rule that minimizes the
number of errors committed on the training sequence $D_{n}$, among the classifiers in $\mathcal{C}$. In other words,

$$
\widehat{L}_{n}\left(\phi_{n}^{*}\right) \leq \widehat{L}_{n}(\phi) \quad \text { for every } \phi \in \mathcal{C},
$$

where $\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi\left(X_{i}\right) \neq Y_{i}\right\}}$. Then for every $n$ and $\epsilon>0$,

$$
\mathbf{P}\left\{\left|\widehat{L}_{n}\left(\phi_{n}^{*}\right)-\mathbf{E}\left\{\widehat{L}_{n}\left(\phi_{n}^{*}\right)\right\}\right|>\epsilon\right\} \leq 2 e^{-2 n \epsilon^{2}} .
$$

The corollary follows immediately from Theorem 9.1 by observing that changing the value of one ( $X_{i}, Y_{i}$ ) pair in the training sequence results in a change of at most $1 / n$ in the value of $\widehat{L}_{n}\left(\phi_{n}^{*}\right)$. The corollary is true even if the vC dimension of $\mathcal{C}$ is infinite (!). The result shows that $\widehat{L}_{n}\left(\phi_{n}^{*}\right)$ is always very close to its expected value with large probability, even if $\mathbf{E}\left\{\widehat{L}_{n}\left(\phi_{n}^{*}\right)\right\}$ is far from $\inf _{\phi \in \mathcal{C}} L(\phi)$ (see also Problem 12.12).

### 12.6 Sample Complexity

In his theory of learning, Valiant (1984) rephrases the empirical classifier selection problem as follows. For $\epsilon, \delta>0$, define an $(\epsilon, \delta)$ learning algorithm as a method that selects a classifier $g_{n}$ from $\mathcal{C}$ using the data $D_{n}$ such that for the selected rule

$$
\sup _{(X, Y)} \mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq \delta,
$$

whenever $n \geq N(\epsilon, \delta)$. Here $N(\epsilon, \delta)$ is the sample complexity of the algorithm, defined as the smallest integer with the above property. Since the supremum is taken over all possible distributions of ( $X, Y$ ), the integer $N(\epsilon, \delta)$ is the number of data pairs that guarantees $\epsilon$ accuracy with $\delta$ confidence for any distribution. Note that we use the notation $g_{n}$ and not $\phi_{n}^{*}$ in the definition above, as the definition does not force us to take the empirical risk minimizer $\phi_{n}^{*}$.

We may use Theorem 12.6 to get an upper bound on the sample complexity of the selection algorithm based on empirical error minimization (i.e., the classifier $\phi_{n}^{*}$ ).

COROLLARY 12.3. The sample complexity of the method based on empirical error minimization is bounded from above by

$$
N(\epsilon, \delta) \leq \max \left(\frac{512 V_{\mathcal{C}}}{\epsilon^{2}} \log \frac{256 V_{\mathcal{C}}}{\epsilon^{2}}, \frac{256}{\epsilon^{2}} \log \frac{8}{\delta}\right) .
$$

The corollary is a direct consequence of Theorem 12.6. The details are left as an exercise (Problem 12.5). The constants may be improved by using refined versions of Theorem 12.5 (see, e.g., Theorem 12.8). The sample size that guarantees the prescribed accuracy and confidence is proportional to the maximum of

$$
\frac{V_{\mathcal{C}}}{\epsilon^{2}} \log \frac{V_{\mathcal{C}}}{\epsilon^{2}} \quad \text { and } \quad \frac{1}{\epsilon^{2}} \log \frac{1}{\delta}
$$

Here we have our first practical interpretation of the vc dimension. Doubling the vC dimension requires that we basically double the sample size to obtain the same accuracy and confidence. Doubling the accuracy however, forces us to quadruple the sample size. On the other hand, the confidence level has little influence on the sample size, as it is hidden behind a logarithmic term, thanks to the exponential nature of the Vapnik-Chervonenkis inequality.

The Vapnik-Chervonenkis bound and the sample complexity $N(\epsilon, \delta)$ also allow us to compare different classes in a unified manner. For example, if we pick $\phi_{n}^{*}$ by minimizing the empirical error over all hyperrectangles of $\mathcal{R}^{18}$, will we need a sample size that exceeds that of the rule that minimizes the empirical error over all linear halfspaces of $\mathcal{R}^{29}$ ? With the sample complexity in hand, it is just a matter of comparing VC dimensions.

As a function of $\epsilon$, the above bound grows as $O\left(\left(1 / \epsilon^{2}\right) \log \left(1 / \epsilon^{2}\right)\right)$. It is possible, interestingly, to get rid of the "log" term, at the expense of increasing the linearity in the vc dimension (see Problem 12.11).

### 12.7 The Zero-Error Case

Theorem 12.5 is completely general, as it applies to any class of classifiers and all distributions. In some cases, however, when we have some additional information about the distribution, it is possible to obtain even better bounds. For example, in the theory of concept learning one commonly assumes that $L^{*}=0$, and that the Bayes decision is contained in $\mathcal{C}$ (see, e.g., Valiant (1984), Blumer, Ehrenfeucht, Haussler, and Warmuth (1989), Natarajan (1991)). The following theorem provides significant improvement. Its various forms have been proved by Devroye and Wagner (1976b), Vapnik (1982), and Blumer, Ehrenfeucht, Haussler, and Warmuth (1989). For a sharper result, see Problem 12.9.

Theorem 12.7. Let $\mathcal{C}$ be a class of decision functions mapping $\mathcal{R}^{d}$ to $\{0,1\}$, and let $\phi_{n}^{*}$ be a function in $\mathcal{C}$ that minimizes the empirical error based on the training sample $D_{n}$. Suppose that $\inf _{\phi \in \mathcal{C}} L(\phi)=0$, i.e., the Bayes decision is contained in $\mathcal{C}$, and $L^{*}=0$. Then

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)>\epsilon\right\} \leq 2 \mathcal{S}(\mathcal{C}, 2 n) 2^{-n \epsilon / 2} .
$$

To contrast this with Theorem 12.6, observe that the exponent in the upper bound for the empirically optimal rule is proportional to $-n \epsilon$ instead of $-n \epsilon^{2}$. To see the significance of this difference, note that Theorem 12.7 implies that the error probability of the selected classifier is within $O(\log n / n)$ of the optimal rule in the class (which equals zero in this case, see Problem 12.8), as opposed to $O(\sqrt{\log n / n})$ from Theorem 12.6. We show in Chapter 14 that this is not a technical coincidence, but since both bounds are essentially tight, it is a mathematical witness to the fact that it is remarkably easier to select a good classifier when $L^{*}=0$. The proof is
based on the random permutation argument developed in the original proof of the Vapnik-Chervonenkis inequality (1971).

Proof. For $n \epsilon \leq 2$, the inequality is clearly true. So, we assume that $n \epsilon>2$. First observe that since $\inf _{\phi \in \mathcal{C}} L(\phi)=0, \widehat{L}_{n}\left(\phi_{n}^{*}\right)=0$ with probability one. It is easily seen that

$$
L\left(\phi_{n}^{*}\right) \leq \sup _{\phi: \widehat{L}_{n}(\phi)=0}\left|L(\phi)-\widehat{L}_{n}(\phi)\right| .
$$

Now, we return to the notation of the previous sections, that is, $Z_{i}$ denotes the pair $\left(X_{i}, Y_{i}\right), v$ denotes its measure, and $\nu_{n}$ is the empirical measure based on $Z_{1}, \ldots, Z_{n}$. Also, $\mathcal{A}$ consists of all sets of the form $A=\{(x, y): \phi(x) \neq y\}$ for $\phi \in \mathcal{C}$. With these notations,

$$
\sup _{\phi: \widehat{L}_{n}(\phi)=0}\left|L(\phi)-\widehat{L}_{n}(\phi)\right|=\sup _{A: v_{n}(A)=0}\left|v_{n}(A)-v(A)\right| .
$$

Step 1. Symmetrization by a ghost sample. The first step of the proof is similar to that of Theorem 12.5. Introduce the auxiliary sample $Z_{1}^{\prime}, \ldots, Z_{n}^{\prime}$ such that the random variables $Z_{1}, \ldots, Z_{n}, Z_{1}^{\prime}, \ldots, Z_{n}^{\prime}$ are i.i.d., and let $v_{n}^{\prime}$ be the empirical measure for $Z_{1}^{\prime}, \ldots, Z_{n}^{\prime}$. Then for $n \epsilon>2$

$$
\mathbf{P}\left\{\sup _{A: v_{n}(A)=0}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 2 \mathbf{P}\left\{\sup _{A: v_{n}(A)=0}\left|v_{n}(A)-v_{n}^{\prime}(A)\right|>\frac{\epsilon}{2}\right\}
$$

The proof of this inequality parallels that of the corresponding one in the proof of Theorem 12.5. However, an important difference is that the condition $n \epsilon^{2}>2$ there is more restrictive than the condition $n \epsilon>2$ here. The details of the proof are left to the reader (see Problem 12.7).
Step 2. Symmetrization by permuting. Note that the distribution of

$$
\sup _{A: v_{n}(A)=0}\left|v_{n}(A)-v_{n}^{\prime}(A)\right|=\sup _{A: \sum_{i=1}^{n} I_{A}\left(Z_{i}\right)=0}\left|\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(Z_{i}\right)-\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(Z_{i}^{\prime}\right)\right|
$$

is the same as the distribution of

$$
\beta(\Pi) \stackrel{\text { def }}{=} \sup _{A: \sum_{i=1}^{n} I_{A}\left(\Pi\left(Z_{i}\right)\right)=0}\left|\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi\left(Z_{i}\right)\right)-\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi\left(Z_{i}^{\prime}\right)\right)\right|,
$$

where $\Pi\left(Z_{1}\right), \ldots, \Pi\left(Z_{n}\right), \Pi\left(Z_{1}^{\prime}\right), \ldots, \Pi\left(Z_{n}^{\prime}\right)$ is an arbitrary permutation of the random variables $Z_{1}, \ldots, Z_{n}, Z_{1}^{\prime}, \ldots, Z_{n}^{\prime}$. The ( $2 n$ )! possible permutations are denoted by

$$
\Pi_{1}, \Pi_{2}, \ldots, \Pi_{(2 n)!}
$$

Therefore,

$$
\begin{gathered}
\mathbf{P}\left\{\sup _{A: v_{n}(A)=0}\left|v_{n}(A)-v_{n}^{\prime}(A)\right|>\frac{\epsilon}{2}\right\} \\
=\mathbf{E}\left\{\frac{1}{(2 n)!} \sum_{j=1}^{(2 n)!} I_{\left\{\beta\left(\Pi_{j}\right)>\frac{\epsilon}{2}\right\}}\right\}
\end{gathered}
$$

$$
\begin{aligned}
=\mathbf{E}\left\{\frac{1}{(2 n)!}\right. & \sum_{j=1}^{(2 n)!} \\
& \left.\sup _{A: \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(Z_{i}\right)\right)=0} I_{\left\{\left|\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(Z_{i}\right)\right)-\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(Z_{i}^{\prime}\right)\right)\right|>\frac{\epsilon}{2}\right\}}\right\} .
\end{aligned}
$$

Step 3. Conditioning. Next we fix $z_{1}, \ldots, z_{n}, z_{1}^{\prime}, \ldots, z_{n}^{\prime}$ and bound the value of the random variable above. Let $\overline{\mathcal{A}} \subset \mathcal{A}$ be a collection of sets such that any two sets in $\overline{\mathcal{A}}$ pick different subsets of $\left\{z_{1}, \ldots, z_{n}, z_{1}^{\prime}, \ldots, z_{n}^{\prime}\right\}$, and its cardinality is $N_{\mathcal{A}}\left(z_{1}, \ldots, z_{n}, z_{1}^{\prime}, \ldots, z_{n}^{\prime}\right)$, that is, all possible subsets are represented exactly once. Then it suffices to take the supremum over $\overline{\mathcal{A}}$ instead of $\mathcal{A}$ :

$$
\begin{aligned}
& \frac{1}{(2 n)!} \sum_{j=1}^{(2 n)!} \sup _{A: \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}\right)\right)=0} I_{\left\{\left.1 \frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}\right)\right)-\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}^{\prime}\right)\right) \right\rvert\,>\frac{\epsilon}{2}\right\}} \\
& \quad=\frac{1}{(2 n)!} \sum_{j=1}^{(2 n)!} \sup _{A \in \overline{\mathcal{A}}: \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}\right)\right)=0} I_{\left\{\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}^{\prime}\right)\right)>\frac{\epsilon}{2}\right\}} \\
& \quad \leq \frac{1}{(2 n)!} \sum_{j=1}^{(2 n)!} \sum_{A \in \overline{\mathcal{A}}: \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}\right)\right)=0} I_{\left\{\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}^{\prime}\right)\right)>\frac{\epsilon}{2}\right\}} \\
& \quad=\frac{1}{(2 n)!} \sum_{j=1}^{(2 n)!} \sum_{A \in \overline{\mathcal{A}}} I_{\left\{\sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}\right)\right)=0\right\}} I_{\left\{\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}^{\prime}\right)\right)>\frac{\epsilon}{2}\right\}} \\
& \quad=\sum_{A \in \overline{\mathcal{A}}} \frac{1}{(2 n)!} \sum_{j=1}^{(2 n)!} I_{\left\{\sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}\right)\right)=0\right\}} I_{\left\{\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}^{\prime}\right)\right)>\frac{\epsilon}{2}\right\}} .
\end{aligned}
$$

Step 4. Counting. Clearly, the expression behind the first summation sign is just the number of permutations of the $2 n$ points $z_{1}, z_{1}^{\prime}, \ldots, z_{n}, z_{n}^{\prime}$ with the property

$$
I_{\left\{\sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}\right)\right)=0\right\}} I_{\left\{\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}^{\prime}\right)\right)>\frac{\epsilon}{2}\right\}}=1,
$$

divided by $(2 n)$ !, the total number of permutations. Observe that if $l=\sum_{i=1}^{n}\left(I_{A}\left(z_{i}\right)+\right.$ $\left.I_{A}\left(z_{i}^{\prime}\right)\right)$ is the total number of points in $A$ among $z_{1}, z_{1}^{\prime}, \ldots, z_{n}, z_{n}^{\prime}$ for a fixed set $A$, then the number of permutations such that

$$
I_{\left\{\sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}\right)\right)=0\right\}} I_{\left\{\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(z_{i}^{\prime}\right)\right)>\frac{\epsilon}{2}\right\}}=1
$$

is zero if $l \leq n \epsilon / 2$. If $l>n \epsilon / 2$, then the fraction of the number of permutations with the above property and the number of all permutations can not exceed

$$
\frac{\left(e_{1}^{2}\right)}{\left(a_{n}^{2}\right)}
$$

To see this, note that for the above product of indicators to be 1 , all the points falling in $A$ have to be in the second half of the permuted sample. Now clearly,

$$
\frac{\binom{n}{l}}{\binom{n}{l}}=\frac{n(n-1) \cdots(n-l+1)}{2 n(2 n-1) \cdots(2 n-l+1)} \leq 2^{-l} \leq 2^{-n \epsilon / 2} .
$$

Summarizing, we have

$$
\begin{aligned}
& \mathbf{P}\left\{\sup _{A: v_{n}(A)=0}\left|v_{n}(A)-v_{n}^{\prime}(A)\right|>\frac{\epsilon}{2}\right\} \\
& \leq \mathbf{E}\left\{\frac{\{A \in \mathbb{A}}{(2 n)!} \sum_{j=1}^{(2 n)!}\right. \\
& \left.\quad \leq \sup _{A: \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(Z_{i}\right)\right)=0} I_{\left\{\left|\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(Z_{i}\right)\right)-\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(Z_{i}^{\prime}\right)\right)\right|>\frac{\epsilon}{2}\right\}}\right\} \\
& \leq \mathbf{E}\left\{\sum_{A \in \overline{\mathcal{A}}} 2^{-n \epsilon / 2}\right\} \\
& \quad=\mathbf{E}\left\{N_{\mathcal{A}}\left(Z_{1}, \ldots, Z_{n}, Z_{1}^{\prime}, \ldots, Z_{n}^{\prime}\right) 2^{-n \epsilon / 2}\right\} \\
& \leq s(\mathcal{A}, 2 n) 2^{-n \epsilon / 2},
\end{aligned}
$$

and the theorem is proved.

Again, we can bound the sample complexity $N(\epsilon, \delta)$ restricted to the class of distributions with $\inf _{\phi \in \mathcal{C}} L(\phi)=0$. Just as Theorem 12.6 implies Corollary 12.3, Theorem 12.7 yields

Corollary 12.4. The sample complexity $N(\epsilon, \delta)$ that guarantees

$$
\sup _{(X, Y): \sin _{\phi \in C} L(\phi)=0} \mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq \delta
$$

for $n \geq N(\epsilon, \delta)$, is bounded by

$$
N(\epsilon, \delta) \leq \max \left(\frac{8 V_{\mathcal{C}}}{\epsilon} \log _{2} \frac{13}{\epsilon}, \frac{4}{\epsilon} \log _{2} \frac{2}{\delta}\right)
$$

A quick comparison with Corollary 12.3 shows that the $\epsilon^{2}$ factors in the denominators there are now replaced by $\epsilon$. For the same accuracy, much smaller samples suffice if we know that $\inf _{\phi \in \mathcal{C}} L(\phi)=0$. Interestingly, the sample complexity is still roughly linear in the vC dimension.
Remark. We also note that under the condition of Theorem 12.7,

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)>\epsilon\right\} \leq 2 \mathbf{E}\left\{N_{\mathcal{A}}\left(Z_{1}, \ldots, Z_{n}, Z_{1}^{\prime}, \ldots, Z_{n}^{\prime}\right)\right\} 2^{-n \epsilon / 2}
$$

$$
=2 \mathbf{E}\left\{N_{\mathcal{A}}\left(Z_{1}, \ldots, Z_{2 n}\right)\right\} 2^{-n \epsilon / 2}
$$

where the class $\mathcal{A}$ of sets is defined in the proof of Theorem 12.7.

Remark. As Theorem 12.7 shows, the difference between the error probability of an empirically optimal classifier and that of the optimal in the class is much smaller if the latter quantity is known to be zero than if no restriction is imposed on the distribution. To bridge the gap between the two bounds, one may put the restriction $\inf _{\phi \in \mathcal{C}} L(\phi) \leq L$ on the distribution, where $L \in(0,1 / 2)$ is a fixed number. Devroye and Wagner (1976b) and Vapnik (1982) obtained such bounds. For example, it follows from a result by Vapnik (1982) that

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq 8 \mathcal{S}(\mathcal{C}, n) e^{-n \epsilon^{2} / 8(L+\epsilon)}
$$

As expected, the bound becomes smaller as $L$ decreases. We face the same phenomenon in Chapter 14, where lower bounds are obtained for the probability above.

### 12.8 Extensions

We mentioned earlier that the constant in the exponent in Theorem 12.5 can be improved at the expense of a more complicated argument. The best possible exponent appears in the following result, whose proof is left as an exercise (Problem 12.15):

Theorem 12.8. (Devroye (1982A)).

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq \operatorname{cs}\left(\mathcal{A}, n^{2}\right) e^{-2 n \epsilon^{2}},
$$

where the constant $c$ does not exceed $4 e^{4 \epsilon+4 \epsilon^{2}} \leq 4 e^{8}, \epsilon \leq 1$.
Even though the coefficient in front is larger than in Theorem 12.5, it becomes very quickly absorbed by the exponential term. We will see in Chapter 13 that for $V_{\mathcal{A}}>2, s(\mathcal{A}, n) \leq n^{V_{\mathcal{A}}}$, so $s\left(\mathcal{A}, n^{2}\right) \leq n^{2 V_{\mathcal{A}}}$. This difference is negligible compared to the difference between the exponential terms, even for moderately large values of $n \epsilon^{2}$.

Both Theorem 12.5 and Theorem 12.8 imply that

$$
\mathbf{E}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|\right\}=O(\sqrt{\log n / n})
$$

(see Problem 12.1). However, it is possible to get rid of the logarithmic term to obtain $O(1 / \sqrt{n})$. For example, for the Kolmogorov-Smirnov statistic we have the following result by Dvoretzky, Kiefer and Wolfowitz (1956), sharpened by Massart (1990):

Theorem 12.9. (Dvoretzky, Kiefer, and Wolfowitz (1956); Massart (1990)). Using the notation of Theorem 12.4, we have for every $n$ and $\epsilon>0$,

$$
\mathbf{P}\left\{\sup _{z \in \mathcal{R}}\left|F(z)-F_{n}(z)\right|>\epsilon\right\} \leq 2 e^{-2 n \epsilon^{2}} .
$$

For the general case, we also have Alexander's bound:
Theorem 12.10. (Alexander (1984)). For $n \epsilon^{2} \geq 64$,

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 16(\sqrt{n} \epsilon)^{4096 V_{\mathcal{A}}} e^{-2 n \epsilon^{2}}
$$

The theorem implies the following (Problem 12.10):
Corollary 12.5 .

$$
\mathbf{E}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|\right\} \leq \frac{8+\sqrt{2048 V_{\mathcal{A}} \log \left(4096 V_{\mathcal{A}}\right)}}{\sqrt{n}}
$$

The bound in Theorem 12.10 is theoretically interesting, since it implies (see Problem 12.10) that for fixed $V_{\mathcal{A}}$ the expected value of the supremum decreases as $a / \sqrt{n}$ instead of $\sqrt{\log n / n}$. However, a quick comparison reveals that Alexander's bound is larger than that of Theorem 12.8, unless $n>2^{6144}$, an astronomically large value. Recently, Talagrand (1994) obtained a very strong result. He proved that there exists a universal constant $c$ such that

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq \frac{c}{\sqrt{n} \epsilon}\left(\frac{c n \epsilon^{2}}{V_{\mathcal{A}}}\right)^{V_{\mathcal{A}}} e^{-2 n \epsilon^{2}}
$$

For more information about these inequalities, see also Vapnik (1982), Gaenssler (1983), Gaenssler and Stute (1979), and Massart (1983).

It is only natural to ask whether the uniform law of large numbers

$$
\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right| \rightarrow 0
$$

holds if we allow $\mathcal{A}$ to be the class of all measurable subsets of $\mathcal{R}^{d}$. In this case the supremum above is called the total variation between the measures $v_{n}$ and $\nu$. The convergence clearly can not hold if $v_{n}$ is the standard empirical measure

$$
v_{n}(A)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{Z_{i} \in A\right\}} .
$$

But is there another empirical measure such that the convergence holds? The somewhat amusing answer is no. As Devroye and Györfi (1992) proved, for any empirical measure $\tilde{v}_{n}$-that is, a function depending on $Z_{1}, \ldots, Z_{n}$ assigning a nonnegative number to any measurable set-there exists a distribution of $Z$ such that for all n

$$
\sup _{A \in \mathcal{A}}\left|\tilde{v}_{n}(A)-v(A)\right|>1 / 4
$$

almost surely. Thus, in this generality, the problem is hopeless. For meaningful results, either $\mathcal{A}$ or $v$ must be restricted. For example, if we assume that $v$ is absolutely continuous with density $f$, and that $\tilde{v}_{n}$ is absolutely continuous too (with density $f_{n}$ ), then by Scheffés theorem (1947),

$$
\sup _{A \in \mathcal{A}}\left|\tilde{v}_{n}(A)-v(A)\right|=\frac{1}{2} \int_{\mathcal{R}^{d}}\left|f_{n}(x)-f(x)\right| d x
$$

(see Problem 12.13). But as we see from Problems 6.2, 10.2, 9.6, and 10.3, there exist density estimators (such as histogram and kernel estimates) such that the $L_{1^{-}}$ error converges to zero almost surely for all possible densities. Therefore, the total variation between the empirical measures derived from these density estimates and the true measure converges to zero almost surely for all distributions with a density. For other large classes of distributions that can be estimated consistently in total variation, we refer to Barron, Györfi, and van der Meulen (1992).

## Problems and Exercises

Problem 12.1. Prove that if a nonnegative random variable $Z$ satisfies $\mathbf{P}\{Z>t\} \leq c e^{-2 n t^{2}}$ for all $t>0$ and some $c>0$, then,

$$
\mathbf{E}\left\{Z^{2}\right\} \leq \frac{\log (c e)}{2 n}
$$

Furthermore,

$$
\mathbf{E} Z \leq \sqrt{\mathbf{E}\left\{Z^{2}\right\}} \leq \sqrt{\frac{\log (c e)}{2 n}}
$$

Hint: Use the identity $\mathbf{E}\left\{Z^{2}\right\}=\int_{0}^{\infty} \mathbf{P}\left\{Z^{2}>t\right\} d t$, and set $\int_{0}^{\infty}=\int_{0}^{u}+\int_{u}^{\infty}$. Bound the first integral by $u$, and the second by the exponential inequality. Find the value of $u$ that minimizes the obtained upper bound.

Problem 12.2. Generalize the arguments of Theorems 4.5 and 4.6 to prove Theorems 12.2 and 12.3.

Problem 12.3. Determine the fingering dimension of classes of classifiers $\mathcal{C}=\{\phi: \phi(x)=$ $\left.I_{\{x \in A\}} ; A \in \mathcal{A}\right\}$ if the class $\mathcal{A}$ is
(1) the class of all closed intervals in $\mathcal{R}$,
(2) the class of all sets obtained as the union of $m$ closed intervals in $\mathcal{R}$,
(3) the class of balls in $\mathcal{R}^{d}$ centered at the origin,
(4) the class of all balls in $\mathcal{R}^{d}$,
(5) the class of sets of the form $\left(-\infty, x_{1}\right] \times \cdots \times\left(-\infty, x_{d}\right]$ in $\mathcal{R}^{d}$, and
(6) the class of all convex polygons in $\mathcal{R}^{2}$.

Problem 12.4. Let $\mathcal{C}$ be a class of classifiers with fingering dimension $k>4$ (independently of $n$ ). Show that $V_{\mathcal{C}} \leq k \log _{2}^{2} k$.

Problem 12.5. Prove that Theorem 12.6 implies Corollary 12.3. Hint: Find $N(\epsilon, \delta)$ such that $8 n^{V_{c}} e^{-n \epsilon^{2} / 128} \leq \delta$ whenever $n>N(\epsilon, \delta)$. To see this, first show that $n^{V_{c}} \leq e^{n \epsilon^{2} / 256}$ is satisfied for $n \geq \frac{51 \overline{2} V_{C}}{\epsilon^{2}} \log \frac{256 V_{C}}{\epsilon^{2}}$, which follows from the fact that $2 \log x \leq x$ if $x \geq e^{2}$. But in this case $8 n^{V_{c}} e^{-n \epsilon^{2} / 128} \leq 8 e^{-n \epsilon^{2} / 256}$. The upper bound does not exceed $\delta$ if $n \geq \frac{256}{\epsilon^{2}} \log \frac{8}{\delta}$.

Problem 12.6. Let $\mathcal{C}$ be a class of classifiers $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$, and let $\phi_{n}^{*}$ be a classifier minimizing the empirical error probability measured on $D_{n}$. Assume that we have an algorithm which selects a classifier $g_{n}$ such that

$$
\mathbf{P}\left\{\widehat{L}_{n}\left(g_{n}\right) \leq \inf _{\phi \in \mathcal{C}} \widehat{L}_{n}(\phi)+\epsilon_{n}\right\} \geq 1-\delta_{n},
$$

where $\left\{\epsilon_{n}\right\}$ and $\left\{\delta_{n}\right\}$ are sequences of positive numbers converging to zero. Show that

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq \delta_{n}+\mathbf{P}\left\{2 \sup _{\phi \in \mathcal{C}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|>\epsilon-\epsilon_{n}\right\}
$$

Find conditions on $\left\{\epsilon_{n}\right\}$ and $\left\{\delta_{n}\right\}$ so that

$$
\mathbb{E} L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)=O\left(\sqrt{\frac{\log n}{n}}\right)
$$

that is, $\mathbf{E} L\left(g_{n}\right)$ converges to the optimum at the same order as the error probability of $\phi_{n}^{*}$.
Problem 12.7. Prove that

$$
\mathbf{P}\left\{\sup _{A: v_{n}(A)=0}\left|v_{n}(A)-v(A)\right|>\epsilon\right\} \leq 2 \mathbf{P}\left\{\sup _{A: v_{n}(A)=0}\left|v_{n}(A)-v_{n}^{\prime}(A)\right|>\frac{\epsilon}{2}\right\}
$$

holds if $n \epsilon>2$. This inequality is needed to complete the proof of Theorem 12.7. Hint: Proceed as in the proof of Theorem 12.5. Introduce $A^{*}$ with $v_{n}\left(A^{*}\right)=0$ and justify the validity of the steps of the following chain of inequalities:

$$
\begin{aligned}
& \mathbf{P}\left\{\sup _{A: v_{n}(A)=0}\left|v_{n}(A)-v_{n}^{\prime}(A)\right|>\epsilon / 2\right\} \\
& \geq \mathbf{E}\left\{I_{\left\{v\left(A^{*}\right)>\epsilon\right\}} \mathbf{P}\left\{\left.v_{n}^{\prime}\left(A^{*}\right) \geq \frac{\epsilon}{2} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\}\right\} \\
& \quad \geq \mathbf{P}\left\{B(n, \epsilon)>\frac{n \epsilon}{2}\right\} \mathbf{P}\left\{\left|v_{n}\left(A^{*}\right)-v\left(A^{*}\right)\right|>\epsilon\right\},
\end{aligned}
$$

where $B(n, \epsilon)$ is a binomial random variable with parameters $n$ and $\epsilon$. Finish the proof by showing that the probability on the right-hand side is greater than or equal to $1 / 2$ if $n \epsilon>2$. (Under the slightly more restrictive condition $n \epsilon>8$, this follows from Chebyshev's inequality.)

Problem 12.8. Prove that Theorem 12.7 implies that if $\inf _{\phi \in \mathcal{C}} L(\phi)=0$, then

$$
\mathbf{E} L\left(\phi_{n}^{*}\right) \leq \frac{2 V_{\mathcal{C}} \log (2 n)+4}{n \log 2} .
$$

We note here that Haussler, Littlestone, and Warmuth (1988) demonstrated the existence of a classifier $\phi_{n}^{*}$ with $\mathbf{E} L\left(\phi_{n}^{*}\right) \leq 2 V_{\mathcal{C}} / n$ when $\inf _{\phi \in \mathcal{C}} L(\phi)=0$. Hint: Use the identity $\mathbf{E} X=\int_{0}^{\infty} \mathbf{P}\{X>t\} d t$ for nonnegative random variables $X$, and employ the fact that $\mathcal{S}(\mathcal{C}, n) \leq n^{V_{C}}$ (see Theorem 13.3).

Problem 12.9. Prove the following version of Theorem 12.7. Let $\phi_{n}^{*}$ be a function that minimizes the empirical error over a class in $\mathcal{C}$. Assume that $\inf _{\phi \in \mathcal{C}} L(\phi)=0$. Then

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)>\epsilon\right\} \leq 2 \mathcal{S}\left(\mathcal{C}, n^{2}\right) e^{-n \epsilon}
$$

(Shawe-Taylor, Anthony, and Biggs (1993)). Hint: Modify the proof of Theorem 12.7 by introducing a ghost sample $Z_{1}^{\prime}, \ldots, Z_{m}^{\prime}$ with size $m$ (to be specified after optimization). Only the first symmetrization step (Problem 12.7) needs adjusting: show that for any $\alpha \in(0,1)$,

$$
\mathbf{P}\left\{\sup _{A: v_{n}(A)-0} v(A)>\epsilon\right\} \leq \frac{1}{1-e^{-m \alpha \epsilon}} \mathbf{P}\left\{\sup _{A: v_{n}(A)=0} v_{m}^{\prime}(A)>(1-\alpha) \epsilon\right\}
$$

where $\nu_{m}^{\prime}$ is the empirical measure based on the ghost sample (use Bernstein's inequalityTheorem 8.2). The rest of the proof is similar to that of Theorem 12.7. Choose $m=n^{2}-n$ and $\alpha=n /(n+m)$.

Problem 12.10. Prove that Alexander's bound (Theorem 12.10) implies that if $\mathcal{A}$ is a Vapnik-Chervonenkis class with vc dimension $V=V_{\mathcal{A}}$, then

$$
\mathbf{E}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|\right\} \leq \frac{8+\sqrt{2048 V \log (4096 V)}}{\sqrt{n}}
$$

Hint: Justify the following steps:
(1) If $\psi$ is a negative decreasing concave function, then

$$
\int_{u}^{\infty} e^{\psi(t)} d t \leq \frac{e^{\psi(u)}}{-\psi^{\prime}(u)}
$$

Hint: Bound $\psi$ by using its Taylor series expansion.
(2) Let $b>0$ be fixed. Then for $u \geq \sqrt{b / 2}$,

$$
\int_{u}^{\infty} t^{b} e^{-2 t^{2}} d t \leq \frac{u^{b-1} e^{-2 u^{2}}}{2}
$$

Hint: Use the previous step.
(3) Let $X$ be a positive random variable for which

$$
\mathbf{P}\{X>u\} \leq a u^{b} e^{-2 u^{2}}, u \geq \sqrt{c}
$$

where $a, b, c$ are positive constants. Then, if $b \geq 2 c \geq e$,

$$
\mathbf{E} X \leq \frac{a}{2}+\sqrt{\frac{b \log b}{2}}
$$

Hint: Use $\mathbf{E} X=\int_{0}^{\infty} \mathbf{P}\{X>u\} d u$, and bound the probability either by one, or by the bound of the previous step.

Problem 12.11. Use Alexander's inequality to obtain the following sample size bound for empirical error minimization:

$$
N(\epsilon, \delta) \leq \max \left(\frac{2^{14} V_{\mathcal{C}} \log \left(2^{14} V_{\mathcal{C}}\right)}{\epsilon^{2}}, \frac{4}{\epsilon^{2}} \log \frac{16}{\delta}\right)
$$

For what values of $\epsilon$ does this bound beat Corollary 12.3?
Problem 12.12. Let $Z_{1}, \ldots, Z_{n}$ be i.i.d. random variables in $\mathcal{R}^{d}$, with measure $v$, and standard empirical measure $v_{n}$. Let $\mathcal{A}$ be an arbitrary class of subsets of $\mathcal{R}^{d}$. Show that as $n \rightarrow \infty$,

$$
\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right| \rightarrow 0 \text { in probability }
$$

if and only if

$$
\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right| \rightarrow 0 \quad \text { with probability one. }
$$

Hint: Use McDiarmid's inequality.
Problem 12.13. Prove Scheffé's theorem (1947): let $\mu$ and $v$ be absolute continuous probability measures on $\mathcal{R}^{d}$ with densities $f$ and $g$, respectively. Prove that

$$
\sup _{A \in \mathcal{A}}|\mu(A)-v(A)|=\frac{1}{2} \int|f(x)-g(x)| d x
$$

where $\mathcal{A}$ is the class of all Borel-measurable sets. Hint: Show that the supremum is achieved for the set $\{x: f(x)>g(x)\}$.

Problem 12.14. Learning based on empirical covering. This problem demonstrates an alternative method of picking a classifier which works as well as empirical error minimization. The method, based on empirical covering of the class of classifiers, was introduced by Buescher and Kumar (1996a). The idea of covering the class goes back to Vapnik (1982). See also Benedek and Itai (1988), Kulkarni (1991), and Dudley, Kulkarni, Richardson, and Zeitouni (1994). Let $\mathcal{C}$ be a class of classifiers $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$. The data set $D_{n}$ is split into two parts, $D_{m}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{m}, Y_{m}\right)\right)$, and $T_{l}=\left(\left(X_{m+1}, Y_{m+1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$, where $n=m+l$. We use the first part $D_{m}$ to cover $\mathcal{C}$ as follows. Define the random variable $N$ as the number of different values the binary vector $\mathbf{b}_{m}(\phi)=\left(\phi\left(X_{1}\right), \ldots, \phi\left(X_{m}\right)\right)$ takes as $\phi$ is varied over $\mathcal{C}$. Clearly, $N \leq \mathcal{S}(\mathcal{C}, m)$. Take $N$ classifiers from $\mathcal{C}$, such that all $N$ possible values of the binary vector $\mathbf{b}_{m}(\phi)$ are represented exactly once. Denote these classifiers by $\tilde{\phi}_{1}, \ldots, \tilde{\phi}_{N}$. Among these functions, pick one that minimizes the empirical error on the second part of the data set $T_{l}$ :

$$
\widehat{L}_{l}\left(\tilde{\phi}_{i}\right)=\frac{1}{l} \sum_{j=1}^{l} I_{\left\{\tilde{\phi}_{i}\left(X_{m+j}\right) \neq Y_{m+j}\right\}}
$$

Denote the selected classifier by $\widehat{\phi}_{n}$. Show that for every $n, m$ and $\epsilon>0$, the difference between the error probability $L\left(\widehat{\phi}_{n}\right)=\mathbf{P}\left\{\widehat{\phi}_{n}(X) \neq Y \mid D_{n}\right\}$ and the minimal error probability in the class satisfies

$$
\mathbf{P}\left\{L\left(\widehat{\phi}_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq 2 \mathcal{S}(\mathcal{C}, m) e^{-(n-m) \epsilon^{2} / 8}+2 S^{4}(\mathcal{C}, 2 m) e^{-m \epsilon \log 2 / 4}
$$

(Buescher and Kumar (1996a)). For example, by taking $m \sim \sqrt{n}$, we get

$$
\mathbf{E}\left\{L\left(\widehat{\phi}_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)\right\} \leq c \sqrt{\frac{V_{\mathcal{C}} \log n}{n}}
$$

where $c$ is a universal constant. The fact that the number of samples $m$ used for covering $\mathcal{C}$ is very small compared to $n$, may make the algorithm computationally more attractive than the method of empirical error minimization. Hint: Use the decomposition

$$
\begin{aligned}
& \mathbf{P}\left\{L\left(\widehat{\phi}_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \\
& \quad \leq \mathbf{P}\left\{L\left(\widehat{\phi}_{n}\right)-\inf _{i=1, \ldots, N} L\left(\tilde{\phi}_{i}\right)>\frac{\epsilon}{2}\right\}+\mathbf{P}\left\{\inf _{i=1, \ldots, N} L\left(\tilde{\phi}_{i}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\frac{\epsilon}{2}\right\}
\end{aligned}
$$

Bound the first term on the right-hand side by using Lemma 8.2 and Hoeffding's inequality:

$$
\mathbf{P}\left\{L\left(\widehat{\phi}_{n}\right)-\inf _{i=1, \ldots, N} L\left(\tilde{\phi}_{i}\right)>\frac{\epsilon}{2}\right\} \leq 2 \mathcal{S}(\mathcal{C}, m) e^{-l \epsilon^{2} / 8}
$$

To bound the second term of the decomposition, observe that

$$
\begin{aligned}
\inf _{i=1, \ldots, N} L\left(\tilde{\phi}_{i}\right)-\inf _{\phi \in \mathcal{C}} L(\phi) & \leq \sup _{\phi, \phi^{\prime} \in \mathcal{C}: \mathbf{b}_{m}(\phi)=b_{m}\left(\phi^{\prime}\right)}\left|L(\phi)-L\left(\phi^{\prime}\right)\right| \\
& \leq \sup _{\phi, \phi^{\prime} \in \mathcal{C}: \mathbf{b}_{m}(\phi)=b_{m}\left(\phi^{\prime}\right)} \mathbf{P}\left\{\phi(X) \neq \phi^{\prime}(X)\right\} \\
& =\sup _{A \in \mathcal{A}: v_{m}(A)=0} v(A),
\end{aligned}
$$

where

$$
\mathcal{A}=\left\{\{x: \phi(x)=1\}: \phi(x)=\left|\phi_{1}(x)-\phi_{2}(x)\right|, \phi_{1}, \phi_{2} \in \mathcal{C}\right\},
$$

and $v_{m}(A)=\frac{1}{m} \sum_{i=1}^{m} I_{\left\{X_{i} \in A\right\}}$. Bound the latter quantity by applying Theorem 12.7. To do this, you will need to bound the shatter coefficients $s(\mathcal{A}, 2 m)$. In Chapter 13 we introduce simple tools for this. For example, it is easy to deduce from parts (ii), (iii), and (iv) of Theorem 13.5, that $s(\mathcal{A}, 2 m) \leq S^{4}(\mathcal{C}, 2 m)$.

Problem 12.15. Prove that for all $\epsilon \in(0,1)$,

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|\nu_{n}(A)-v(A)\right|>\epsilon\right\} \leq c s\left(\mathcal{A}, n^{2}\right) e^{-2 n \epsilon^{2}}
$$

where $c \leq 4 e^{4 \epsilon+4 \epsilon^{2}}$ (Devroye (1982a)). Hint: Proceed as indicated by the following steps:
(1) Introduce an i.i.d. ghost sample $Z_{1}^{\prime}, \ldots, Z_{m}^{\prime}$ of size $m=n^{2}-n$, where $Z_{1}^{\prime}$ is distributed as $Z_{1}$. Denote the corresponding empirical measure by $\nu_{m}^{\prime}$. As in the proof of the first step of Theorem 12.4, prove that for $\alpha, \epsilon \in(0,1)$,

$$
\begin{aligned}
& \mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v_{m}^{\prime}(A)\right|>(1-\alpha) \epsilon\right\} \\
& \quad \geq\left(1-\frac{1}{4 \alpha^{2} \epsilon^{2} m}\right) \mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-v(A)\right|>\epsilon\right\}
\end{aligned}
$$

(2) Introduce $n^{2}$ ! permutations $\Pi_{1}, \ldots, \Pi_{n+m}$ of the $n+m$ random variables as in Step 2 of the proof of Theorem 12.5. Show that

$$
\begin{aligned}
& \frac{1}{n^{2}!} \sum_{j=1}^{n^{2}!} \sup _{A \in \mathcal{A}} I_{\left\{\left|\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(Z_{i}\right)\right)-\frac{1}{m} \sum_{i=1}^{m} I_{A}\left(\Pi_{j}\left(Z_{i}^{\prime}\right)\right)\right|>(1-\alpha) \epsilon\right\}} \\
& \quad \leq \mathcal{S}\left(\mathcal{A}, n^{2}\right) \max _{A \in \mathcal{A}} \frac{1}{n^{2}!} \sum_{j=1}^{n^{2}!} I_{\left\{\left|\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(Z_{i}\right)\right)-\frac{1}{m} \sum_{i=1}^{m} I_{A}\left(\Pi_{j}\left(Z_{i}^{\prime}\right)\right)\right|>(1-\alpha) \epsilon\right\}} .
\end{aligned}
$$

(3) Show that for each $A \in \mathcal{A}$,

$$
\frac{1}{n^{2}!} \sum_{j=1}^{n^{2}!} I_{\left\{\left|\frac{1}{n} \sum_{i=1}^{n} I_{A}\left(\Pi_{j}\left(Z_{i}\right)\right)-\frac{1}{m} \sum_{i=1}^{m} I_{A}\left(\Pi_{j}\left(Z_{i}^{\prime}\right)\right)\right|>(1-\alpha) \epsilon\right\}} \leq 2 e^{-2 n \epsilon^{2}+4 \alpha n \epsilon^{2}+4 \epsilon^{2}}
$$

by using Hoeffding's inequality for sampling without replacement from $n^{2}$ binaryvalued elements (see Theorem A.25). Choose $\alpha=1 /(n \epsilon)$.

## 13

## Combinatorial Aspects of Vapnik-Chervonenkis Theory

### 13.1 Shatter Coefficients and VC Dimension

In this section we list a few interesting properties of shatter coefficients $s(\mathcal{A}, n)$ and of the vc dimension $V_{\mathcal{A}}$ of a class of sets $\mathcal{A}$. We begin with a property that makes things easier. In Chapter 12 we noted the importance of classes of the form

$$
\overline{\mathcal{A}}=\left\{A \times\{0\} \cup A^{c} \times\{1\} ; A \in \mathcal{A}\right\} .
$$

(The sets $A$ are of the form $\{x: \phi(x)=1\}$, and the sets in $\overline{\mathcal{A}}$ are sets of pairs $(x, y)$ for which $\phi(x) \neq y$.) Recall that if $\mathcal{C}$ is a class of classifiers $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$, then by definition, $\mathcal{S}(\mathcal{C}, n)=s(\overline{\mathcal{A}}, n)$ and $V_{\mathcal{C}}=V_{\overline{\mathcal{A}}}$. The first result states that $\mathcal{S}(\mathcal{C}, n)=s(\mathcal{A}, n)$, so it suffices to investigate properties of $\mathcal{A}$, a class of subsets of $\mathcal{R}^{d}$.

Theorem 13.1. For every $n$ we have $s(\overline{\mathcal{A}}, n)=s(\mathcal{A}, n)$, and therefore $V_{\overline{\mathcal{A}}}=V_{\mathcal{A}}$.
Proof. Let $N$ be a positive integer. We show that for any $n$ pairs from $\mathcal{R}^{d} \times$ $\{0,1\}$, if $N$ sets from $\overline{\mathcal{A}}$ pick $N$ different subsets of the $n$ pairs, then there are $N$ corresponding sets in $\mathcal{A}$ that pick $N$ different subsets of $n$ points in $\mathcal{R}^{d}$, and vice versa. Fix $n$ pairs $\left(x_{1}, 0\right), \ldots,\left(x_{m}, 0\right),\left(x_{m+1}, 1\right), \ldots,\left(x_{n}, 1\right)$. Note that since ordering does not matter, we may arrange any $n$ pairs in this manner. Assume that for a certain set $A \in \mathcal{A}$, the corresponding set $\bar{A}=A \times\{0\} \bigcup A^{c} \times\{1\} \in \overline{\mathcal{A}}$ picks out the pairs $\left(x_{1}, 0\right), \ldots,\left(x_{k}, 0\right),\left(x_{m+1}, 1\right), \ldots,\left(x_{m+l}, 1\right)$, that is, the set of these pairs is the intersection of $\overline{\mathcal{A}}$ and the $n$ pairs. Again, we can assume without loss of generality that the pairs are ordered in this way. This means that $A$ picks from the set
$\left\{x_{1}, \ldots, x_{n}\right\}$ the subset $\left\{x_{1}, \ldots, x_{k}, x_{m+l+1}, \ldots, x_{n}\right\}$, and the two subsets uniquely determine each other. This proves $s(\overline{\mathcal{A}}, n) \leq s(\mathcal{A}, n)$. To prove the other direction, notice that if $A$ picks a subset of $k$ points $x_{1}, \ldots, x_{k}$, then the corresponding set $\bar{A} \in \overline{\mathcal{A}}$ picks the pairs with the same indices from $\left\{\left(x_{1}, 0\right), \ldots,\left(x_{k}, 0\right)\right\}$. Equality of the VC dimensions follows from the equality of the shatter coefficients.

The following theorem, attributed to Vapnik and Chervonenkis (1971) and Sauer (1972), describes the relationship between the vC dimension and shatter coefficients of a class of sets. This is the most important tool for obtaining useful upper bounds on the shatter coefficients in terms of the VC dimension.

Theorem 13.2. If $\mathcal{A}$ is a class of sets with vC dimension $V_{\mathcal{A}}$, then for every $n$

$$
s(\mathcal{A}, n) \leq \sum_{i=0}^{V_{\mathcal{A}}}\binom{n}{i}
$$

Proof. Recall the definition of the shatter coefficients

$$
s(\mathcal{A}, n)=\max _{\left(x_{1}, \ldots, x_{n}\right)} N_{\mathcal{A}}\left(x_{1}, \ldots, x_{n}\right)
$$

where

$$
N_{\mathcal{A}}\left(x_{1}, \ldots, x_{n}\right)=\left|\left\{\left\{x_{1}, \ldots, x_{n}\right\} \bigcap A ; A \in \mathcal{A}\right\}\right|
$$

Clearly, it suffices to prove that for every $x_{1}, \ldots, x_{n}$,

$$
N_{\mathcal{A}}\left(x_{1}, \ldots, x_{n}\right) \leq \sum_{i=0}^{V_{\mathcal{A}}}\binom{n}{i}
$$

But since $N_{\mathcal{A}}\left(x_{1}, \ldots, x_{n}\right)$ is just the shatter coefficient of the class of finite sets

$$
\left\{\left\{x_{1}, \ldots, x_{n}\right\} \bigcap A ; A \in \mathcal{A}\right\},
$$

we need only to prove the theorem for finite sets. We assume without loss of generality that $\mathcal{A}$ is a class of subsets of $\left\{x_{1}, \ldots, x_{n}\right\}$ with vC dimension $V_{\mathcal{A}}$. Note that in this case $s(\mathcal{A}, n)=|\mathcal{A}|$.

We prove the theorem by induction with respect to $n$ and $V_{\mathcal{A}}$. The statement is obviously true for $n=1$ for any class with $V_{\mathcal{A}} \geq 1$. It is also true for any $n \geq 1$ if $V_{\mathcal{A}}=0$, since $s(\mathcal{A}, n)=1$ for all $n$ in this case. Thus, we assume $V_{\mathcal{A}} \geq 1$. Assume that the statement is true for all $k<n$ for all classes of subsets of $\left\{x_{1}, \ldots, x_{k}\right\}$ with VC dimension not exceeding $V_{\mathcal{A}}$, and for $n$ for all classes with VC dimension smaller than $V_{\mathcal{A}}$. Define the following two classes of subsets of $\left\{x_{1}, \ldots, x_{n}\right\}$ :

$$
\mathcal{A}^{\prime}=\left\{A-\left\{x_{n}\right\} ; A \in \mathcal{A}\right\}
$$

and

$$
\widehat{\mathcal{A}}=\left\{\widehat{A} \in \mathcal{A}: x_{n} \notin \widehat{A}, \widehat{A} \cup\left\{x_{n}\right\} \in \mathcal{A}\right\} .
$$

Note that both $\mathcal{A}^{\prime}$ and $\widehat{\mathcal{A}}$ contain subsets of $\left\{x_{1}, \ldots, x_{n-1}\right\}$. $\widehat{\mathcal{A}}$ contains all sets $\widehat{A}$ that are members of $\mathcal{A}$ such that $\widehat{A} \cup\left\{x_{n}\right\}$ is also in $\mathcal{A}$, but $x_{n} \notin \widehat{A}$.

Then $|\mathcal{A}|=\left|\mathcal{A}^{\prime}\right|+|\widehat{\mathcal{A}}|$. To see this, write

$$
\mathcal{A}^{\prime}=\left\{A-\left\{x_{n}\right\}: x_{n} \in A, A \in \mathcal{A}\right\} \bigcup\left\{A-\left\{x_{n}\right\}: x_{n} \notin A, A \in \mathcal{A}\right\}=\mathcal{B}_{1} \cup \mathcal{B}_{2} .
$$

Thus,

$$
\begin{aligned}
\left|\mathcal{A}^{\prime}\right| & =\left|\mathcal{B}_{1}\right|+\left|\mathcal{B}_{2}\right|-\left|\mathcal{B}_{1} \cap \mathcal{B}_{2}\right| \\
& =\left|\left\{A-\left\{x_{n}\right\}: x_{n} \in A, A \in \mathcal{A}\right\}\right|+\left|\left\{A-\left\{x_{n}\right\}: x_{n} \notin A, A \in \mathcal{A}\right\}\right|-\mid \widehat{\mathcal{A} \mid} \\
& =\left|\left\{A: x_{n} \in A, A \in \mathcal{A}\right\}\right|+\left|\left\{A: x_{n} \notin A, A \in \mathcal{A}\right\}\right|-|\widehat{\mathcal{A}}| \\
& =|\mathcal{A}|-|\widehat{\mathcal{A}}| .
\end{aligned}
$$

Since $\left|\mathcal{A}^{\prime}\right| \leq|\mathcal{A}|$, and $\mathcal{A}^{\prime}$ is a class of subsets of $\left\{x_{1}, \ldots, x_{n-1}\right\}$, the induction hypothesis implies that

$$
\left|\mathcal{A}^{\prime}\right|=s\left(\mathcal{A}^{\prime}, n-1\right) \leq \sum_{i=0}^{V_{\mathcal{A}}}\binom{n-1}{i} .
$$

Next we show that $V_{\widehat{\mathcal{A}}} \leq V_{\mathcal{A}}-1$, which will imply

$$
|\widehat{\mathcal{A}}|=s(\widehat{\mathcal{A}}, n-1) \leq \sum_{i=0}^{V_{A}-1}\binom{n-1}{i}
$$

by the induction hypothesis. To see this, consider a set $S \subset\left\{x_{1}, \ldots, x_{n-1}\right\}$ that is shattered by $\widehat{\mathcal{A}}$. Then $S \cup\left\{x_{n}\right\}$ is shattered by $\mathcal{A}$. To prove this we have to show that any set $S^{\prime} \subset S$ and $S^{\prime} \cup\left\{x_{n}\right\}$ is the intersection of $S \cup\left\{x_{n}\right\}$ and a set from $\mathcal{A}$. Since $S$ is shattered by $\widehat{\mathcal{A}}$, if $S^{\prime} \subset S$, then there exists a set $\widehat{A} \in \widehat{\mathcal{A}}$ such that $S^{\prime}=S \cap \widehat{A}$. But since by definition $x_{n} \notin \widehat{A}$, we must have

$$
S^{\prime}=\left(S \cup\left\{x_{n}\right\}\right) \cap \widehat{A}
$$

and

$$
S^{\prime} \cup\left\{x_{n}\right\}=\left(S \cup\left\{x_{n}\right\}\right) \cap\left(\widehat{A} \cup\left\{x_{n}\right\}\right) .
$$

Since by the definition of $\widehat{\mathcal{A}}$ both $\widehat{A}$ and $\widehat{A} \cup\left\{x_{n}\right\}$ are in $\mathcal{A}$, we see that $S \cup\left\{x_{n}\right\}$ is indeed shattered by $\mathcal{A}$. But any set that is shattered by $\mathcal{A}$ must have cardinality not exceeding $V_{\mathcal{A}}$, therefore $|S| \leq V_{\mathcal{A}}-1$. But $S$ was an arbitrary set shattered by $\widehat{\mathcal{A}}$, which means $V_{\widehat{\mathcal{A}}} \leq V_{\mathcal{A}}-1$. Thus, we have shown that

$$
s(\mathcal{A}, n)=|\mathcal{A}|=\left|\mathcal{A}^{\prime}\right|+|\widehat{\mathcal{A}}| \leq \sum_{i=0}^{V_{\mathcal{A}}}\binom{n-1}{i}+\sum_{i=0}^{V_{A}-1}\binom{n-1}{i} .
$$

Straightforward application of the identity $\binom{n}{i}=\binom{n-1}{i}+\binom{n-1}{i-1}$ shows that

$$
\sum_{i=0}^{V_{\mathcal{A}}}\binom{n}{i}=\sum_{i=0}^{V_{A}}\binom{n-1}{i}+\sum_{i=0}^{V_{A}-1}\binom{n-1}{i}
$$

Theorem 13.2 has some very surprising implications. For example, it follows immediately from the binomial theorem that $s(\mathcal{A}, n) \leq(n+1)^{V_{\mathcal{A}}}$. This means that a shatter coefficient falls in one of two categories: either $s(\mathcal{A}, n)=2^{n}$ for all $n$, or $s(\mathcal{A}, n) \leq(n+1)^{V_{\mathcal{A}}}$, which happens if the VC dimension of $\mathcal{A}$ is finite. We cannot have $s(\mathcal{A}, n) \sim 2^{\sqrt{n}}$, for example. If $V_{\mathcal{A}}<\infty$, the upper bound in Theorem 12.5 decreases exponentially quickly with $n$. Other sharper bounds are given below.

Theorem 13.3. For all $n>2 V$,

$$
s(\mathcal{A}, n) \leq \sum_{i=0}^{V_{\mathcal{A}}}\binom{n}{i} \leq\left(\frac{e n}{V_{\mathcal{A}}}\right)^{V_{\mathcal{A}}}
$$

For $V_{\mathcal{A}}>2, s(\mathcal{A}, n) \leq n^{V_{\mathcal{A}}}$, and for all $V_{\mathcal{A}}, s(\mathcal{A}, n) \leq n^{V_{\mathcal{A}}}+1$.
Theorem 13.3 follows from Theorem 13.4 below. We leave the details as an exercise (see Problem 13.2).

Theorem 13.4. For all $n \geq 1$ and $V_{\mathcal{A}}<n / 2$,

$$
s(\mathcal{A}, n) \leq e^{n \mathcal{H}\left(\frac{V_{A}}{n}\right)}
$$

where $\mathcal{H}(x)=-x \log x-(1-x) \log (1-x)$ for $x \in(0,1)$, and $\mathcal{H}(0)=\mathcal{H}(1)=0$ is the binary entropy function.

Theorem 13.4 is a consequence of Theorem 13.2, and the inequality below. A different, probabilistic proof is sketched in Problem 13.3 (see also Problem 13.4).

Lemma 13.1. For $k<n / 2$,

$$
\sum_{i=0}^{k}\binom{n}{i} \leq e^{n \mathcal{H}\left(\frac{k}{n}\right)}
$$

Proof. Introduce $\lambda=k / n \leq 1 / 2$. By the binomial theorem,

$$
\begin{aligned}
1 & =(\lambda+(1-\lambda))^{n} \\
& \geq \sum_{i=1}^{\lambda n}\binom{n}{i} \lambda^{i}(1-\lambda)^{n-i} \\
& \geq \sum_{i=1}^{\lambda n}\binom{n}{i}\left(\frac{\lambda}{1-\lambda}\right)^{\lambda n}(1-\lambda)^{n} \quad(\text { since } \lambda /(1-\lambda) \leq 1) \\
& =e^{-n \mathcal{H}(\lambda)} \sum_{i=0}^{k}\binom{n}{i}
\end{aligned}
$$

the desired inequality.

REMARK. The binary entropy function $\mathcal{H}(x)$ plays a central role in information theory (see, e.g., Csiszár and Körner (1981), Cover and Thomas (1991)). Its main properties are the following: $\mathcal{H}(x)$ is symmetric around $1 / 2$, where it takes its maximum. It is continuous, concave, strictly monotone increasing in $[0,1 / 2]$, decreasing in $[1 / 2,1]$, and equals zero for $x=0$, and $x=1$.

Next we present some simple results about shatter coefficients of classes that are obtained by combinations of classes of sets.

## Theorem 13.5.

(i) If $\mathcal{A}=\mathcal{A}_{1} \cup \mathcal{A}_{2}$, then $s(\mathcal{A}, n) \leq s\left(\mathcal{A}_{1}, n\right)+s\left(\mathcal{A}_{2}, n\right)$.
(ii) Given a class $\mathcal{A}$ define $\mathcal{A}_{c}=\left\{A^{c} ; A \in \mathcal{A}\right\}$. Then $s\left(\mathcal{A}_{c}, n\right)=s(\mathcal{A}, n)$.
(iii) For $\mathcal{A}=\{\widehat{A} \cap \tilde{A} ; \widehat{A} \in \widehat{\mathcal{A}}, \tilde{A} \in \tilde{\mathcal{A}}\}, s(\mathcal{A}, n) \leq s(\widehat{\mathcal{A}}, n) s(\tilde{\mathcal{A}}, n)$.
(iv) For $\mathcal{A}=\{\widehat{A} \cup \tilde{A} ; \widehat{A} \in \widehat{\mathcal{A}}, \tilde{A} \in \tilde{\mathcal{A}}\}, s(\mathcal{A}, n) \leq s(\widehat{\mathcal{A}}, n) s(\tilde{\mathcal{A}}, n)$.
(v) For $\mathcal{A}=\{\widehat{A} \times \tilde{A} ; \widehat{A} \in \widehat{\mathcal{A}}, \tilde{A} \in \tilde{\mathcal{A}}\}, s(\mathcal{A}, n) \leq s(\widehat{\mathcal{A}}, n) s(\tilde{\mathcal{A}}, n)$.

Proof. (i), (ii), and (v) are trivial. To prove (iii), fix $n$ points $x_{1}, \ldots, x_{n}$, and assume that $\widehat{\mathcal{A}}$ picks $N \leq s(\widehat{\mathcal{A}}, n)$ subsets $C_{1}, \ldots, C_{N}$. Then $\tilde{\mathcal{A}}$ picks from $C_{i}$ at most $s\left(\tilde{\mathcal{A}},\left|C_{i}\right|\right)$ subsets. Therefore, sets of the form $\widehat{A} \cap \tilde{A}$ pick at most

$$
\sum_{i=1}^{N} s\left(\tilde{\mathcal{A}},\left|C_{i}\right|\right) \leq s(\widehat{\mathcal{A}}, n) s(\tilde{\mathcal{A}}, n)
$$

subsets. Here we used the obvious monotonicity property $s(\mathcal{A}, n) \leq s(\mathcal{A}, n+m)$. To prove (iv), observe that

$$
\{\widehat{A} \cup \tilde{A} ; \widehat{A} \in \widehat{\mathcal{A}}, \tilde{A} \in \tilde{\mathcal{A}}\}=\left\{\left(\widehat{A}^{c} \cap \tilde{A}^{c}\right)^{c} ; \widehat{A} \in \widehat{\mathcal{A}}, \tilde{A} \in \tilde{\mathcal{A}}\right\}
$$

The statement now follows from (ii) and (iii).

### 13.2 Shatter Coefficients of Some Classes

Here we calculate shatter coefficients of some simple but important examples of classes of subsets of $\mathcal{R}^{d}$. We begin with a simple observation.

Theorem 13.6. If $\mathcal{A}$ contains finitely many sets, then $V_{\mathcal{A}} \leq \log _{2}|\mathcal{A}|$, and $s(\mathcal{A}, n) \leq|\mathcal{A}|$ for every $n$.

Proof. The first inequality follows from the fact that at least $2^{n}$ sets are necessary to shatter $n$ points. The second inequality is trivial.

In the next example, it is interesting to observe that the bound of Theorem 13.2 is tight.

## Theorem 13.7.

(i) If $\mathcal{A}$ is the class of all half lines: $\mathcal{A}=\{(-\infty, x] ; x \in \mathcal{R}\}$, then $V_{\mathcal{A}}=1$, and

$$
s(\mathcal{A}, n)=n+1=\binom{n}{0}+\binom{n}{1}
$$

(ii) If $\mathcal{A}$ is the class of all intervals in $\mathcal{R}$, then $V_{\mathcal{A}}=2$, and

$$
s(\mathcal{A}, n)=\frac{n(n+1)}{2}+1=\binom{n}{0}+\binom{n}{1}+\binom{n}{2}
$$

Proof. (i) is easy. To see that $V_{\mathcal{A}}=2$ in (ii), observe that if we fix three different points in $\mathcal{R}$, then there is no interval that does not contain the middle point, but does contain the other two. The shatter coefficient can be calculated by counting that there are at most $n-k+1$ sets in $\left\{A \cap\left\{x_{1}, \ldots, x_{n}\right\} ; A \in \mathcal{A}\right\}$ such that $\left|A \cap\left\{x_{1}, \ldots, x_{n}\right\}\right|=k$ for $k=1, \ldots, n$, and one set (namely $\emptyset$ ) such that $\mid A \cap$ $\left\{x_{1}, \ldots, x_{n}\right\} \mid=0$. This gives altogether

$$
1+\sum_{k=1}^{n}(n-k+1)=\frac{n(n+1)}{2}+1
$$

Now we can generalize the result above for classes of intervals and rectangles in $\mathcal{R}^{d}$ :

## Theorem 13.8.

(i) If $\mathcal{A}=\left\{\left(-\infty, x_{1}\right] \times \cdots \times\left(-\infty, x_{d}\right]\right\}$, then $V_{\mathcal{A}}=d$.
(ii) If $\mathcal{A}$ is the class of all rectangles in $\mathcal{R}^{d}$, then $V_{\mathcal{A}}=2 d$.

Proof. We prove (ii). The first part is left as an exercise (Problem 13.5). We have to show that there are $2 d$ points that can be shattered by $\mathcal{A}$, but for any set of $2 d+1$ points there is a subset of it that can not be picked by sets in $\mathcal{A}$. To see the first part just consider the following $2 d$ points:

$$
\begin{aligned}
& (1,0,0, \ldots, 0),(0,1,0,0, \ldots, 0), \ldots,(0,0, \ldots, 0,1) \\
& (-1,0,0, \ldots, 0),(0,-1,0,0, \ldots, 0), \ldots,(0,0, \ldots, 0,-1)
\end{aligned}
$$



Figure 13.1. $2^{4}=16$ rectangles shatter 4 points in $\mathcal{R}^{2}$.

On the other hand, for any given set of $2 d+1$ points we can choose a subset of at most $2 d$ points with the property that it contains a point with largest first coordinate, a point with smallest first coordinate, a point with largest second coordinate, and so forth. Clearly, there is no set in $\mathcal{A}$ that contains these points, but does not contain the others (Figure 13.2).


Figure 13.2. No 5 points can be shattered by rectangles in $\mathcal{R}^{2}$.

Theorem 13.9. (Steele (1975), Dudley (1978)). Let $\mathcal{G}$ be a finite-dimensional vector space of real functions on $\mathcal{R}^{d}$. The class of sets

$$
\mathcal{A}=\{\{x: g(x) \geq 0\}: g \in \mathcal{G}\}
$$

has vc dimension $V_{\mathcal{A}} \leq r$, where $r=$ dimension $(\mathcal{G})$.
Proof. It suffices to show that no set of size $m=1+r$ can be shattered by sets of the form $\{x: g(x) \geq 0\}$. Fix $m$ arbitrary points $x_{1}, \ldots, x_{m}$, and define the linear mapping $L: \mathcal{G} \rightarrow \mathcal{R}^{m}$ as

$$
L(g)=\left(g\left(x_{1}\right), \ldots, g\left(x_{m}\right)\right) .
$$

Then the image of $\mathcal{G}, L(\mathcal{G})$, is a linear subspace of $\mathcal{R}^{m}$ of dimension not exceeding the dimension of $\mathcal{G}$, that is, $m-1$. Then there exists a nonzero vector $\gamma=\left(\gamma_{1}, \ldots, \gamma_{m}\right) \in \mathcal{R}^{m}$, that is orthogonal to $L(\mathcal{G})$, that is, for every $g \in \mathcal{G}$

$$
\gamma_{1} g\left(x_{1}\right)+\ldots+\gamma_{m} g\left(x_{m}\right)=0
$$

We can assume that at least one of the $\gamma_{i}$ 's is negative. Rearrange this equality so that terms with nonnegative $\gamma_{i}$ stay on the left-hand side:

$$
\sum_{i: \gamma_{i} \geq 0} \gamma_{i} g\left(x_{i}\right)=\sum_{i: \gamma_{i}<0}-\gamma_{i} g\left(x_{i}\right) .
$$

Now, suppose that there exists a $g \in \mathcal{G}$ such that the set $\{x: g(x) \geq 0\}$ picks exactly the $x_{i}$ 's on the left-hand side. Then all terms on the left-hand side are nonnegative, while the terms on the right-hand side must be negative, which is a contradiction, so $x_{1}, \ldots, x_{m}$ cannot be shattered, and the proof is completed.

REMARK. Theorem 13.2 implies that the shatter coefficients of the class of sets in Theorem 13.9 are bounded as follows:

$$
s(\mathcal{A}, n) \leq \sum_{i=0}^{r}\binom{n}{i}
$$

In many cases it is possible to get sharper estimates. Let

$$
\mathcal{G}=\left\{\sum_{i=1}^{r} a_{i} \psi_{i}: a_{1}, \ldots, a_{r} \in \mathcal{R}\right\}
$$

be the linear space of functions spanned by some fixed functions $\psi_{1}, \ldots, \psi_{r}$ : $\mathcal{R}^{d} \rightarrow \mathcal{R}$, and define $\Psi(x)=\left(\psi_{1}(x), \ldots, \psi_{r}(x)\right)$. Cover (1965) showed that if for some $x_{1}, \ldots, x_{n} \in \mathcal{R}^{d}$, every $r$-element subset of $\Psi\left(x_{1}\right), \ldots, \Psi\left(x_{n}\right)$ is linearly independent, then the $n$-th shatter coefficient of the class of sets $\mathcal{A}=$ $\{\{x: g(x) \geq 0\}: g \in \mathcal{G}\}$ actually equals

$$
s(\mathcal{A}, n)=2 \sum_{i=0}^{r-1}\binom{n-1}{i}
$$

(see Problem 13.6). By using the last identity in the proof of Theorem 13.2, it is easily seen that the difference between the bound obtained from Theorem 13.9 and the true value is $\binom{n-1}{r}$. Using Cover's result for the shatter coefficients we can actually improve Theorem 13.9 , in that the vc dimension of the class $\mathcal{A}$ equals $r$. To see this, note that

$$
s(\mathcal{A}, r)=2 \sum_{i=0}^{r-1}\binom{r-1}{i}=2 \cdot 2^{r-1}=2^{r}
$$

while

$$
s(\mathcal{A}, r+1)=2 \sum_{i=0}^{r-1}\binom{r}{i}=2 \sum_{i=0}^{r}\binom{r}{i}-2\binom{r}{r}=2 \cdot 2^{r}-2<2^{r+1}
$$

Therefore no $r+1$ points are shattered. It is interesting that Theorem 13.9 above combined with Theorem 13.2 and Theorem 13.3 gives the bound

$$
s(\mathcal{A}, n) \leq n^{r}+1
$$

when $r>2$. Cover's result, however, improves it to

$$
s(\mathcal{A}, n) \leq 2(n-1)^{r-1}+2
$$

Perhaps the most important class of sets is the class of halfspaces in $\mathcal{R}^{d}$, that is, sets containing points falling on one side of a hyperplane. The shatter coefficients of this class can be obtained from the results above:

Corollary 13.1. Let $\mathcal{A}$ be the class of halfspaces, that is, subsets of $\mathcal{R}^{d}$ of the form $\{x: a x \geq b\}$, where $a \in \mathcal{R}^{d}, b \in \mathcal{R}$ take all possible values. Then $V_{\mathcal{A}}=d+1$, and

$$
s(\mathcal{A}, n)=2 \sum_{i=0}^{d}\binom{n-1}{i} \leq 2(n-1)^{d}+2 .
$$

Proof. This is an immediate consequence of the remark above if we take $\mathcal{G}$ to be the linear space spanned by the functions

$$
\phi_{1}(x)=x^{(1)}, \phi_{2}(x)=x^{(2)}, \ldots, \phi_{d}(x)=x^{(d)}, \text { and } \phi_{d+1}(x)=1
$$

where $x^{(1)}, \ldots, x^{(d)}$ denote the $d$ components of the vector $x$.

It is equally simple now to obtain an upper bound on the vc dimension of the class of all closed balls in $\mathcal{R}^{d}$ (see Cover (1965), Devroye (1978), or Dudley (1979) for more information).

Corollary 13.2. Let $\mathcal{A}$ be the class of all closed balls in $\mathcal{R}^{d}$, that is, subsets of $\mathcal{R}^{d}$ of the form

$$
\left\{x=\left(x^{(1)}, \ldots, x^{(d)}\right): \sum_{i=1}^{d}\left|x^{(i)}-a_{i}\right|^{2} \leq b\right\},
$$

where $a_{1}, \ldots, a_{d}, b \in \mathcal{R}$ take all possible values. Then $V_{\mathcal{A}} \leq d+2$.
Proof. If we write

$$
\sum_{i=1}^{d}\left|x^{(i)}-a_{i}\right|^{2}-b=\sum_{i=1}^{d}\left|x^{(i)}\right|^{2}-2 \sum_{i=1}^{d} x^{(i)} a_{i}+\sum_{i=1}^{d} a_{i}^{2}-b
$$

then it is clear that Theorem 13.9 yields the result by setting $\mathcal{G}$ to be the linear space spanned by

$$
\phi_{1}(x)=\sum_{i=1}^{d}\left|x^{(i)}\right|^{2}, \phi_{2}(x)=x^{(1)}, \ldots, \phi_{d+1}(x)=x^{(d)}, \text { and } \phi_{d+2}(x)=1
$$

It follows from Theorems 13.9 and 13.5 (iii) that the class of all polytopes with a bounded number of faces has finite vc dimension. The next negative example demonstrates that this boundedness is necessary.

Theorem 13.10. If $\mathcal{A}$ is the class of all convex polygons in $\mathcal{R}^{2}$, then $V_{\mathcal{A}}=\infty$.


Figure 13.3. Any subset of $n$ points on the unit circle can be picked by a convex polygon.

Proof. Let $x_{1}, \ldots, x_{n} \in \mathcal{R}^{2}$ lie on the unit circle. Then it is easy to see that for any subset of these (different) points there is a polygon that picks that subset.

### 13.3 Linear and Generalized Linear Discrimination Rules

Recall from Chapter 4 that a linear classification rule classifies $x$ into one of the two classes according to whether

$$
a_{0}+\sum_{i=1}^{d} a_{i} x^{(i)}
$$

is positive or negative, where $x^{(1)}, \ldots, x^{(d)}$ denote the components of $x \in \mathcal{R}^{d}$. The coefficients $a_{i}$ are determined by the training sequence. These decisions dichotomize the space $\mathcal{R}^{d}$ by virtue of a halfspace, and assign class 1 to one halfspace, and class 0 to the other. Points on the border are treated as belonging to class 0 . Consider a classifier that adjusts the coefficients by minimizing the number of
errors committed on $D_{n}$. In the terminology of Chapter $12, \mathcal{C}$ is the collection of all linear classifiers.


Figure 13.4. An empirically optimal linear classifier.

Glick (1976) pointed out that for the error probability $L\left(\phi_{n}^{*}\right)$ of this classifier, $L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi) \rightarrow 0$ almost surely. However, from Theorems 12.6, 13.1 and Corollary 13.1, we can now provide more details:

Theorem 13.11. For alln and $\epsilon>0$, the error probability $L\left(\phi_{n}^{*}\right)$ of the empirically optimal linear classifier satisfies

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq 8 n^{d+1} e^{-n \epsilon^{2} / 128} .
$$

Comparing the above inequality with Theorem 4.5 , note that there we selected $\phi_{n}^{*}$ by a specific algorithm, while this result holds for any linear classifier whose empirical error is minimal.

Generalized linear classification rules (see Duda and Hart (1973)) are defined by

$$
g_{n}(x)= \begin{cases}0 & \text { if } a_{0}+\sum_{i=1}^{d^{*}} a_{i} \psi_{i}(x) \geq 0 \\ 1 & \text { otherwise }\end{cases}
$$

where $d^{*}$ is a positive integer, the functions $\psi_{1}, \ldots, \psi_{d^{*}}$ are fixed, and the coefficients $a_{0}, \ldots, a_{d^{*}}$ are functions of the data $D_{n}$. These include for example all quadratic discrimination rules in $\mathcal{R}^{d}$ when we choose all functions that are either components of $x$, or squares of components of $x$, or products of two components of $x$. That is, the functions $\psi_{i}(x)$ are of the form either $x^{(j)}$, or $x^{(j)} x^{(k)}$. In all, $d^{*}=2 d+d(d-1) / 2$. The argument used for linear discriminants remain valid, and we obtain

Theorem 13.12. Let $\mathcal{C}$ be the class of generalized linear discriminants (i.e., the coefficients vary, the basis functions $\psi_{i}$ are fixed). For the error probability $L\left(\phi_{n}^{*}\right)$ of the empirically optimal classifier, for all $d^{*}>1, n$ and $\epsilon>0$, we have

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq 8 n^{d^{*}+1} e^{-n \epsilon^{2} / 128} .
$$

Also, for $n>2 d^{*}+1$,

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq 8 e^{n \mathcal{H}\left(\frac{d^{*}+1}{n}\right)} e^{-n \epsilon^{2} / 128}
$$

The second inequality is obtained by using the bound of Theorem 13.4 for the shatter coefficients. Note nevertheless that unless $d^{*}$ (and therefore $\mathcal{C}$ ) is allowed to increase with $n$, there is no hope of obtaining universal consistency. The question of universal consistency will be addressed in Chapters 17 and 18.

### 13.4 Convex Sets and Monotone Layers

Classes of infinite vc dimension are not hopeless by any means. In this section, we offer examples that will show how they may be useful in pattern recognition. The classes of interest to us for now are

$$
\begin{aligned}
\mathcal{C}= & \left\{\text { all convex sets of } \mathcal{R}^{2}\right\} \\
\mathcal{L}= & \left\{\text { all monotone layers of } \mathcal{R}^{2}\right. \text {, i.e., all sets of the form } \\
& \left.\left\{\left(x_{1}, x_{2}\right): x_{2} \leq \psi\left(x_{1}\right)\right\} \text { for some nonincreasing function } \psi\right\}
\end{aligned}
$$

In discrimination, this corresponds to making decisions of the form $\phi(x)=I_{\{x \in C\}}$, $C \in \mathcal{C}$, or $\phi(x)=I_{\{x \notin C\}}, C \in \mathcal{C}$, and similarly for $\mathcal{L}$. Decisions of these forms are important in many situations. For example, if $\eta(x)$ is monotone decreasing in both components of $x \in \mathcal{R}^{2}$, then the Bayes rule is of the form $g^{*}(x)=I_{\{x \in L\}}$ for some $L \in \mathcal{L}$. We have pointed out elsewhere (Theorem 13.10, and Problem 13.19) that $V_{\mathcal{C}}=V_{\mathcal{L}}=\infty$. To see this, note that any set of $n$ points on the unit circle is shattered by $\mathcal{C}$, while any set of $n$ points on the antidiagonal $x_{2}=-x_{1}$ is shattered by $\mathcal{L}$. Nevertheless, shattering becomes unlikely if $X$ has a density. Our starting point here is the bound obtained while proving Theorem 12.5:

$$
\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|\mu_{n}(A)-\mu(A)\right|>\epsilon\right\} \leq 8 \mathbf{E}\left\{N_{\mathcal{A}}\left(X_{1}, \ldots, X_{n}\right)\right\} e^{-n \epsilon^{2} / 32}
$$

where $N_{\mathcal{A}}\left(X_{1}, \ldots, X_{n}\right)$ is the number of sets in $\left\{A \cap\left\{X_{1}, \ldots, X_{n}\right\}: A \in \mathcal{A}\right\}$. The following theorem is essential:

Theorem 13.13. If $X$ has a density $f$ on $\mathcal{R}^{2}$, then

$$
\mathbf{E}\left\{N_{\mathcal{A}}\left(X_{1}, \ldots, X_{n}\right)\right\}=2^{o(n)}
$$

when $\mathcal{A}$ is either $\mathcal{C}$ or $\mathcal{L}$.
This theorem, a proof of which is a must for the reader, implies the following:

Corollary 13.3. Let $X$ have a density $f$ on $\mathcal{R}^{2}$. Let $\phi_{n}^{*}$ be picked by minimizing the empirical error over all classifiers of the form

$$
\phi(x)= \begin{cases}1 & \text { if } x \in A \\ 0 & \text { if } x \notin A,\end{cases}
$$

where $A$ or $A^{c}$ is in $\mathcal{C}($ or $\mathcal{L})$. Then

$$
L\left(\phi_{n}^{*}\right) \rightarrow \inf _{\phi=I_{C} \text { for } C \text { or } C^{c} \text { in } \mathcal{C}} L(\phi)
$$

with probability one (and similarly for $\mathcal{L}$ ).

Proof. This follows from the inequality of Lemma 8.2, Theorem 13.13, and the Borel-Cantelli lemma.

REmark. Theorem 13.13 and the corollary may be extended to $\mathcal{R}^{d}$, but this generalization holds nothing new and will only result in tedious notations.

Proof of Theorem 13.13. We show the theorem for $\mathcal{L}$, and indicate the proof for $\mathcal{C}$. Take two sequences of integers, $m$ and $r$, where $m \sim \sqrt{n}$ and $r \sim m^{1 / 3}$, so that $m \rightarrow \infty$, yet $r^{2} / m \rightarrow 0$, as $n \rightarrow \infty$. Consider the set $[-r, r]^{2}$ and partition each side into $m$ equal intervals, thus obtaining an $m \times m$ grid of square cells $C_{1}, \ldots, C_{m^{2}}$. Denote $C_{0}=\mathcal{R}^{2}-[-r, r]^{2}$. Let $N_{0}, N_{1}, \ldots, N_{m^{2}}$ be the number of points among $X_{1}, \ldots, X_{n}$ that belong to these cells. The vector ( $N_{0}, N_{1}, \ldots, N_{m^{2}}$ ) is clearly multinomially distributed.

Let $\psi$ be a nonincreasing function $\mathcal{R} \rightarrow \mathcal{R}$ defining a set in $\mathcal{L}$ by $L=\left\{\left(x_{1}, x_{2}\right)\right.$ : $\left.x_{2} \leq \psi\left(x_{1}\right)\right\}$. Let $C(\psi)$ be the collection of all cell sets cut by $\psi$, that is, all cells with a nonempty intersection with both $L$ and $L^{c}$. The collection $C(\psi)$ is shaded in Figure 13.5.


Figure 13.5. A monotone layer and bordering cells.

We bound $N_{\mathcal{L}}\left(X_{1}, \ldots, X_{n}\right)$ from above, conservatively, as follows:

$$
\begin{equation*}
\sum_{\text {all possible } C(\psi)}\left(\prod_{C_{i} \in C(\psi)} 2^{N_{i}}\right) 2^{N_{0}} . \tag{13.1}
\end{equation*}
$$

The number of different collections $C(\psi)$ cannot exceed $2^{2 m}$ because each cell in $C(\psi)$ may be obtained from its predecessor cell by either moving right on the same row or moving down one cell in the same column. For a particular collection, denoting $p_{i}=\mathbf{P}\left\{X \in C_{i}\right\}$, we have

$$
\begin{aligned}
\mathbf{E} & \left\{\left(\prod_{C_{i} \in C(\psi)} 2^{N_{i}}\right) 2^{N_{0}}\right\} \\
& =\left(\sum_{C_{i} \in C(\psi)} 2 p_{i}+2 p_{0}+1-\sum_{C_{i} \in C(\psi)} p_{i}-p_{0}\right)^{n}
\end{aligned}
$$

(by applying Lemma A.7)
$\leq \exp \left(n\left(\sum_{c_{i} \in C(\psi)} p_{i}+p_{0}\right)\right)$
$\leq \exp \left(n\left(\begin{array}{c}\left.\left.\sup ^{\text {all sets } A \text { with } \lambda(A) \leq 8 r^{2} / m} \int_{A} f+\int_{C_{0}} f\right)\right), ~(1)\end{array}\right.\right.$
(since $\left.\lambda\left(\bigcup_{i: C_{i} \in C(\psi)} C_{i}\right) \leq 2 m(2 r / m)^{2}=8 r^{2} / m\right)$
$=e^{o(n)}$,
because $r^{2} / m \rightarrow 0$ and $r \rightarrow \infty$ and by the absolute continuity of $X$. As this estimate is uniform over all collections $C(\psi)$, we see that the expected value of (13.1) is not more than $2^{2 m} e^{o(n)}=e^{o(n)}$. The argument for the collection $\mathcal{C}$ of convex sets is analogous.

REMARK. The theorem implies that if $\mathcal{A}$ is the class of all convex sets, then $\sup _{A \in \mathcal{A}}\left|\mu_{n}(A)-\mu(A)\right| \rightarrow 0$ with probability one whenever $\mu$ has a density. This is a special case of a result of Ranga Rao (1962). Assume now that the density $f$ of $X$ is bounded and of bounded support. Then in the proof above we may take $r$ fixed so that $[-r, r]^{2}$ contains the support of $f$. Then the estimate of Theorem 13.13 is

$$
\mathbf{E}\left\{N_{\mathcal{A}}\left(X_{1}, \ldots, X_{n}\right)\right\} \leq 2^{2 m} \cdot e^{8 n\|f\|_{\infty} r^{2} / m}
$$

(where $\|f\|_{\infty}$ denotes the bound on $f$ )

$$
\begin{aligned}
= & 2^{4 r \sqrt{n\|f\|_{\infty}}} e^{4 r \sqrt{n\|f\|_{\infty}}} \\
& \left(\text { if we take } m=2 r \sqrt{n\|f\|_{\infty}}\right) \\
\leq & e^{\alpha \sqrt{n}}
\end{aligned}
$$

for a constant $\alpha$. This implies by Corollary 12.1 that

$$
\mathbf{E}\left\{L\left(\phi_{n}^{*}\right)-\inf _{\phi=I_{C} \text { for } C \text { or } C^{c} \text { in } \mathcal{C}} L(\phi)\right\}=O\left(n^{-1 / 4}\right)
$$

This latter inequality was proved by Steele (1975). To see that it cannot be extended to arbitrary densities, observe that the data points falling on the convex hull (or upper layer) of the points $X_{1}, \ldots, X_{n}$ can always be shattered by convex sets (or monotone layers, respectively). Thus, $N_{\mathcal{A}}\left(X_{1}, \ldots, X_{n}\right)$ is at least $2^{M_{n}}$, where $M_{n}$ is the number of points among $X_{1}, \ldots, X_{n}$, falling on the convex hull (or upper layer) of $X_{1}, \ldots, X_{n}$. Thus,

$$
\mathbf{E}\left\{N_{\mathcal{A}}\left(X_{1}, \ldots, X_{n}\right)\right\} \geq \mathbf{E}\left\{2^{M_{n}}\right\} \geq 2^{\mathbf{E}\left\{M_{n}\right\}}
$$

But it follows from results of Carnal (1970) and Devroye (1991b) that for each $a<1$, there exists a density such that

$$
\limsup _{n \rightarrow \infty} \frac{\mathbf{E}\left\{M_{n}\right\}}{n^{a}}>1
$$

The important point of this interlude is that with infinite VC dimension, we may under some circumstances get expected error rates that are $o(1)$ but larger than $1 / \sqrt{n}$. However, the bounds are sometimes rather loose. The reason is the looseness of the Vapnik-Chervonenkis inequality when the collections $\mathcal{A}$ become very big. To get such results for classes with infinite vc dimension it is necessary to impose some conditions on the distribution. We will prove this in Chapter 14.

## Problems and Exercises

Problem 13.1. Show that the inequality of Theorem 13.2 is tight, that is, exhibit a class $\mathcal{A}$ of sets such that for each $n, s(\mathcal{A}, n)=\sum_{i=0}^{V_{\mathcal{A}}}\binom{n}{i}$.

Problem 13.2. Show that for all $n>2 V_{\mathcal{A}}$

$$
\sum_{i=0}^{V_{\mathcal{A}}}\binom{n}{i} \leq 2\left(\frac{n^{V_{\mathcal{A}}}}{V_{\mathcal{A}}!}\right) \leq\left(\frac{e n}{V_{\mathcal{A}}}\right)^{V_{\mathcal{A}}}
$$

and that

$$
\sum_{i=0}^{V_{\mathcal{A}}}\binom{n}{i} \leq n^{V_{\mathcal{A}}}
$$

if $V_{\mathcal{A}}>2$. Hint: There are several ways to prove the statements. One can proceed directly by using the recurrence $\sum_{i=0}^{V_{\mathcal{A}}}\binom{n}{i}=\sum_{i=0}^{V_{\mathcal{A}}}\binom{n-1}{i}+\sum_{i=0}^{V_{\mathcal{A}}-1}\binom{n-1}{i}$. A simpler way to prove $\sum_{i=0}^{V_{\mathcal{A}}}\binom{n}{i} \leq\left(\frac{e n}{V_{\mathcal{A}}}\right)^{V_{\mathcal{A}}}$ is by using Theorem 13.4. The third inequality is an immediate consequence of the first two.

Problem 13.3. Give an alternative proof of Lemma 13.1 by completing the following probabilistic argument. Observe that for $k^{\prime}=n-k$,

$$
\sum_{i=0}^{k}\binom{n}{i}=\sum_{i=k^{\prime}}^{n}\binom{n}{i}=2^{n} \mathbf{P}\left\{B_{n} \geq k^{\prime}\right\}
$$

where $B_{n}$ is a binomial random variable with parameters $n$ and $1 / 2$. Then Chernoff's bounding technique (see the proof of Theorem 8.1) may be used to bound this probability: for all $s>0$,

$$
\mathbf{P}\left\{B_{n} \geq k^{\prime}\right\} \leq e^{-s k^{\prime}} \mathbf{E}\left\{e^{s B_{n}}\right\}=\exp \left(-n\left(s k^{\prime} / n-\log \left(\frac{e^{s}+1}{2}\right)\right)\right)
$$

Take the derivative of the exponent with respect to $s$ to minimize the upper bound. Substitute the obtained value into the bound to get the desired inequality.

Problem 13.4. Let $B$ be a binomial random variable with parameters $n$ and $p$. Prove that for $k>n p$

$$
\mathbf{P}\{B \geq k\} \leq \exp \left(-n\left(\mathcal{H}(k / n)+\frac{k}{n} \log p+\frac{n-k}{n} \log (1-p)\right)\right)
$$

Hint: Use Chernoff's bounding technique as in the previous problem.
Problem 13.5. Prove part (i) of Theorem 13.8.
Problem 13.6. Prove that for the class of sets defined in the remark following Theorem 13.9,

$$
s(\mathcal{A}, n)=2 \sum_{i=0}^{r-1}\binom{n-1}{i}
$$

(Cover (1965)). Hint: Proceed by induction with respect to $n$ and $r$. In particular, show that the recurrence $s\left(\mathcal{A}_{r}, n\right)=s\left(\mathcal{A}_{r}, n-1\right)+s\left(\mathcal{A}_{r-1}, n-1\right)$ holds, where $\mathcal{A}_{r}$ denotes the class of sets defined as $\{x: g(x) \geq 0\}$ where $g$ runs through a vector space spanned by the first $r$ of the sequence of functions $\psi_{1}, \psi_{2}, \ldots$

Problem 13.7. Let $\mathcal{A}$ and $\mathcal{B}$ be two families of subsets of $\mathcal{R}^{d}$. Assume that for some $r \geq 2$,

$$
s(\mathcal{A}, n) \leq n^{r} s(\mathcal{B}, n)
$$

for all $n \geq 1$. Show that if $V_{\mathcal{B}}>2$, then $V_{\mathcal{A}} \leq 2\left(V_{\mathcal{B}}+r-1\right) \log \left(V_{\mathcal{B}}+r-1\right)$. Hint: By Theorem 13.3, $s(\mathcal{A}, n) \leq n^{V_{\mathcal{B}}+r-1}$. Clearly, $V_{\mathcal{A}}$ is not larger than any $k$ for which $k^{V_{\mathcal{B}}+r-1}<2^{k}$.

Problem 13.8. Determine the vc dimension of the class of subsets of the real line such that each set in the class can be written as a union of $k$ intervals.

Problem 13.9. Determine the vc dimension of the collection of all polygons with $k$ vertices in the plane.

Problem 13.10. What is the vc dimension of the collection of all ellipsoids of $\mathcal{R}^{d}$ ?
Problem 13.11. Determine the vc dimension of the collection of all subsets of $\{1, \ldots, k\}^{d}$, where $k$ and $d$ are fixed. How does the answer change if we restrict the subsets to those of cardinality $l \leq k^{d}$ ?

Problem 13.12. Let $\mathcal{A}$ consist of all simplices of $\mathcal{R}^{d}$, that is, all sets of the form

$$
\left\{x: x=\sum_{i=1}^{d+1} \lambda_{i} x_{i} \text { for some } \lambda_{1}, \ldots, \lambda_{d+1}: \sum_{i=1}^{d+1} \lambda_{i}=1, \lambda_{i} \geq 0\right\},
$$

where $x_{1}, \ldots, x_{d+1}$ are fixed points of $\mathcal{R}^{d}$. Determine the vC dimension of $\mathcal{A}$.
Problem 13.13. Let $A\left(x_{1}, \ldots, x_{k}\right)$ be the set of all $x \in \mathcal{R}$ that are of the form

$$
x=\psi_{1}(i) x_{1}+\psi_{2}(i) x_{2}+\cdots+\psi_{k}(i) x_{k}, \quad i=1,2, \ldots,
$$

where $x_{1}, \ldots, x_{k}$ are fixed numbers and $\psi_{1}, \ldots, \psi_{k}$ are fixed functions on the integers. Let $\mathcal{A}=\left\{A\left(x_{1}, \ldots, x_{k}\right): x_{1}, \ldots, x_{k} \in \mathcal{R}\right\}$. Determine the vC dimension of $\mathcal{A}$.

Problem 13.14. In some sense, vc dimension measures the "size" of a class of sets. However it has little to do with cardinalities, as this exercise demonstrates. Exhibit a class of subsets of the integers with uncountably many sets, yet vc dimension 1. (This property was pointed out to us by András Faragó.) Note: On the other hand, the class of all subsets of integers cannot be written as a countable union of classes with finite vc dimension (Theorem 18.6). Hint: Find a class of subsets of the reals with the desired properties, and make a proper correspondence between sets of integers and sets in the class.

Problem 13.15. Show that if a class of sets $\mathcal{A}$ is linearly ordered by inclusion, that is, for any pair of sets $A, B \in \mathcal{A}$ either $A \subset B$ or $B \subset A$ and $|\mathcal{A}| \geq 2$, then $V_{\mathcal{A}}=1$. Conversely, assume that $V_{\mathcal{A}}=1$ and for every set $B$ with $|B|=2$,

$$
\emptyset, B \in\{A \cap B: A \in \mathcal{A}\} .
$$

Prove that then $\mathcal{A}$ is linearly ordered by inclusion (Dudley (1984)).
Problem 13.16. We say that four sets $A, B, C, D$ form a diamond if $A \subset B \subset C, A \subset$ $D \subset C$, but $B \not \subset D$ and $D \not \subset C$. Let $\mathcal{A}$ be a class of sets. Show that $V_{\mathcal{A}} \geq 2$ if and only if for some set $R$, the class $\{A \cap R: A \in \mathcal{A}\}$ includes a diamond (Dudley (1984)).

Problem 13.17. Let $\mathcal{A}$ be a class of sets, and define its density by

$$
D_{\mathcal{A}}=\inf \left\{r>0: \sup _{n \geq 1} \frac{s(\mathcal{A}, n)}{n^{r}}<\infty\right\} .
$$

Verify the following properties:
(1) $D_{\mathcal{A}} \leq V_{\mathcal{A}}$;
(2) For each positive integer $k$, there exists a class $\mathcal{A}$ of sets such that $V_{\mathcal{A}}=k$, yet $D_{\mathcal{A}}=0$; and
(3) $D_{\mathcal{A}}<\infty$ if and only if $V_{\mathcal{A}}<\infty$
(Assouad (1983a)).
Problem 13.18. Continued. Let $\mathcal{A}$ and $\mathcal{A}^{\prime}$ be classes of sets, and define $\mathcal{B}=\mathcal{A} \cup \mathcal{A}^{\prime}$. Show that
(1) $D_{\mathcal{B}}=\max \left(D_{\mathcal{A}}, D_{\mathcal{A}^{\prime}}\right)$;
(2) $V_{\mathcal{B}} \leq V_{\mathcal{A}}+V_{\mathcal{A}^{\prime}}+1$; and
(3) For every pair of positive integers $k, m$ there exist classes $\mathcal{A}$ and $\mathcal{A}^{\prime}$ such that $V_{\mathcal{A}}=k, V_{\mathcal{A}^{\prime}}=m$, and $V_{\mathcal{B}}=k+m+1$
(Assouad (1983a)).
Problem 13.19. A set $A \subset \mathcal{R}^{d}$ is called a monotone layer if $x \in A$ implies that $y \in A$ for all $y \leq x$ (i.e., each component of $x$ is not larger than the corresponding component of $y$ ). Show that the class of all monotone layers has infinite VC dimension.

Problem 13.20. Let $(X, Y)$ be a pair of random variables in $\mathcal{R}^{2} \times\{0,1\}$ such that $Y=I_{\{X \in A\}}$, where $A$ is a convex set. Let $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$ be an i.i.d. training sequence, and consider the classifier

$$
g_{n}(x)= \begin{cases}1 & \text { if } x \text { is in the convex hull of the } X_{i} \text { 's with } Y=1 \\ 0 & \text { otherwise. }\end{cases}
$$

Find a distribution for which $g_{n}$ is not consistent, and find conditions for consistency. Hint: Recall Theorem 13.10 and its proof.

## 14

## Lower Bounds for Empirical Classifier Selection

In Chapter 12 a classifier was selected by minimizing the empirical error over a class of classifiers $\mathcal{C}$. With the help of the Vapnik-Chervonenkis theory we have been able to obtain distribution-free performance guarantees for the selected rule. For example, it was shown that the difference between the expected error probability of the selected rule and the best error probability in the class behaves at least as well as $O\left(\sqrt{V_{\mathcal{C}} \log n / n}\right)$, where $V_{\mathcal{C}}$ is the Vapnik-Chervonenkis dimension of $\mathcal{C}$, and $n$ is the size of the training data $D_{n}$. (This upper bound is obtained from Theorem 12.5. Corollary 12.5 may be used to replace the $\log n$ term with $\log V_{\mathcal{C}}$.) Two questions arise immediately: Are these upper bounds (at least up to the order of magnitude) tight? Is there a much better way of selecting a classifier than minimizing the empirical error? This chapter attempts to answer these questions. As it turns out, the answer is essentially affirmative for the first question, and negative for the second.

These questions were also asked in the learning theory setup, where it is usually assumed that the error probability of the best classifier in the class is zero (see Blumer, Ehrenfeucht, Haussler, and Warmuth (1989), Haussler, Littlestone, and Warmuth (1988), and Ehrenfeucht, Haussler, Kearns, and Valiant (1989)). In this case, as the bound of Theorem 12.7 implies, the error of the rule selected by minimizing the empirical error is within $O\left(V_{\mathcal{C}} \log n / n\right)$ of that of the best in the class (which equals zero, by assumption). We will see that essentially there is no way to beat this upper bound either.

### 14.1 Minimax Lower Bounds

Let us formulate exactly what we are interested in. Let $\mathcal{C}$ be a class of decision functions $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$. The training sequence $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$ is used to select the classifier $g_{n}(X)=g_{n}\left(X, D_{n}\right)$ from $\mathcal{C}$, where the selection is based on the data $D_{n}$. We emphasize here that $g_{n}$ can be an arbitrary function of the data, we do not restrict our attention to empirical error minimization, where $g_{n}$ is a classifier in $\mathcal{C}$ that minimizes the number errors committed on the data $D_{n}$.

As before, we measure the performance of the selected classifier by the difference between the error probability $L\left(g_{n}\right)=\mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}$ of the selected classifier and that of the best in the class. To save space further on, denote this optimum by

$$
L_{\mathcal{C}} \stackrel{\text { def }}{=} \inf _{\phi \in \mathcal{C}} \mathbf{P}\{\phi(X) \neq Y\} .
$$

In particular, we seek lower bounds for

$$
\sup \mathbf{P}\left\{L\left(g_{n}\right)-L_{\mathcal{C}}>\epsilon\right\},
$$

and

$$
\sup \mathbf{E} L\left(g_{n}\right)-L_{\mathcal{C}},
$$

where the supremum is taken over all possible distributions of the pair $(X, Y)$. A lower bound for one of these quantities means that no matter what our method of picking a rule from $\mathcal{C}$ is, we may face a distribution such that our method performs worse than the bound. This view may be criticized as too pessimistic. However, it is clearly a perfectly meaningful question to pursue, as typically we have no other information available than the training data, so we have to be prepared for the worst situation.

Actually, we investigate a stronger problem, in that the supremum is taken over all distributions with $L_{\mathcal{C}}$ kept at a fixed value between zero and $1 / 2$. We will see that the bounds depend on $n, V_{\mathcal{C}}$, and $L_{\mathcal{C}}$ jointly. As it turns out, the situations for $L_{\mathcal{C}}>0$ and $L_{\mathcal{C}}=0$ are quite different. Because of its relative simplicity, we first treat the case $L_{\mathcal{C}}=0$. All the proofs are based on a technique called "the probabilistic method." The basic idea here is that the existence of a "bad" distribution is proved by considering a large class of distributions, and bounding the average behavior over the class.

Lower bounds on the probabilities $\mathbf{P}\left\{L\left(g_{n}\right)-L_{\mathcal{C}}>\epsilon\right\}$ may be translated into lower bounds on the sample complexity $N(\epsilon, \delta)$. We obtain lower bounds for the size of the training sequence such that for any classifier, $\mathbf{P}\left\{L\left(g_{n}\right)-L_{\mathcal{C}}>\epsilon\right\}$ cannot be smaller than $\delta$ for all distributions if $n$ is smaller than this bound.

### 14.2 The Case $L_{\mathcal{C}}=0$

In this section we obtain lower bounds under the assumption that the best classifier in the class has zero error probability. In view of Theorem 12.7 we see that the
situation here is different from when $L_{\mathcal{C}}>0$; there exist methods of picking a classifier from $\mathcal{C}$ (e.g., minimization of the empirical error) such that the error probability decreases to zero at a rate of $O\left(V_{\mathcal{C}} \log n / n\right)$. We obtain minimax lower bounds close to the upper bounds obtained for empirical error minimization. For example, Theorem 14.1 shows that if $L_{\mathcal{C}}=0$, then the expected error probability cannot decrease faster than a sequence proportional to $V_{\mathcal{C}} / n$ for some distributions.

Theorem 14.1. (Vapnik and Chervonenkis (1974c); Haussler, Littlestone, and Warmuth (1988)). Let $\mathcal{C}$ be a class of discrimination functions with vc dimension $V$. Let $\mathcal{X}$ be the set of all random variables $(X, Y)$ for which $L_{\mathcal{C}}=0$. Then, for every discrimination rule $g_{n}$ based upon $X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}$, and $n \geq V-1$,

$$
\sup _{(X, Y) \in \mathcal{X}} \mathbf{E} L_{n} \geq \frac{V-1}{2 e n}\left(1-\frac{1}{n}\right)
$$

Proof. The idea is to construct a family $\mathcal{F}$ of $2^{V-1}$ distributions within the distributions with $L_{\mathcal{C}}=0$ as follows: first find points $x_{1}, \ldots, x_{V}$ that are shattered by $\mathcal{C}$. Each distribution in $\mathcal{F}$ is concentrated on the set of these points. A member in $\mathcal{F}$ is described by $V-1$ bits, $b_{1}, \ldots, b_{V-1}$. For convenience, this is represented as a bit vector $b$. Assume $V-1 \leq n$. For a particular bit vector, we let $X=x_{i}(i<V)$ with probability $1 / n$ each, while $X=x_{V}$ with probability $1-(V-1) / n$. Then set $Y=f_{b}(X)$, where $f_{b}$ is defined as follows:

$$
f_{b}(x)= \begin{cases}b_{i} & \text { if } x=x_{i}, i<V \\ 0 & \text { if } x=x_{V}\end{cases}
$$

Note that since $Y$ is a function of $X$, we must have $L^{*}=0$. Also, $L_{\mathcal{C}}=0$, as the set $\left\{x_{1}, \ldots, x_{V}\right\}$ is shattered by $\mathcal{C}$, i.e., there is a $g \in \mathcal{C}$ with $g\left(x_{i}\right)=f_{b}\left(x_{i}\right)$ for $1 \leq i \leq V$. Clearly,

$$
\begin{aligned}
\sup _{(X, Y): L_{\mathcal{C}}} & =0 \\
& \mathbf{E}\left\{L_{n}-L_{\mathcal{C}}\right\} \\
& \sup _{(X, Y) \in \mathcal{F}} \mathbf{E}\left\{L_{n}-L_{\mathcal{C}}\right\} \\
& =\sup _{b} \mathbf{E}\left\{L_{n}-L_{\mathcal{C}}\right\} \\
& \geq \mathbf{E}\left\{L_{n}-L_{\mathcal{C}}\right\}
\end{aligned}
$$

$$
\text { (where } b \text { is replaced by } B \text {, uniformly distributed over }\{0,1\}^{V-1} \text { ) }
$$

$$
=\mathbf{E}\left\{L_{n}\right\}
$$

$$
=\mathbf{P}\left\{g_{n}\left(X, X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right) \neq f_{B}(X)\right\}
$$

The last probability may be viewed as the error probability of the decision function $g_{n}: \mathcal{R}^{d} \times\left(\mathcal{R}^{d} \times\{0,1\}\right)^{n} \rightarrow\{0,1\}$ in predicting the value of the random variable $f_{B}(X)$ based on the observation $Z_{n}=\left(X, X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right)$. Naturally, this probability is bounded from below by the Bayes probability of error

$$
L^{*}\left(Z_{n}, f_{B}(X)\right)=\inf _{g_{n}} \mathbf{P}\left\{g_{n}\left(Z_{n}\right) \neq f_{B}(X)\right\}
$$

corresponding to the decision problem $\left(Z_{n}, f_{B}(X)\right)$. By the results of Chapter 2,

$$
L^{*}\left(Z_{n}, f_{B}(X)\right)=\mathbf{E}\left\{\min \left(\eta^{*}\left(Z_{n}\right), 1-\eta^{*}\left(Z_{n}\right)\right)\right\}
$$

where $\eta^{*}\left(Z_{n}\right)=\mathbf{P}\left\{f_{B}(X)=1 \mid Z_{n}\right\}$. Observe that

$$
\eta^{*}\left(Z_{n}\right)= \begin{cases}1 / 2 & \text { if } X \neq X_{1}, \ldots, X \neq X_{n}, X \neq x_{V} \\ 0 \text { or } 1 & \text { otherwise }\end{cases}
$$

Thus, we see that

$$
\begin{aligned}
\sup _{(X, Y): L_{\mathcal{C}}=0} \mathbf{E}\left\{L_{n}-L_{\mathcal{C}}\right\} & \geq L^{*}\left(Z_{n}, f_{B}(X)\right) \\
& =\frac{1}{2} \mathbf{P}\left\{X \neq X_{1}, \ldots, X \neq X_{n}, X \neq x_{V}\right\} \\
& =\frac{1}{2} \sum_{i=1}^{V-1} \mathbf{P}\left\{X=x_{i}\right\}\left(1-\mathbf{P}\left\{X=x_{i}\right\}\right)^{n} \\
& =\frac{V-1}{2 n}(1-1 / n)^{n} \\
& \geq \frac{V-1}{2 e n}\left(1-\frac{1}{n}\right) \quad\left(\text { since }(1-1 / n)^{n-1} \downarrow 1 / e\right)
\end{aligned}
$$

This concludes the proof.

Minimax lower bounds on the probability $\mathbf{P}\left\{L_{n} \geq \epsilon\right\}$ can also be obtained. These bounds have evolved through several papers: see Ehrenfeucht, Haussler, Kearns and Valiant (1989); and Blumer, Ehrenfeucht, Haussler and Warmuth (1989). The tightest bounds we are aware of thus far are given by the next theorem.

Theorem 14.2. (Devroye and Lugosi (1995)). Let $\mathcal{C}$ be a class of discrimination functions with vc dimension $V \geq 2$. Let $\mathcal{X}$ be the set of all random variables ( $X, Y$ ) for which $L_{\mathcal{C}}=0$. Assume $\epsilon \leq 1 / 4$. Assume $n \geq V-1$. Then for every discrimination rule $g_{n}$ based upon $X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}$,

$$
\sup _{(X, Y) \in \mathcal{X}} \mathbf{P}\left\{L_{n} \geq \epsilon\right\} \geq \frac{1}{e \sqrt{\pi V}}\left(\frac{2 n e \epsilon}{V-1}\right)^{(V-1) / 2} e^{-4 n \epsilon /(1-4 \epsilon)}
$$

If on the other hand $n \geq 15$ and $n \leq(V-1) /(12 \epsilon)$, then

$$
\sup _{(X, Y) \in \mathcal{X}} \mathbf{P}\left\{L_{n} \geq \epsilon\right\} \geq \frac{1}{10} .
$$

Proof. We randomize as in the proof of Theorem 14.1. The difference now is that we pick $x_{1}, \ldots, x_{V-1}$ with probability $p$ each. Thus, $\mathbf{P}\left\{X=x_{V}\right\}=1-p(V-1)$.

We inherit the notation from the proof of Theorem 14.1. For a fixed $b$, denote the error probability by

$$
L_{n}(b)=\mathbf{P}\left\{g_{n}\left(X, X_{1}, f_{b}\left(X_{1}\right), \ldots, X_{n}, f_{b}\left(X_{n}\right)\right) \neq f_{b}(X) \mid X_{1}, \ldots, X_{n}\right\} .
$$

We now randomize and replace $b$ by $B$. Clearly,

$$
\begin{aligned}
\sup _{(X, Y): L_{c}=0} \mathbf{P}\left\{L_{n} \geq \epsilon\right\} & \geq \sup _{b} \mathbf{P}\left\{L_{n}(b) \geq \epsilon\right\} \\
& \geq \mathbf{E}\left\{\mathbf{P}\left\{L_{n}(B) \geq \epsilon \mid B\right\}\right\} \\
& =\mathbf{P}\left\{L_{n}(B) \geq \epsilon\right\} .
\end{aligned}
$$

As in the proof of Theorem 14.1, observe that $L_{n}(B)$ cannot be smaller than the Bayes risk corresponding to the decision problem $\left(Z_{n}, f_{B}(X)\right.$ ), where

$$
Z_{n}=\left(X, X_{1}, f_{B}\left(X_{1}\right), \ldots, X_{n}, f_{B}\left(X_{n}\right)\right) .
$$

Thus,

$$
L_{n}(B) \geq \mathbf{E}\left\{\min \left(\eta^{*}\left(Z_{n}\right), 1-\eta^{*}\left(Z_{n}\right)\right) \mid X_{1}, \ldots, X_{n}\right\} .
$$

As in the proof of Theorem 14.1, we see that

$$
\begin{aligned}
& \mathbf{E}\left\{\min \left(\eta^{*}\left(Z_{n}\right), 1-\eta^{*}\left(Z_{n}\right)\right) \mid X_{1}, \ldots, X_{n}\right\} \\
& \\
& =\frac{1}{2} \mathbf{P}\left\{X \neq X_{1}, \ldots, X \neq X_{n}, X \neq x_{V} \mid X_{1}, \ldots, X_{n}\right\} \\
& \geq \frac{1}{2} p \sum_{i=1}^{V-1} I_{\left\{x_{i} \neq X_{1}, \ldots, x_{i} \neq X_{n}\right\}} .
\end{aligned}
$$

For fixed $X_{1}, \ldots, X_{n}$, we denote by $J$ the collection $\left\{j: 1 \leq j \leq V-1, \cap_{i=1}^{n}\left\{X_{i} \neq\right.\right.$ $\left.x_{j}\right\}$. This is the collection of empty cells $x_{i}$. We summarize:

$$
\begin{aligned}
\sup _{(X, Y): L c=0} \mathbf{P}\left\{L_{n} \geq \epsilon\right\} & \geq \mathbf{P}\left\{\frac{1}{2} p \sum_{i=1}^{V-1} I_{\left\{x_{i} \neq X_{1}, \ldots, x_{i} \neq X_{n}\right\}} \geq \epsilon\right\} \\
& \geq \mathbf{P}\{|J| \geq 2 \epsilon / p\} .
\end{aligned}
$$

We consider two choices for $p$ :
Choice A. Take $p=1 / n$, and assume $12 n \epsilon \leq V-1, \epsilon<1 / 2$. Note that for $n \geq 15, \mathbf{E}|J|=(V-1)(1-p)^{n} \geq(V-1) / 3$. Also, since $0 \leq|J| \leq V-1$, we have Var $|J| \leq(V-1)^{2} / 4$. By the Chebyshev-Cantelli inequality (Theorem A.17),

$$
\begin{aligned}
\mathbf{P}\{|J| \geq 2 n \epsilon\} & =1-\mathbf{P}\{|J|<2 n \epsilon\} \\
& \geq 1-\mathbf{P}\{|J|<(V-1) / 6\} \\
& =1-\mathbf{P}\{|J|-\mathbf{E}|J| \leq-(V-1) / 6\}
\end{aligned}
$$

$$
\begin{aligned}
& \geq 1-\frac{\operatorname{Var}|J|}{\operatorname{Var}|J|+(V-1)^{2} / 36} \\
& \geq 1-\frac{(V-1)^{2} / 4}{(V-1)^{2} / 4+(V-1)^{2} / 36} \\
& =\frac{1}{10} .
\end{aligned}
$$

This proves the second inequality for $\sup \mathbf{P}\left\{L_{n} \geq \epsilon\right\}$.
Choice B. Assume that $\epsilon \leq 1 / 4$. By the pigeonhole principle, $|J| \geq 2 \epsilon / p$ if the number of points $X_{i}, 1 \leq i \leq n$, that are not equal to $x_{V}$ does not exceed $V-1-2 \epsilon / p$. Therefore, we have a further lower bound:

$$
\mathbf{P}\{|J| \geq 2 \epsilon / p\} \geq \mathbf{P}\{\operatorname{Binomial}(n,(V-1) p) \leq V-1-2 \epsilon / p\} .
$$

Define $v=\lceil(V-1) / 2\rceil$. Take $p=2 \epsilon / v$. By assumption, $n \geq 2 v-1$. Then the lower bound is

$$
\begin{aligned}
& \mathbf{P}\{\operatorname{Binomial}(n, 4 \epsilon) \leq v\} \\
& \geq\binom{n}{v}(4 \epsilon)^{v}(1-4 \epsilon)^{n-v} \\
& \geq \frac{1}{e \sqrt{2 \pi v}}\left(\frac{4 e \epsilon(n-v+1)}{v(1-4 \epsilon)}\right)^{v}(1-4 \epsilon)^{n} \\
& \quad\left(\text { since }\binom{n}{v} \geq\left(\frac{(n-v+1) e}{v}\right)^{v} \frac{1}{e \sqrt{2 \pi v}}\right. \text { by Stirling's formula) } \\
& \geq \frac{1}{e \sqrt{2 \pi v}}\left(\frac{4 e \epsilon(n-v+1)}{v}\right)^{v}(1-4 \epsilon)^{n} \\
& \geq \frac{1}{e \sqrt{2 \pi v}}\left(\frac{4 e n \epsilon}{v}\right)^{v}(1-4 \epsilon)^{n}\left(1-\frac{v-1}{n}\right)^{v} \\
& \geq \frac{1}{e \sqrt{2 \pi v}}\left(\frac{2 e n \epsilon}{v}\right)^{v} e^{-4 n \epsilon /(1-4 \epsilon)} \quad(\text { since } n \geq 2(v-1)) \\
& \quad(\text { use } 1-x \geq \exp (-x /(1-x))) .
\end{aligned}
$$

This concludes the proof.

### 14.3 Classes with Infinite VC Dimension

The results presented in the previous section may also be applied to classes with infinite vc dimension. For example, it is not hard to derive the following result from Theorem 14.1:

Theorem 14.3. Assume that $V_{\mathcal{C}}=\infty$. For every n, $\delta>0$ and classification rule $g_{n}$, there is a distribution with $L_{\mathcal{C}}=0$ such that

$$
\mathbf{E} L\left(g_{n}\right)>\frac{1}{2 e}-\delta
$$

For the proof, see Problem 14.2. Thus, when $V_{\mathcal{C}}=\infty$, distribution-free nontrivial performance guarantees for $L\left(g_{n}\right)-L_{\mathcal{C}}$ do not exist. This generalizes Theorem 7.1, where a similar result is shown if $\mathcal{C}$ is the class of all measurable discrimination functions. We have also seen in Theorem 7.2, that if $\mathcal{C}$ is the class of all measurable classifiers, then no universal rate of convergence exists. However, we will see in Chapter 18 that for some classes with infinite vc dimension, it is possible to find a classification rule such that $L\left(g_{n}\right)-L^{*} \leq c \sqrt{\log n / n}$ for any distribution such that the Bayes classifier is in $\mathcal{C}$. The constant $c$, however, necessarily depends on the distribution, as is apparent from Theorem 14.3.

Infinite vc dimension means that the class $\mathcal{C}$ shatters finite sets of any size. On the other hand, if $\mathcal{C}$ shatters infinitely many points, then no universal rate of convergence exists. This may be seen by an appropriate modification of Theorem 7.2, as follows. See Problem 14.3 for the proof.

Theorem 14.4. Let $\left\{a_{n}\right\}$ be a sequence of positive numbers converging to zero with $1 / 16 \geq a_{1} \geq a_{2} \geq \ldots$.. Let $\mathcal{C}$ be a class of classifiers with the property that there exists a set $A \subset \mathcal{R}^{d}$ of infinite cardinality such that for any subset $B$ of $A$, there exists $\phi \in \mathcal{C}$ such that $\phi(x)=1$ if $x \in B$ and $\phi(x)=0$ if $x \in A-B$. Then for every sequence of classification rules, there exists a distribution of $(X, Y)$ with $L_{\mathcal{C}}=0$, such that

$$
\mathbf{E} L_{n} \geq a_{n}
$$

for all $n$.
Note that the basic difference between this result and all others in this chapter is that in Theorem 14.4 the "bad" distribution does not vary with $n$. This theorem shows that selecting a classifier from a class shattering infinitely many points is essentially as hard as selecting one from the class of all classifiers.

### 14.4 The Case $L_{\mathcal{C}}>0$

In the more general case, when the best decision in the class $\mathcal{C}$ has positive error probability, the upper bounds derived in Chapter 12 for the expected error probability of the classifier obtained by minimizing the empirical risk are much larger than when $L_{\mathcal{C}}=0$. In this section we show that these upper bounds are necessarily large, and they may be tight for some distributions. Moreover, there is no other classifier that performs significantly better than empirical error minimization.

Theorem 14.5 below gives a lower bound for $\sup _{(X, Y): L_{\mathcal{C}} \text { fixed }} \mathbf{E} L\left(g_{n}\right)-L_{\mathcal{C}}$. As a function of $n$ and $V_{\mathcal{C}}$, the bound decreases basically as in the upper bound obtained from Theorem 12.6. Interestingly, the lower bound becomes smaller as $L_{\mathcal{C}}$
decreases, as should be expected. The bound is largest when $L_{\mathcal{C}}$ is close to $1 / 2$. The constants in the bound may be tightened at the expense of more complicated expressions. The theorem is essentially due to Devroye and Lugosi (1995), though the proof given here is different. Similar bounds-without making the dependence on $L_{\mathcal{C}}$ explicit—have been proved by Vapnik and Chervonenkis (1974c) and Simon (1993).

Theorem 14.5. Let $\mathcal{C}$ be a class of discrimination functions with vC dimension $V \geq 2$. Let $\mathcal{X}$ be the set of all random variables $(X, Y)$ for which for fixed $L \in(0,1 / 2)$,

$$
L=\inf _{g \in \mathcal{C}} \mathbf{P}\{g(X) \neq Y\}
$$

Then, for every discrimination rule $g_{n}$ based upon $X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}$,

$$
\sup _{(X, Y) \in \mathcal{X}} \mathbf{E}\left(L_{n}-L\right) \geq \sqrt{\frac{L(V-1)}{24 n}} e^{-8} \quad \text { if } n \geq \frac{V-1}{2 L} \max \left(9,1 /(1-2 L)^{2}\right)
$$

Proof. Again we consider the finite family $\mathcal{F}$ from the previous section. The notation $b$ and $B$ is also as above. $X$ now puts mass $p$ at $x_{i}, i<V$, and mass $1-(V-1) p$ at $x_{V}$. This imposes the condition $(V-1) p \leq 1$, which will be satisfied. Next introduce the constant $c \in(0,1 / 2)$. We no longer have $Y$ as a function of $X$. Instead, we have a uniform [0,1] random variable $U$ independent of $X$ and define

$$
Y= \begin{cases}1 & \text { if } U \leq \frac{1}{2}-c+2 c b_{i}, X=x_{i}, i<V \\ 0 & \text { otherwise }\end{cases}
$$

Thus, when $X=x_{i}, i<V, Y$ is 1 with probability $1 / 2-c$ or $1 / 2+c$. A simple argument shows that the best rule for $b$ is the one which sets

$$
f_{b}(x)= \begin{cases}1 & \text { if } x=x_{i}, i<V, b_{i}=1 \\ 0 & \text { otherwise }\end{cases}
$$

Also, observe that

$$
L=(V-1) p(1 / 2-c)
$$

Noting that $\left|2 \eta\left(x_{i}\right)-1\right|=c$ for $i<V$, for fixed $b$, by the equality in Theorem 2.2 , we may write

$$
L_{n}-L \geq \sum_{i=1}^{V-1} 2 p c I_{\left\{g_{n}\left(x_{i}, X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right)=1-f_{b}\left(x_{i}\right)\right\}}
$$

It is sometimes convenient to make the dependence of $g_{n}$ upon $b$ explicit by considering $g_{n}\left(x_{i}\right)$ as a function of $x_{i}, X_{1}, \ldots, X_{n}, U_{1}, \ldots, U_{n}$ (an i.i.d. sequence of uniform $[0,1]$ random variables), and $b_{i}$. The proof given here is based on
the ideas used in the proofs of Theorems 14.1 and 14.2. We replace $b$ by a uniformly distributed random $B$ over $\{0,1\}^{V-1}$. After this randomization, denote $Z_{n}=\left(X, X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right)$. Thus,

$$
\begin{aligned}
\sup _{(X, Y) \in \mathcal{F}} \mathbf{E}\left\{L_{n}-L\right\} & =\sup _{b} \mathbf{E}\left\{L_{n}-L\right\} \\
& \geq \mathbf{E}\left\{L_{n}-L\right\} \quad \text { (with random } B \text { ) } \\
& \geq \sum_{i=1}^{V-1} 2 p c \mathbf{E} I_{\left\{g_{n}\left(x_{i}, X_{1}, \ldots, Y_{n}\right)=1-f_{B}\left(x_{i}\right)\right\}} \\
& =2 c \mathbf{P}\left\{g_{n}\left(Z_{n}\right) \neq f_{B}(X)\right\} \\
& \geq 2 c L^{*}\left(Z_{n}, f_{B}(X)\right),
\end{aligned}
$$

where, as before, $L^{*}\left(Z_{n}, f_{B}(X)\right)$ denotes the Bayes probability of error of predicting the value of $f_{B}(X)$ based on observing $Z_{n}$. All we have to do is to find a suitable lower bound for

$$
L^{*}\left(Z_{n}, f_{B}(X)\right)=\mathbf{E}\left\{\min \left(\eta^{*}\left(Z_{n}\right), 1-\eta^{*}\left(Z_{n}\right)\right)\right\},
$$

where $\eta^{*}\left(Z_{n}\right)=\mathbf{P}\left\{f_{B}(X)=1 \mid Z_{n}\right\}$. Observe that

$$
\eta^{*}\left(Z_{n}\right)= \begin{cases}1 / 2 & \text { if } X \neq X_{1}, \ldots, X \neq X_{n} \text { and } X \neq x_{V} \\ \mathbf{P}\left\{B_{i}=1 \mid Y_{i_{1}}, \ldots, Y_{i_{k}}\right\} & \text { if } X=X_{i_{1}}=\cdots=X_{i_{k}}=x_{i}, i<V\end{cases}
$$

Next we compute $\mathbf{P}\left\{B_{i}=1 \mid Y_{i_{1}}=y_{1}, \ldots, Y_{i_{k}}=y_{k}\right\}$ for $y_{1}, \ldots, y_{k} \in\{0,1\}$. Denoting the numbers of zeros and ones by $k_{0}=\left|\left\{j \leq k: y_{j}=0\right\}\right|$ and $k_{1}=\mid\{j \leq$ $\left.k: y_{j}=1\right\} \mid$, we see that

$$
\begin{aligned}
& \mathbf{P}\left\{B_{i}=1 \mid Y_{i_{1}}=y_{1}, \ldots, Y_{i_{k}}=y_{k}\right\} \\
& \quad=\frac{(1-2 c)^{k_{1}}(1+2 c)^{k_{0}}}{(1-2 c)^{k_{1}}(1+2 c)^{k_{0}}+(1+2 c)^{k_{1}}(1-2 c)^{k_{0}}} .
\end{aligned}
$$

Therefore, if $X=X_{i_{1}}=\cdots=X_{i_{k}}=x_{i}, i<V$, then

$$
\begin{aligned}
& \min \left(\eta^{*}\left(Z_{n}\right), 1-\eta^{*}\left(Z_{n}\right)\right) \\
&=\frac{\min \left((1-2 c)^{k_{1}}(1+2 c)^{k_{0}},(1+2 c)^{k_{1}}(1-2 c)^{k_{0}}\right)}{(1-2 c)^{k_{1}}(1+2 c)^{k_{0}}+(1+2 c)^{k_{1}}(1-2 c)^{k_{0}}} \\
& \quad=\frac{\min \left(1,\left(\frac{1+2 c}{1-2 c}\right)^{k_{1}-k_{0}}\right)}{1+\left(\frac{1+2 c}{1-2 c}\right)^{k_{1}-k_{0}}} \\
& \quad=\frac{1}{1+\left(\frac{1+2 c}{1-2 c}\right)^{\left|k_{1}-k_{0}\right|}} .
\end{aligned}
$$

In summary, denoting $a=(1+2 c) /(1-2 c)$, we have

$$
\begin{aligned}
L^{*}\left(Z_{n}, f_{B}(X)\right)= & \mathbf{E}\left\{\frac{1}{1+a^{\mid \sum_{j: X_{j}=X}\left(2 Y_{j}-1\right)} \mid}\right\} \\
\geq & \mathbf{E}\left\{\frac{1}{2 a^{\mid \sum_{j: X_{j}=X}\left(2 Y_{j}-1\right)} \mid}\right\} \\
\geq & \frac{1}{2} \sum_{i=1}^{V-1} \mathbf{P}\left\{X=x_{i}\right\} \mathbf{E}\left\{a^{-\left|\sum_{j: X_{j}=x_{i}}\left(2 Y_{j}-1\right)\right|}\right\} \\
\geq & \frac{1}{2}(V-1) p a^{-\mathbf{E}}\left\{\left|\sum_{j: X_{j}=x_{i}}\left(2 Y_{j}-1\right)\right|\right\} \\
& \text { (by Jensen's inequality). }
\end{aligned}
$$

Next we bound $\mathbf{E}\left\{\left|\sum_{j: X_{j}=x_{i}}\left(2 Y_{j}-1\right)\right|\right\}$. Clearly, if $B(k, q)$ denotes a binomial random variable with parameters $k$ and $q$,

$$
\mathbf{E}\left\{\left|\sum_{j: X_{j}=x_{i}}\left(2 Y_{j}-1\right)\right|\right\}=\sum_{k=0}^{n}\binom{n}{k} p^{k}(1-p)^{n-k} \mathbf{E}\{|2 B(k, 1 / 2-c)-k|\}
$$

However, by straightforward calculation we see that

$$
\begin{aligned}
\mathbf{E}\{|2 B(k, 1 / 2-c)-k|\} & \leq \sqrt{\mathbf{E}\left\{(2 B(k, 1 / 2-c)-k)^{2}\right\}} \\
& =\sqrt{k\left(1-4 c^{2}\right)+4 k^{2} c^{2}} \\
& \leq 2 k c+\sqrt{k}
\end{aligned}
$$

Therefore, applying Jensen's inequality once again, we get

$$
\sum_{k=0}^{n}\binom{n}{k} p^{k}(1-p)^{n-k} \mathbf{E}\{|2 B(k, 1 / 2-c)-k|\} \leq 2 n p c+\sqrt{n p}
$$

Summarizing what we have obtained so far, we have

$$
\begin{aligned}
\sup _{b} \mathbf{E}\left\{L_{n}-L\right\} \geq & 2 c L^{*}\left(Z_{n}, f_{B}(X)\right) \\
\geq & 2 c \frac{1}{2}(V-1) p a^{-2 n p c-\sqrt{n p}} \\
\geq & c(V-1) p e^{-2 n p c(a-1)-(a-1) \sqrt{n p}} \\
& \left(\text { by the inequality } 1+x \leq e^{x}\right) \\
= & c(V-1) p e^{-8 n p c^{2} /(1-2 c)-4 c \sqrt{n p} /(1-2 c)} .
\end{aligned}
$$

A rough asymptotic analysis shows that the best asymptotic choice for $c$ is given by

$$
c=\frac{1}{\sqrt{4 n p}}
$$

Then the constraint $L=(V-1) p(1 / 2-c)$ leaves us with a quadratic equation in $c$. Instead of solving this equation, it is more convenient to take $c=\sqrt{(V-1) /(8 n L)}$. If $2 n L /(V-1) \geq 9$, then $c \leq 1 / 6$. With this choice for $c$, using $L=(V-1) p(1 / 2-$ $c$ ), straightforward calculation provides

$$
\sup _{(X, Y) \in \mathcal{F}} \mathbf{E}\left(L_{n}-L\right) \geq \sqrt{\frac{(V-1) L}{24 n}} e^{-8} .
$$

The condition $p(V-1) \leq 1$ implies that we need to ask that $n \geq(V-1) /(2 L(1-$ $2 L)^{2}$ ). This concludes the proof of Theorem 14.5. $\square$

Next we obtain a probabilistic bound. Its proof below is based upon Hellinger distances, and its methodology is essentially due to Assouad (1983b). For refinements and applications, we refer to Birgé (1983; 1986) and Devroye (1987).

Theorem 14.6. (Devroye and Lugosi (1995)). Let $\mathcal{C}$ be a class of discrimination functions with VC dimension $V \geq 2$. Let $\mathcal{X}$ be the set of all random variables $(X, Y)$ for which for fixed $L \in(0,1 / 4]$,

$$
L=\inf _{g \in \mathcal{C}} \mathbf{P}\{g(X) \neq Y\}
$$

Then, for every discrimination rule $g_{n}$ based upon $X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}$, and any $\epsilon \leq L$,

$$
\sup _{(X, Y) \in \mathcal{X}} \mathbf{P}\left\{L_{n}-L \geq \epsilon\right\} \geq \frac{1}{4} e^{-4 n \epsilon^{2} / L}
$$

PROOF. The method of randomization here is similar to that in the proof of Theorem 14.5. Using the same notation as there, it is clear that

$$
\begin{aligned}
& \sup _{(X, Y) \in \mathcal{X}} \mathbf{P}\left\{L_{n}-L \geq \epsilon\right\} \\
& \geq \mathbf{E} I_{\left\{\sum_{i=1}^{V-1} 2 p c I_{\left\{g n\left(x_{i}, X_{1}, \ldots, Y_{n}\right)=1-f_{B}\left(x_{i}\right)\right\}} \geq \epsilon\right\}} \\
& =2^{-(V-1)} \sum_{\substack{\left.\left(x_{1}^{\prime} 1, \ldots, x_{n}^{\prime}, y_{1}, \ldots, y_{n}\right) \\
\epsilon\left(\mid x_{1}, \ldots, x_{V}\right) \times(0,1)\right)^{n}}} \\
& \sum_{b} I_{\left\{\sum_{i=1}^{V-1} 2 p c I_{\left\{g_{n}\left(x_{i} ; x_{1}^{\prime}, y_{1}, \ldots, x_{n}^{\prime}, y_{n}\right)=1-f_{b}\left(x_{i}\right)\right\}} \geq \epsilon\right\}} \prod_{j=1}^{n} p_{b}\left(x_{j}^{\prime}, y_{j}\right) .
\end{aligned}
$$

First observe that if

$$
\begin{equation*}
\epsilon /(2 p c) \leq(V-1) / 2 \tag{14.1}
\end{equation*}
$$

then

$$
\begin{aligned}
& I_{\left\{\sum_{i=1}^{V-1} 2 p c I_{\left\{g n\left(x_{i} ; x_{1}^{\prime}, y_{1}, \ldots x_{n}^{\prime}, y_{n}\right)=1-f_{b}\left(x_{i}\right)\right]} \geq \epsilon\right\}} \\
& \left.\quad+\quad I_{\left\{\sum_{i=1}^{V-1} 2 p c I_{\{g n}\left(x_{i} ; x_{1}^{\prime}, y_{1}, \ldots, x_{n}^{\prime}, y_{n}\right)=1-f_{b}\left(x_{i}\right)\right\}} \geq \epsilon\right\}
\end{aligned}
$$

where $b^{c}$ denotes the binary vector $\left(1-b_{1}, \ldots, 1-b_{V-1}\right)$, that is, the complement of $b$. Therefore, for $\epsilon \leq p c(V-1)$, the last expression in the lower bound above is bounded from below by

$$
\begin{aligned}
& \frac{1}{2^{V-1}} \sum_{\substack{\left.\left(x_{1}^{\prime}, \ldots, x_{n}^{\prime}, y_{1}, \ldots, y_{n}\right) \\
\epsilon\left(x_{1}, \ldots, x_{V}\right] \times(0,1)\right)^{n}}} \sum_{b} \frac{1}{2} \min \left(\prod_{j=1}^{n} p_{b}\left(x_{j}^{\prime}, y_{j}\right), \prod_{j=1}^{n} p_{b^{c}}\left(x_{j}^{\prime}, y_{j}\right)\right) \\
& \geq \frac{1}{2^{V+1}} \sum_{b}\left(\sum_{\substack{\left(x_{1}^{\prime}, \ldots, x_{n}^{\prime}, y_{1}, \ldots, y_{n}\right) \\
\epsilon\left(x_{1}, \ldots, x_{V}\right] \times\{0,1)^{n}}} \sqrt{\left.\prod_{j=1}^{n} p_{b}\left(x_{j}^{\prime}, y_{j}\right) \times \prod_{j=1}^{n} p_{b^{c}}\left(x_{j}^{\prime}, y_{j}\right)\right)^{2}}\right.
\end{aligned}
$$

(by LeCam's inequality, Lemma 3.1)

$$
=\frac{1}{2^{V+1}} \sum_{b}\left(\sum_{(x, y)} \sqrt{p_{b}(x, y) p_{b^{c}}(x, y)}\right)^{2 n} .
$$

It is easy to see that for $x=x_{V}$,

$$
p_{b}(x, y)=p_{b^{c}}(x, y)=\frac{1-(V-1) p}{2}
$$

and for $x=x_{i}, i<V$,

$$
p_{b}(x, y) p_{b^{c}}(x, y)=p^{2}\left(\frac{1}{4}-c^{2}\right)
$$

Thus, we have the equality

$$
\sum_{(x, y)} \sqrt{p_{b}(x, y) p_{b^{c}}(x, y)}=1-(V-1) p+2(V-1) p \sqrt{\frac{1}{4}-c^{2}}
$$

Summarizing, since $L=p(V-1)(1 / 2-c)$, we have

$$
\begin{aligned}
\sup _{(X, Y) \in \mathcal{X}} \mathbf{P}\left\{L_{n}-L \geq \epsilon\right\} & \geq \frac{1}{4}\left(1-\frac{L}{\frac{1}{2}-c}\left(1-\sqrt{1-4 c^{2}}\right)\right)^{2 n} \\
& \geq \frac{1}{4}\left(1-\frac{L}{\frac{1}{2}-c} 4 c^{2}\right)^{2 n} \\
& \geq \frac{1}{4} \exp \left(-\frac{16 n L c^{2}}{1-2 c} /\left(1-\frac{8 L c^{2}}{1-2 c}\right)\right)
\end{aligned}
$$

where we used the inequality $1-x \geq e^{-x /(1-x)}$ again. We may choose $c$ as $\frac{\epsilon}{2 L+2 \epsilon}$. It is easy to verify that condition (14.1) holds. Also, $p(V-1) \leq 1$. From the
condition $L \geq \epsilon$ we deduce that $c \leq 1 / 4$. The exponent in the expression above may be bounded as

$$
\begin{aligned}
\frac{\frac{16 n L c^{2}}{1-2 c}}{1-\frac{8 L c^{2}}{1-2 c}} & =\frac{16 n L c^{2}}{1-2 c-8 L c^{2}} \\
& =\frac{\frac{4 n \epsilon^{2}}{L+\epsilon}}{1-\frac{2 \epsilon^{2}}{L+\epsilon}} \quad(\text { by substituting } c=\epsilon /(2 L+2 \epsilon)) \\
& \leq \frac{4 n \epsilon^{2}}{L} .
\end{aligned}
$$

Thus,

$$
\sup _{(X, Y) \in \mathcal{X}} \mathbf{P}\left\{L_{n}-L \geq \epsilon\right\} \geq \frac{1}{4} \exp \left(-4 n \epsilon^{2} / L\right)
$$

as desired.

### 14.5 Sample Complexity

We may rephrase the probability bounds above in terms of the sample complexity of algorithms for selecting a classifier from a class. Recall that for given $\epsilon, \delta>0$, the sample complexity of a selection rule $g_{n}$ is the smallest integer $N(\epsilon, \delta)$ such that

$$
\sup \mathbf{P}\left\{L\left(g_{n}\right)-L_{\mathcal{C}} \geq \epsilon\right\} \leq \delta
$$

for all $n \geq N(\epsilon, \delta)$. The supremum is taken over a class of distributions of ( $X, Y$ ). Here we are interested in distributions such that $L_{\mathcal{C}}$ is fixed.

We start with the case $L_{\mathcal{C}}=0$, by checking the implications of Theorem 14.2 for the sample complexity $N(\epsilon, \delta)$. First Blumer, Ehrenfeucht, Haussler, and Warmuth (1989) showed that for any algorithm,

$$
N(\epsilon, \delta) \geq C\left(\frac{1}{\epsilon} \log \left(\frac{1}{\delta}\right)+V_{\mathcal{C}}\right)
$$

where $C$ is a universal constant. In Ehrenfeucht, Haussler, Kearns, and Valiant (1989), the lower bound was partially improved to

$$
N(\epsilon, \delta) \geq \frac{V_{\mathcal{C}}-1}{32 \epsilon}
$$

when $\epsilon \leq 1 / 8$ and $\delta \leq 1 / 100$. It may be combined with the previous bound. Theorem 14.2 provides the following bounds:

Corollary 14.1. Let $\mathcal{C}$ be a class of discrimination functions with vc dimension $V \geq 2$. Let $\mathcal{X}$ be the set of all random variables $(X, Y)$ for which $L_{\mathcal{C}}=0$. Assume
$\epsilon \leq 1 / 4$, and denote $v=\lceil(V-1) / 2\rceil$. Then for every discrimination rule $g_{n}$ based upon $X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}$, when $\epsilon \leq 1 / 8$ and

$$
\log \left(\frac{1}{\delta}\right) \geq\left(\frac{4 v}{e}\right)(e \sqrt{2 \pi v})^{1 / v}
$$

then

$$
N(\epsilon, \delta) \geq \frac{1}{8 \epsilon} \log \left(\frac{1}{\delta}\right)
$$

Finally, for $\delta \leq 1 / 10$, and $\epsilon<1 / 2$,

$$
N(\epsilon, \delta) \geq \frac{V-1}{12 \epsilon}
$$

Proof. The second bound follows trivially from the second inequality of Theorem 14.2. By the first inequality there,

$$
\begin{aligned}
\sup _{(X, Y) \in \mathcal{X}} \mathbf{P}\left\{L_{n} \geq \epsilon\right\} & \geq \frac{1}{e \sqrt{2 \pi v}}\left(\frac{2 n e \epsilon}{V-1}\right)^{(V-1) / 2} e^{-4 n \epsilon /(1-4 \epsilon)} \\
& \geq \frac{1}{e \sqrt{2 \pi v}}\left(\frac{2 e n \epsilon}{v}\right)^{v} e^{-8 n \epsilon} \quad(\text { since } \epsilon \leq 1 / 8) \\
& \geq \frac{(8 \epsilon)^{v}}{\log ^{v}(1 / \delta)} n^{v} e^{-8 n \epsilon}
\end{aligned}
$$

(since we assume $\left.\log \left(\frac{1}{\delta}\right) \geq\left(\frac{4 v}{e}\right)(e \sqrt{2 \pi v})^{1 / v}\right)$.
The function $n^{v} e^{-8 n \epsilon}$ varies unimodally in $n$, and achieves a peak at $n=v /(8 \epsilon)$. For $n$ below this threshold, by monotonicity, we apply the bound at $n=v /(8 \epsilon)$. It is easy to verify that the value of the bound at $v /(8 \epsilon)$ is always at least $\delta$. If on the other hand, $(1 / 8 \epsilon) \log (1 / \delta) \geq n \geq v /(8 \epsilon)$, the lower bound achieves its minimal value at $(1 / 8 \epsilon) \log (1 / \delta)$, and the value there is $\delta$. This proves the first bound.

Corollary 14.1 shows that for any classifier, at least

$$
\max \left(\frac{1}{8 \epsilon} \log \left(\frac{1}{\delta}\right), \frac{V_{\mathcal{C}}-1}{4 \epsilon}\right)
$$

training samples are necessary to achieve $\epsilon$ accuracy with $\delta$ confidence for all distributions. Apart from a $\log \left(\frac{1}{\epsilon}\right)$ factor, the order of magnitude of this expression is the same as that of the upper bound for empirical error minimization, obtained in Corollary 12.4. That the upper and lower bounds are very close, has two important messages. On the one hand it gives a very good estimate for the number of training samples needed to achieve a certain performance. On the other hand it shows that there is essentially no better method than minimizing the empirical error probability.

In the case $L_{C}>0$, we may derive lower bounds for $N(\epsilon, \delta)$ from Theorems 14.5 and 14.6 :

Corollary 14.2. Let $\mathcal{G}$ be a class of discrimination functions with vC dimension $V \geq 2$. Let $\mathcal{X}$ be the set of all random variables $(X, Y)$ for which for fixed $L \in(0,1 / 2)$,

$$
L=\inf _{g \in \mathcal{G}} \mathbf{P}\{g(X) \neq Y\}
$$

Then, for every discrimination rule $g_{n}$ based upon $X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}$,

$$
N(\epsilon, \delta) \geq \frac{L(V-1) e^{-10}}{32} \times \min \left(\frac{1}{\delta^{2}}, \frac{1}{\epsilon^{2}}\right)
$$

Also, and in particular, for $\epsilon \leq L \leq 1 / 4$,

$$
N(\epsilon, \delta) \geq \frac{L}{4 \epsilon^{2}} \log \frac{1}{4 \delta} .
$$

Proof. The first bound may be obtained easily from the expectation-bound of Theorem 14.5 (see Problem 14.1). Setting the bound of Theorem 14.6 equal to $\delta$ provides the second bound on $N(\epsilon, \delta)$.

These bounds may of course be combined. They show that $N(\epsilon, \delta)$ is bounded from below by terms like $\left(1 / \epsilon^{2}\right) \log (1 / \delta)$ (independent of $V_{\mathcal{C}}$ ) and $\left(V_{\mathcal{C}}-1\right) / \epsilon^{2}$, as $\delta$ is typically much smaller than $\epsilon$. By comparing these bounds with the upper bounds of Corollary 12.3 , we see that the only difference between the orders of magnitude is a $\log (1 / \epsilon)$-factor, so all remarks made for the case $L_{\mathcal{C}}=0$ remain valid.

Interestingly, all bounds depend on the class $\mathcal{C}$ only through its vC dimension. This fact suggests that when studying distribution-free properties of $L\left(g_{n}\right)-L_{\mathcal{C}}$, the VC dimension is the most important characteristic of the class. Also, all bounds are linear in the VC dimension, which links it conveniently to sample size.

REMARK. It is easy to see from the proofs that all results remain valid if we allow randomization in the rules $g_{n}$. $\square$

## Problems and Exercises

Problem 14.1. Show that Theorem 14.5 implies that for every discrimination rule $g_{n}$ based upon $D_{n}$,

$$
N(\epsilon, \delta) \geq \frac{L(V-1) e^{-10}}{32} \times \min \left(\frac{1}{\delta^{2}}, \frac{1}{\epsilon^{2}}\right) .
$$

Hint: Assume that $\mathbf{P}\left\{L_{n}-L>\epsilon\right\}<\delta$. Then clearly, $\mathbf{E}\left\{L_{n}-L\right\} \leq \epsilon+\delta$.
Problem 14.2. Prove Theorem 14.3: First apply the proof method of Theorem 14.1 to show that for every $n$, and $g_{n}$, there is a distribution with $L_{\mathcal{C}}=0$ such that $\mathbf{E} L\left(g_{n}\right) \geq(1-1 / n) /(2 e)$. Use a monotonicity argument to finish the proof.

Problem 14.3. Prove Theorem 14.4 by modifying the proof of Theorem 7.2.

## 15

## The Maximum Likelihood Principle

In this chapter we explore the various uses of the maximum likelihood principle in discrimination. In general, the principle is only applicable if we have some a priori knowledge of the problem at hand. We offer definitions, consistency results, and examples that highlight the advantages and shortcomings.

### 15.1 Maximum Likelihood: The Formats

Sometimes, advance information takes a very specific form (e.g., "if $Y=1, X$ is normal ( $\mu, \sigma^{2}$ )"). Often, it is rather vague (e.g., "we believe that $X$ has a density," or " $\eta(x)$ is thought to be a monotone function of $x \in \mathcal{R}$ ").

If we have information in set format, the maximum likelihood principle is less appropriate. Here we know that the Bayes rule $g^{*}(x)$ is of the form $g(x)=I_{\{x \in A\}}$, where $A \in \mathcal{A}$ and $\mathcal{A}$ is a class of sets of $\mathcal{R}^{d}$. We refer to the chapters on empirical risk minimization (see Chapter 12 and also Chapter 18) for this situation.

If we know that the true (unknown) $\eta$ belongs to a class $\mathcal{F}$ of functions that map $\mathcal{R}^{d}$ to [0, 1], then we say that we are given information in regression format. With each $\eta^{\prime} \in \mathcal{F}$ we associate a set $A=\left\{x: \eta^{\prime}(x)>1 / 2\right\}$ and a discrimination rule $g(x)=I_{\{x \in A\}}$. The class of these rules is denoted by $\mathcal{C}$. Assume that we somehow could estimate $\eta$ by $\eta_{n}$. Then it makes sense to use the associated rule $g_{n}(x)=$ $I_{\left\{\eta_{n}(x)>1 / 2\right\}}$. The maximum likelihood method suggests a way of picking the $\eta_{n}$ from $\mathcal{F}$ that in some sense is most likely given the data. It is fully automatic-the user does not have to pick any parameters-but it does require a serious implementation effort in many cases. In a sense, the regression format is more powerful than the
set format, as there is more information in knowing a function $\eta$ than in knowing the indicator function $I_{\{\eta>1 / 2]}$. Still, no structure is assumed on the part of $X$, and none is needed to obtain consistency results.
A third format, even more detailed, is that in which we know that the distribution of ( $X, Y$ ) belongs to a class $\mathcal{D}$ of distributions on $\mathcal{R}^{d} \times\{0,1\}$. For a given distribution, we know $\eta$, so we may once again deduce a rule $g$ by setting $g(x)=I_{\{\eta(x)>1 / 2]}$. This distribution format is even more powerful, as the positions $X_{1}, \ldots, X_{n}$ alone in some cases may determine the unknown parameters in the model. This situation fits in squarely with classical parameter estimation in mathematical statistics. Once again, we may apply the maximum likelihood principle to select a distribution from $\mathcal{D}$. Unfortunately, as we move to more restrictive and stronger formats, the number of conditions under which the maximum likelihood principle is consistent increases as well. We will only superficially deal with the distribution format (see Chapter 16 for more detail).

### 15.2 The Maximum Likelihood Method: Regression Format

Given $X_{1}, \ldots, X_{n}$, the probability of observing $Y_{1}=y_{1}, \ldots, Y_{n}=y_{n}$ is

$$
\prod_{i=1}^{n} \eta\left(X_{i}\right)^{y_{i}}\left(1-\eta\left(X_{i}\right)\right)^{1-y_{i}}
$$

If $\eta$ is unknown but belongs to a family $\mathcal{F}$ of functions, we may wish to select that $\eta^{\prime}$ from $\mathcal{F}$ (if it exists) for which that likelihood product is maximal. More formally, we select $\eta^{\prime}$ so that the logarithm

$$
\mathcal{L}_{n}\left(\eta^{\prime}\right)=\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i} \log \eta^{\prime}\left(X_{i}\right)+\left(1-Y_{i}\right) \log \left(1-\eta^{\prime}\left(X_{i}\right)\right)\right)
$$

is maximal. If the family $\mathcal{F}$ is too rich, this will overfit, and consequently, the selected function $\eta_{n}$ has a probability of error

$$
L\left(\eta_{n}\right)=\mathbf{P}\left\{I_{\left\{\eta_{n}(X)>1 / 2\right\}} \neq Y \mid D_{n}\right\}
$$

that does not tend to $L^{*}$. For convenience, we assume here that there exists an element of $\mathcal{F}$ maximizing $\mathcal{L}_{n}$.

We do not assume here that the class $\mathcal{F}$ is very small. Classes in which each $\eta^{\prime}$ in $\mathcal{F}$ is known up to one or a few parameters are loosely called parametric. Sometimes $\mathcal{F}$ is defined via a generic description such as: $\mathcal{F}$ is the class of all $\eta^{\prime}: \mathcal{R}^{d} \rightarrow[0,1]$ that are Lipschitz with constant $c$. Such classes are called nonparametric. In certain cases, the boundary between parametric and nonparametric is unclear. We will be occupied with the consistency question: does $L\left(\eta_{n}\right) \rightarrow \inf _{\eta^{\prime} \in \mathcal{F}} L\left(\eta^{\prime}\right)$ with probability one for all distribution of $(X, Y)$ ? (Here $L\left(\eta^{\prime}\right)=\mathbf{P}\left\{I_{\left\{\eta^{\prime}(X)>1 / 2\right\}} \neq Y\right\}$ is
the probability of error of the natural rule that corresponds to $\eta^{\prime}$.) If, additionally, $\mathcal{F}$ is rich enough or our prior information is good enough, we may have inf $\eta_{\eta^{\prime} \in \mathcal{F}} L\left(\eta^{\prime}\right)=$ $L^{*}$, but that is not our concern here, as $\mathcal{F}$ is given to us.

We will not be concerned with the computational problems related to the maximization of $\mathcal{L}_{n}\left(\eta^{\prime}\right)$ over $\mathcal{F}$. Gradient methods or variations of them are sometimes used-refer to McLachlan (1992) for a bibliography. It suffices to say that in simple cases, an explicit form for $\eta_{n}$ may be available. An example follows.

Our first lemma illustrates that the maximum likelihood method should only be used when the true (but unknown) $\eta$ indeed belongs to $\mathcal{F}$. Recall that the same was not true for empirical risk minimization over vc classes (see Chapter 12).

Lemma 15.1. Consider the class $\mathcal{F}$ with two functions $\eta_{A} \equiv 0.45, \eta_{B} \equiv 0.95$. Let $\eta_{n}$ be the function picked by the maximum likelihood method. There exists a distribution for $(X, Y)$ with $\eta \notin \mathcal{F}$ such that with probability one, as $n \rightarrow \infty$,

$$
L\left(\eta_{n}\right) \rightarrow \max _{\eta^{\prime} \in \mathcal{F}} L\left(\eta^{\prime}\right)>\min _{\eta^{\prime} \in \mathcal{F}} L\left(\eta^{\prime}\right)
$$

Thus, maximum likelihood picks the wrong classifier.
Proof. Define the distribution of $(X, Y)$ on $\{0,1\} \times\{0,1\}$ by $\mathbf{P}\{X=0, Y=0\}=$ $\mathbf{P}\{X=1, Y=0\}=2 / 9, \mathbf{P}\{X=0, Y=1\}=1 / 9$, and $\mathbf{P}\{X=1, Y=1\}=4 / 9$. Then one may quickly verify that

$$
L\left(\eta_{A}\right)=\frac{5}{9}, \quad L\left(\eta_{B}\right)=\frac{4}{9} .
$$

(Note that $L^{*}=1 / 3$, but this is irrelevant here.) Within $\mathcal{F}, \eta_{B}$ is the better for our distribution. By the strong law of large numbers, we have

$$
\mathcal{L}_{n}\left(\eta_{A}\right) \rightarrow \mathbf{E}\left\{Y \log \eta_{A}(X)+(1-Y) \log \left(1-\eta_{A}(X)\right)\right\}
$$

with probability one (and similarly for $\eta_{B}$ ). If one works out the values, it is seen that with probability one, $\eta_{n} \equiv \eta_{A}$ for all $n$ large enough. Hence, $L\left(\eta_{n}\right) \rightarrow$ $\max _{\eta^{\prime} \in \mathcal{F}} L\left(\eta^{\prime}\right)$ with probability one.

Remark. Besides the clear theoretical hazard of not capturing $\eta$ in $\mathcal{F}$, the maximum likelihood method runs into a practical problem with "infinity." For example, take $\mathcal{F}=\left\{\eta_{A} \equiv 0, \eta_{B} \equiv 1\right\}$, and assume $\eta \equiv 1 / 3$. For all $n$ large enough, both classes are represented in the data sample with probability one. This implies that $\mathcal{L}_{n}\left(\eta_{A}\right)=\mathcal{L}_{n}\left(\eta_{B}\right)=-\infty$. The maximum likelihood estimate $\eta_{n}$ is ill-defined, while any reasonable rule should quickly be able to pick $\eta_{A}$ over $\eta_{B}$.

The lemma shows that when $\eta \notin \mathcal{F}$, the maximum likelihood method is not even capable of selecting one of two choices. The situation changes dramatically when $\eta \in \mathcal{F}$. For finite classes $\mathcal{F}$, nothing can go wrong. Noting that whenever $\eta \in \mathcal{F}$, we have $L^{*}=\inf _{\eta^{\prime} \in \mathcal{F}} L\left(\eta^{\prime}\right)$, we may now expect that $L\left(\eta_{n}\right) \rightarrow L^{*}$ in probability, or with probability one.

Theorem 15.1. If $|\mathcal{F}|=k<\infty$, and $\eta \in \mathcal{F}$, then the maximum likelihood method is consistent, that is,

$$
L\left(\eta_{n}\right) \rightarrow L^{*} \quad \text { in probability } .
$$

Proof. For a fixed distribution of $(X, Y)$, we rank the members $\eta^{(1)}, \ldots, \eta^{(k)}$ of $\mathcal{F}$ by increasing values of $L\left(\eta^{(i)}\right)$. Put $\eta^{(1)} \equiv \eta$. Let $i_{0}$ be the largest index for which $L\left(\eta^{(i)}\right)=L^{*}$. Let $\eta_{n}$ be the maximum likelihood choice from $\mathcal{F}$. For any $\alpha$, we have

$$
\begin{aligned}
& \mathbf{P}\left\{L\left(\eta_{n}\right) \neq L^{*}\right\} \\
& \quad \leq \mathbf{P}\left\{\mathcal{L}_{n}\left(\eta^{(1)}\right) \leq \max _{i>i_{0}} \mathcal{L}_{n}\left(\eta^{(i)}\right)\right\} \\
& \quad \leq \mathbf{P}\left\{\mathcal{L}_{n}\left(\eta^{(1)}\right) \leq \alpha\right\}+\sum_{i>i_{0}} \mathbf{P}\left\{\mathcal{L}_{n}\left(\eta^{(i)}\right) \geq \alpha\right\}
\end{aligned}
$$

Define the entropy of $(X, Y)$ by

$$
\mathcal{E}=\mathcal{E}(\eta)=-\mathbf{E}\{\eta(X) \log \eta(X)+(1-\eta(X)) \log (1-\eta(X))\}
$$

(see Chapter 3). Recall that $0 \leq \mathcal{E} \leq \log 2$. We also need the negative divergences

$$
D_{i}=\mathbf{E}\left\{\eta(X) \log \frac{\eta^{(i)}(X)}{\eta(X)}+(1-\eta(X)) \log \frac{1-\eta^{(i)}(X)}{1-\eta(X)}\right\},
$$

which are easily seen to be nonpositive for all $i$ (by Jensen's inequality). Furthermore, $D_{i}=0$ if and only if $\eta^{(i)}(X)=\eta(X)$ with probability one. Observe that for $i>i_{0}$, we cannot have this. Let $\theta=\max _{i>i_{0}} D_{i}$. (If $i_{0}=k$, this set is empty, but then the theorem is trivially true.)

It is advantageous to take $\alpha=-\mathcal{E}+\theta / 2$. Observe that

$$
\mathbf{E}\left\{\mathcal{L}_{n}\left(\eta^{(i)}\right)\right\}=-\mathcal{E}+D_{i} \begin{cases}=-\mathcal{E} & \text { if } i=1 \\ \leq-\mathcal{E}+\theta & \text { if } i>i_{0}\end{cases}
$$

Thus,

$$
\mathbf{P}\left\{L\left(\eta_{n}\right) \neq L^{*}\right\} \leq \mathbf{P}\left\{\mathcal{L}_{n}\left(\eta^{(1)}\right) \leq-\mathcal{E}+\frac{\theta}{2}\right\}+\sum_{i>i_{0}} \mathbf{P}\left\{\mathcal{L}_{n}\left(\eta^{(i)}\right) \geq-\mathcal{E}+\frac{\theta}{2}\right\}
$$

By the law of large numbers, we see that both terms converge to zero. Note, in particular, that it is true even if for some $i>i_{0}, D_{i}=-\infty$.

For infinite classes, many things can go wrong. Assume that $\mathcal{F}$ is the class of all $\eta^{\prime}: \mathcal{R}^{2} \rightarrow[0,1]$ with $\eta^{\prime}=I_{A}$ and $A$ is a convex set containing the origin. Pick $X$ uniformly on the perimeter of the unit circle. Then the maximum likelihood estimate $\eta_{n}$ matches the data, as we may always find a closed polygon $P_{n}$ with
vertices at the $X_{i}$ 's with $Y_{i}=1$. For $\eta_{n}=I_{P_{n}}$, we have $\mathcal{L}_{n}\left(\eta_{n}\right)=0$ (its maximal value), yet $L\left(\eta_{n}\right)=\mathbf{P}\{Y=1\}=p$ and $L^{*}=0$. The class $\mathcal{F}$ is plainly too rich.

For distributions of $X$ on the positive integers, maximum likelihood does not behave as poorly, even though it must pick among infinitely many possibilities. Assume $\mathcal{F}$ is the class of all $\eta^{\prime}$, but we know that $X$ puts all its mass on the positive integers. Then maximum likelihood tries to maximize

$$
\prod_{i=1}^{\infty}\left(\eta^{\prime}(i)\right)^{N_{1, i}}\left(1-\eta^{\prime}(i)\right)^{N_{0, i}}
$$

over $\eta^{\prime} \in \mathcal{F}$, where $N_{1, i}=\sum_{j=1}^{n} I_{\left\{X_{j}=i, Y_{j}=1\right\}}$ and $N_{0, i}=\sum_{j=1}^{n} I_{\left\{X_{j}=i, Y_{j}=0\right\}}$. The maximization is to be done over $\operatorname{all}\left(\eta^{\prime}(1), \eta^{\prime}(2), \ldots\right)$ from $\otimes_{i=1}^{\infty}[0,1]$. Fortunately, this is turned into a maximization for each individual $i$-usually we are not so lucky. Thus, if $N_{0, i}+N_{1, i}=0$, we set $\eta_{n}(i)=0$ (arbitrarily), while if $N_{0, i}+N_{1, i}>0$, we pick $\eta_{n}(i)$ as the value $u$ that maximizes

$$
N_{1, i} \log u+N_{0, i} \log (1-u)
$$

Setting the derivative with respect to $u$ equal to zero shows that

$$
\eta_{n}(i)=\frac{N_{1, i}}{N_{1, i}+N_{0, i}}
$$

In other words, $\eta_{n}$ is the familiar histogram estimate with bin width less than one half. It is known to be universally consistent (see Theorem 6.2). Thus, maximum likelihood may work for large $\mathcal{F}$ if we restrict the distribution of $(X, Y)$ a bit.

### 15.3 Consistency

Finally, we are ready for the main consistency result for $\eta_{n}$ when $\mathcal{F}$ may have infinitely many elements. The conditions of the theorem involve the bracketing metric entropy of $\mathcal{F}$, defined as follows: for every distribution of $X$, and $\epsilon>0$, let $\mathcal{F}_{X, \epsilon}$ be a set of functions such that for each $\eta^{\prime} \in \mathcal{F}$, there exist functions $\eta_{L}^{\prime}, \eta_{U}^{\prime} \in \mathcal{F}_{X, \epsilon}$ such that for all $x \in \mathcal{R}^{d}$

$$
\eta_{L}^{\prime}(x) \leq \eta^{\prime}(x) \leq \eta_{U}^{\prime}(x)
$$

and

$$
\mathbf{E}\left\{\eta_{U}^{\prime}(X)-\eta_{L}^{\prime}(X)\right\} \leq \epsilon
$$

That is, every $\eta^{\prime} \in \mathcal{F}$ is "bracketed" between two members of $\mathcal{F}_{X, \epsilon}$ whose $L_{1}(\mu)$ distance is not larger than $\epsilon$. Let $N(X, \epsilon)$ denote the cardinality of the smallest such $\mathcal{F}_{X, \epsilon}$. If $N(X, \epsilon)<\infty$,

$$
\log N(X, \epsilon)
$$

is called the bracketing $\epsilon$-entropy of $\mathcal{F}$, corresponding to $X$.

Theorem 15.2. Let $\mathcal{F}$ be a class of regression functions $\mathcal{R}^{d} \rightarrow[0,1]$. Assume that for every distribution of $X$ and $\epsilon>0, N(X, \epsilon)<\infty$. Then the maximum likelihood choice $\eta_{n}$ satisfies

$$
\lim _{n \rightarrow \infty} L\left(\eta_{n}\right)=L^{*} \quad \text { in probability }
$$

for all distributions of $(X, Y)$ with $\eta \in \mathcal{F}$.
Thus, consistency is guaranteed if $\mathcal{F}$ has a finite bracketing $\epsilon$-entropy for every $X$ and $\epsilon$. We provide examples in the next section. For the proof, first we need a simple corollary of Lemma 3.2:

Corollary 15.1. Let $\eta, \eta^{\prime}: \mathcal{R}^{d} \rightarrow[0,1]$, and let $(X, Y)$ be an $\mathcal{R}^{d} \times\{0,1\}$-valued random variable pair with $\mathbf{P}\{Y=1 \mid X=x\}=\eta(x)$. Define

$$
L^{*}=\mathbf{E}\{\min (\eta(X), 1-\eta(X))\}, \quad L\left(\eta^{\prime}\right)=\mathbf{P}\left\{I_{\left\{\eta^{\prime}(X)>1 / 2\right\}} \neq Y\right\}
$$

Then

$$
\begin{aligned}
\mathbf{E}\left\{\sqrt{\eta^{\prime}(X) \eta(X)}+\sqrt{\left(1-\eta^{\prime}(X)\right)(1-\eta(X))}\right\} & \leq 1-\frac{\left(2 \mathbf{E}\left|\eta(X)-\eta^{\prime}(X)\right|\right)^{2}}{8} \\
& \leq 1-\frac{1}{8}\left(L\left(\eta^{\prime}\right)-L^{*}\right)^{2}
\end{aligned}
$$

Proof. The first inequality follows by Lemma 3.2, and the second by Theorem 2.2.

Now, we are ready to prove the theorem. Some of the ideas used here appear in Wong and Shen (1992).
Proof of Theorem 15.2. Look again at the proof for the case $|\mathcal{F}|<\infty$ (Theorem 15.1). Define $\mathcal{E}$ as there. Let $\mathcal{F}^{(\epsilon)}$ be the collection of those $\eta^{\prime} \in \mathcal{F}$ with $L\left(\eta^{\prime}\right)>$ $L^{*}+\epsilon$, recalling that $L\left(\eta^{\prime}\right)=\mathbf{P}\left\{I_{\left\{\eta^{\prime}(X)>1 / 2\right\}} \neq Y\right\}$. For every $\alpha$,

$$
\begin{aligned}
\mathbf{P}\left\{L\left(\eta_{n}\right)-L^{*}>\epsilon\right\} & \leq \mathbf{P}\left\{\sup _{\eta^{\prime} \in \mathcal{F}^{(\epsilon)}} \mathcal{L}_{n}\left(\eta^{\prime}\right) \geq \mathcal{L}_{n}(\eta)\right\} \\
& \leq \mathbf{P}\left\{\mathcal{L}_{n}(\eta) \leq \alpha\right\}+\mathbf{P}\left\{\sup _{\eta^{\prime} \in \mathcal{F}(\epsilon)} \mathcal{L}_{n}\left(\eta^{\prime}\right) \geq \alpha\right\}
\end{aligned}
$$

For reasons that will be obvious later, we take $\alpha=-\mathcal{E}-\epsilon^{2} / 16$. Thus,

$$
\begin{aligned}
& \mathbf{P}\left\{L\left(\eta_{n}\right)-L^{*}>\epsilon\right\} \\
& \quad \leq \mathbf{P}\left\{\mathcal{L}_{n}(\eta) \leq-\mathcal{E}-\frac{\epsilon^{2}}{16}\right\}+\mathbf{P}\left\{\sup _{\eta^{\prime} \in \mathcal{F}^{(\epsilon)}} \mathcal{L}_{n}\left(\eta^{\prime}\right) \geq-\mathcal{E}-\frac{\epsilon^{2}}{16}\right\}
\end{aligned}
$$

$$
\begin{aligned}
= & \mathbf{P}\left\{\mathcal{L}_{n}(\eta) \leq-\mathcal{E}-\frac{\epsilon^{2}}{16}\right\} \\
& +\mathbf{P}\left\{\left(\sup _{\eta^{\prime} \in \mathcal{F}^{(\epsilon)}} \mathcal{L}_{n}\left(\eta^{\prime}\right)-\mathcal{L}_{n}(\eta)\right)+\left(\mathcal{L}_{n}(\eta)+\mathcal{E}\right) \geq \frac{\epsilon^{2}}{16}-\frac{\epsilon^{2}}{8}\right\} \\
\leq & \mathbf{P}\left\{\left|\mathcal{L}_{n}(\eta)+\mathcal{E}\right| \geq \frac{\epsilon^{2}}{16}\right\}+\mathbf{P}\left\{\sup _{\eta^{\prime} \in \mathcal{F}^{(\epsilon)}} \mathcal{L}_{n}\left(\eta^{\prime}\right)-\mathcal{L}_{n}(\eta) \geq-\frac{\epsilon^{2}}{8}\right\}
\end{aligned}
$$

Noting that $\mathbf{E}\left\{\mathcal{L}_{n}(\eta)\right\}=-\mathcal{E}$, the law of large numbers implies that the first term on the right-hand side converges to zero for every $\epsilon$. (See Problem 15.4 for more information.)

Next we bound the second term. For a fixed distribution, let $\mathcal{F}_{X, \delta}$ be the smallest set of functions such that for each $\eta^{\prime} \in \mathcal{F}$, there exist $\eta_{L}^{\prime}, \eta_{U}^{\prime} \in \mathcal{F}_{X, \delta}$ with

$$
\eta_{L}^{\prime}(x) \leq \eta^{\prime}(x) \leq \eta_{U}^{\prime}(x), \quad x \in \mathcal{R}^{d}
$$

and

$$
\mathbf{E}\left\{\eta_{U}^{\prime}(X)-\eta_{L}^{\prime}(X)\right\} \leq \delta
$$

where $\delta>0$ will be specified 1ater. By assumption, $N(X, \delta)=\left|\mathcal{F}_{X, \delta}\right|<\infty$. We have,

$$
\begin{aligned}
& \mathbf{P}\left\{\sup _{\eta^{\prime} \in \mathcal{F}(\epsilon)} \mathcal{L}_{n}\left(\eta^{\prime}\right)-\mathcal{L}_{n}(\eta) \geq-\frac{\epsilon^{2}}{8}\right\} \\
& \quad=\mathbf{P}\left\{\sup _{\eta^{\prime} \in \mathcal{F}(\epsilon)} \prod_{i=1}^{n} \frac{\eta^{\prime}\left(X_{i}\right)^{Y_{i}}\left(1-\eta^{\prime}\left(X_{i}\right)\right)^{1-Y_{i}}}{\eta\left(X_{i}\right)^{Y_{i}}\left(1-\eta\left(X_{i}\right)\right)^{1-Y_{i}}} \geq e^{-n \epsilon^{2} / 8}\right\} \\
& \quad \leq \mathbf{P}\left\{\sup _{\eta^{\prime} \in \mathcal{F}(\epsilon)} \prod_{i=1}^{n} \frac{\eta_{U}^{\prime}\left(X_{i}\right)^{Y_{i}}\left(1-\eta_{L}^{\prime}\left(X_{i}\right)\right)^{1-Y_{i}}}{\eta\left(X_{i}\right)^{Y_{i}}\left(1-\eta\left(X_{i}\right)\right)^{1-Y_{i}}} \geq e^{-n \epsilon^{2} / 8}\right\}
\end{aligned}
$$

$$
\text { (where } \eta_{L}^{\prime}, \eta_{U}^{\prime} \in \mathcal{F}_{X, \delta}, \eta_{L}^{\prime} \leq \eta^{\prime} \leq \eta_{U}^{\prime} \text { and } \mathbf{E}\left\{\eta_{U}^{\prime}(X)-\eta_{L}^{\prime}(X)\right\}<\delta \text { ) }
$$

$$
\leq \quad N^{2}(X, \delta) \sup _{\eta^{\prime} \in \mathcal{F}^{(\epsilon)}} \mathbf{P}\left\{\prod_{i=1}^{n} \frac{\eta_{U}^{\prime}\left(X_{i}\right)^{Y_{i}}\left(1-\eta_{L}^{\prime}\left(X_{i}\right)\right)^{1-Y_{i}}}{\eta\left(X_{i}\right)^{Y_{i}}\left(1-\eta\left(X_{i}\right)\right)^{1-Y_{i}}} \geq e^{-n \epsilon^{2} / 8}\right\}
$$

(by the union bound)

$$
\begin{aligned}
= & N^{2}(X, \delta) \sup _{\eta^{\prime} \in \mathcal{F}(\epsilon)} \mathbf{P}\left\{\prod_{i=1}^{n} \sqrt{\frac{\eta_{U}^{\prime}\left(X_{i}\right)^{Y_{i}}\left(1-\eta_{L}^{\prime}\left(X_{i}\right)\right)^{1-Y_{i}}}{\eta\left(X_{i}\right)^{Y_{i}}\left(1-\eta\left(X_{i}\right)\right)^{1-Y_{i}}}} \geq e^{-n \epsilon^{2} / 16}\right\} \\
\leq & N^{2}(X, \delta) \\
& \times \sup _{\eta^{\prime} \in \mathcal{F}(\epsilon)}\left(\mathbf{E}\left\{\sqrt{\eta_{U}^{\prime}(X) \eta(X)}+\sqrt{\left(1-\eta_{L}^{\prime}(X)\right)(1-\eta(X))}\right\}\right)^{n} e^{n \epsilon^{2} / 16},
\end{aligned}
$$

where in the last step we used Markov's inequality and independence. We now find a good bound for

$$
\mathbf{E}\left\{\sqrt{\eta_{U}^{\prime}(X) \eta(X)}+\sqrt{\left(1-\eta_{L}^{\prime}(X)\right)(1-\eta(X))}\right\}
$$

for each $\eta^{\prime} \in \mathcal{F}^{(\epsilon)}$ as follows:

$$
\begin{aligned}
& \mathbf{E}\left\{\sqrt{\eta_{U}^{\prime}(X) \eta(X)}+\sqrt{\left(1-\eta_{L}^{\prime}(X)\right)(1-\eta(X))}\right\} \\
& \\
& \leq \quad \mathbf{E}\left\{\sqrt{\eta^{\prime}(X) \eta(X)}+\sqrt{\left(1-\eta^{\prime}(X)\right)(1-\eta(X))}\right. \\
& \\
& \\
& \left.\quad+\sqrt{\eta(X)\left(\eta_{U}^{\prime}(X)-\eta^{\prime}(X)\right)}+\sqrt{(1-\eta(X))\left(\eta^{\prime}(X)-\eta_{L}^{\prime}(X)\right)}\right\}
\end{aligned}
$$

$$
\text { (here we used } \sqrt{a+b} \leq \sqrt{a}+\sqrt{b} \text { ) }
$$

$$
\leq \mathbf{E}\left\{\sqrt{\eta^{\prime}(X) \eta(X)}+\sqrt{\left(1-\eta^{\prime}(X)\right)(1-\eta(X))}\right\}
$$

$$
+\sqrt{\mathbf{E}\{\eta(X)\}} \sqrt{\mathbf{E}\left\{\eta_{U}^{\prime}(X)-\eta^{\prime}(X)\right\}}
$$

$$
+\sqrt{\mathbf{E}\{1-\eta(X)\}} \sqrt{\mathbf{E}\left\{\eta^{\prime}(X)-\eta_{L}^{\prime}(X)\right\}}
$$

(by the Cauchy-Schwarz inequality)
$\leq \mathbf{E}\left\{\sqrt{\eta^{\prime}(X) \eta(X)}+\sqrt{\left(1-\eta^{\prime}(X)\right)(1-\eta(X))}\right\}+2 \sqrt{\delta}$
$\leq 1-\frac{\epsilon^{2}}{8}+2 \sqrt{\delta} \quad$ (by Corollary 15.1)
Summarizing, we obtain

$$
\begin{aligned}
& \mathbf{P}\left\{\sup _{\eta^{\prime} \in \mathcal{F}(\epsilon)} \mathcal{L}_{n}\left(\eta^{\prime}\right)-\mathcal{L}_{n}(\eta) \geq-\frac{\epsilon^{2}}{8}\right\} \\
& \quad \leq N^{2}(X, \delta)\left(1-\frac{\epsilon^{2}}{8}+2 \sqrt{\delta}\right)^{n} e^{n \epsilon^{2} / 16} \\
& \quad \leq N^{2}(X, \delta) e^{-n \epsilon^{2} / 16+2 n \sqrt{\delta}} \quad\left(\text { since } 1+x \leq e^{x}\right)
\end{aligned}
$$

By taking $\delta<\left(\frac{\epsilon^{2}}{32}\right)^{2}$, we see that the probability converges to zero exponentially rapidly, which concludes the proof.

### 15.4 Examples

In this section we show by example how to apply the previous consistency results. In all cases, we assume that $\eta \in \mathcal{F}$ and we are concerned with the weak convergence
of $L\left(\eta_{n}\right)$ to $L^{*}$ for all such distributions of $(X, Y)$. The classes are as follows:

$$
\begin{aligned}
& \mathcal{F}_{1}=\left\{\eta=I_{[a, b]},-\infty \leq a \leq b \leq \infty\right\} \\
& \mathcal{F}_{1}^{\prime}=\left\{\eta=c I_{[a, b]}, c \in[0,1],-\infty \leq a \leq b \leq \infty\right\} \\
& \mathcal{F}_{2}=\left\{\eta=I_{\left[a_{1}, b_{1}\right] \times \cdots \times\left[a_{d}, b_{d}\right]},-\infty \leq a_{i} \leq b_{i} \leq \infty, 1 \leq i \leq d\right\} \\
& \mathcal{F}_{3}=\{\eta: \eta(x)=c x /(1+c x): x \geq 0, c \geq 0\} \\
& \mathcal{F}_{4}=\{\eta \text { is monotone decreasing on }[0,1]\} \\
& \mathcal{F}_{5}=\left\{\eta: \eta(x)=\frac{1}{1+\|x-m\|^{2}}, m \in \mathcal{R}^{d}\right\} \\
& \mathcal{F}_{6}=\left\{\eta: \eta(x)=\sin ^{2}(\theta x), \theta \in \mathcal{R}\right\} \\
& \mathcal{F}_{7}=\left\{\eta: \eta(x)=\frac{e^{\alpha_{0}+\alpha^{T} x}}{1+e^{\alpha_{0}+\alpha^{T} x}}, \alpha_{0} \in \mathcal{R}, x, \alpha \in \mathcal{R}^{d}\right\} .
\end{aligned}
$$

These are all rather simple, yet they will illustrate various points. Of these classes, $\mathcal{F}_{4}$ is nonparametric; yet, it behaves "better" than the one-parameter class $\mathcal{F}_{6}$, for example. We emphasize again that we are interested in consistency for all distributions of $X$.

For $\mathcal{F}_{1}, \eta_{n}$ will agree with the samples, that is, $\eta_{n}\left(X_{i}\right)=Y_{i}$ for all $i$, and therefore, $\eta_{n} \in \mathcal{F}_{1}$ is any function of the form $I_{[a, b]}$ with

$$
X^{(l-1)}<a \leq X^{(l)}, \quad X^{(r)} \leq b<X^{(r+1)}
$$

where $X^{(1)} \leq \cdots \leq X^{(n)}$ are the order statistics for $X_{1}, \ldots, X_{n}$ (with $X^{(0)}=$ $\left.-\infty, X^{(n+1)}=\infty\right), X^{(l)}$ is the smallest data point with $Y^{(l)}=1$, and $X^{(r)}$ is the largest data point with $Y^{(r)}=1$. As $L^{*}=0$, we claim that

$$
\begin{aligned}
\mathbf{E}\left\{L\left(\eta_{n}\right)\right\} & \leq \mathbf{P}\left\{X^{(l-1)}<X<X^{(l)}\right\}+\mathbf{P}\left\{X^{(r)}<X<X^{(r+1)}\right\} \\
& \leq \frac{4}{n+1}
\end{aligned}
$$

The rule is simply excellent, and has universal performance guarantees.
Remark. Note that in this case, maximum likelihood minimizes the empirical risk over the class of classifiers $\mathcal{C}=\left\{\phi=I_{[a, b]}, a \leq b\right\}$. As $\mathcal{C}$ has vc dimension $V_{\mathcal{C}}=2$ (Theorem 13.7), and $\inf _{\phi \in \mathcal{C}} L(\phi)=0$, Theorem 12.7 implies that for all $\epsilon>0$,

$$
\mathbf{P}\left\{L\left(\eta_{n}\right)>\epsilon\right\} \leq 8 n^{2} 2^{-n \epsilon / 2},
$$

and that

$$
\mathbf{E}\left\{L\left(\eta_{n}\right)\right\} \leq \frac{4 \log n+4}{n \log 2}
$$

(see Problem 12.8). With the analysis given here, we have gotten rid of the $\log n$ factor.

The class $\mathcal{F}_{1}^{\prime}$ is much more interesting. Here you will observe a dramatic difference with empirical risk minimization, as the parameter $c$ plays a key role that will aid a lot in the selection. Note that the likelihood product is zero if $N_{1}=0$ or if $Y_{i}=1$ for some $i$ with $X_{i} \notin[a, b]$, and it is

$$
c^{N_{1}}(1-c)^{N_{0}} \text { otherwise }
$$

where $N_{0}$ is the number of ( $X_{i}, Y_{i}$ ) pairs with $a \leq X_{i} \leq b, Y_{i}=0$, and $N_{1}$ is the number of ( $X_{i}, Y_{i}$ ) pairs with $a \leq X_{i} \leq b, Y_{i}=1$. For fixed $N_{0}, N_{1}$, this is maximal when

$$
c=\frac{N_{1}}{N_{0}+N_{1}}
$$

Resubstitution yields that the likelihood product is zero if $N_{1}=0$ or if $Y_{i}=1$ for some $i$ with $X_{i} \notin[a, b]$, and equals

$$
\exp \left\{N_{1} \log \frac{N_{1}}{N_{0}+N_{1}}+N_{0} \log \frac{N_{0}}{N_{0}+N_{1}}\right\} \text { if otherwise. }
$$

Thus, we should pick $\eta_{n}=c I_{[a, b]}$ such that $c=N_{1} /\left(N_{0}+N_{1}\right)$, and $[a, b]$ maximizes the divergence

$$
N_{1} \log \frac{N_{1}}{N_{0}+N_{1}}+N_{0} \log \frac{N_{0}}{N_{0}+N_{1}}
$$

See Problem 15.5.
For $\mathcal{F}_{2}$, by a similar argument as for $\mathcal{F}_{1}$, let $\left[A_{1}, B_{1}\right] \times \cdots \times\left[A_{d}, B_{d}\right]$ be the smallest rectangle of $\mathcal{R}^{d}$ that encloses all $X_{i}$ 's for which $Y_{i}=1$. Then, we know that $\eta_{n}=I_{\left[A_{1}, B_{1}\right] \times \cdots \times\left[A_{d}, B_{d}\right]}$ agrees with the data. Furthermore, $L^{*}=0$, and

$$
\mathbf{E}\left\{L\left(\eta_{n}\right)\right\} \leq \frac{4 d}{n+1}
$$

(see Problem 15.6).
The logarithm of the likelihood product for $\mathcal{F}_{3}$ is

$$
\sum_{i: Y_{i}=1} \log \left(c X_{i}\right)-\sum_{i=1}^{n} \log \left(1+c X_{i}\right)
$$

Setting the derivative with respect to $c$ equal to zero in the hope of obtaining an equation that must be satisfied by the maximum, we see that $c$ must satisfy

$$
N_{1}=\sum_{i=1}^{n} \frac{c X_{i}}{1+c X_{i}}=\sum_{i=1}^{n} \frac{X_{i}}{1 / c+X_{i}}
$$

unless $X_{1}=\cdots=X_{n}=0$, where $N_{1}=\sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}}$. This equation has a unique solution for $c$. The rule corresponding to $\eta_{n}$ is of the form $g(x)=I_{\{c x>1\}}=I_{\{x>1 / c\}}$. Equivalently,

$$
g(x)= \begin{cases}1 & \text { if } N_{1}>\sum_{i=1}^{n} \frac{X_{i}}{x+X_{i}} \\ 0 & \text { otherwise }\end{cases}
$$

This surprising example shows that we do not even have to know the parameter $c$ in order to describe the maximum likelihood rule. In quite a few cases, this shortcut makes such rules very appealing indeed. Furthermore, as the condition of Theorem 15.2 is fulfilled, $\eta \in \mathcal{F}_{3}$ implies that the rule is consistent as well.

For $\mathcal{F}_{4}$, it is convenient to order the $X_{i}$ 's from small to large and to identify $k$ consecutive groups, each group consisting of any number of $X_{i}$ 's with $Y_{i}=0$, and one $X_{i}$ with $Y_{i}=1$. Thus, $k=\sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}}$. Then, a moment's thought shows that $\eta_{n}$ must be piecewise constant, taking values $1 \geq a_{1} \geq \cdots \geq a_{k} \geq 0$ on the consecutive groups, and the value zero past the $k$-th group (which can only consist of $X_{i}$ 's with $Y_{i}=0$. The likelihood product thus is of the form

$$
a_{1}\left(1-a_{1}\right)^{n_{1}} a_{2}\left(1-a_{2}\right)^{n_{2}} \cdots a_{k}\left(1-a_{k}\right)^{n_{k}}
$$

where $n_{1}, \ldots, n_{k}$ are the cardinalities of the groups minus one (i.e., the number of $Y_{j}=0$-elements in the $i$-th group). Finally, we have

$$
\left(a_{1}, \ldots, a_{k}\right)=\underset{1 \geq a_{1} \geq \cdots \geq a_{k} \geq 0}{\arg \max } \prod_{j=1}^{k} a_{j}\left(1-a_{j}\right)^{n_{j}}
$$

To check consistency, we see that for every $X$ and $\epsilon>0$,

$$
N(X, \epsilon) \leq 4^{\left[\frac{3}{\epsilon}\right\rceil}
$$

(see Problem 15.7). Thus, the condition of Theorem 15.2 is satisfied, and $L\left(\eta_{n}\right) \rightarrow$ $L^{*}$ in probability, whenever $\eta \in \mathcal{F}_{4}$.

In $\mathcal{F}_{5}$, the maximum likelihood method will attempt to place $m$ at the center of the highest concentration in $\mathcal{R}^{d}$ of $X_{i}$ 's with $Y_{i}=1$, while staying away from $X_{i}$ 's with $Y_{i}=0$. Certainly, there are computational problems, but it takes little thought to verify that the conditions of Theorem 15.2 are satisfied. Explicit description of the rule is not necessary for some theoretical analysis!

Class $\mathcal{F}_{6}$ is a simple one-parameter class that does not satisfy the condition of Theorem 15.2. In fact, maximum likelihood fails here for the following reason: Let $X$ be uniform on $[0,1]$. Then the likelihood product is

$$
\prod_{i: Y_{i}=1} \sin ^{2}\left(\theta X_{i}\right) \times \prod_{i: Y_{i}=0} \cos ^{2}\left(\theta X_{i}\right)
$$

This product reaches a degenerate global maximum (1) as $\theta \rightarrow \infty$, regardless of the true (unknown) value of $\theta$ that gave rise to the data. See Problem 15.9.

Class $\mathcal{F}_{7}$ is used in the popular logistic discrimination problem, reviewed and studied by Anderson (1982), see also McLachlan (1992, Chapter 8). It is particularly important to observe that with this model,

$$
\{\eta>1 / 2\} \equiv\left\{e^{\alpha_{0}+\alpha^{T} x}>1 / 2\right\} \equiv\left\{\alpha^{T} x>\beta\right\}
$$

where $\beta=-\alpha_{0}-\log 2$. Thus, $\mathcal{F}_{7}$ subsumes linear discriminants. It does also force a bit more structure on the problem, making $\eta$ approach zero or one as we move away from the separating hyperplane. Day and Kerridge (1967) point out that model $\mathcal{F}_{7}$ is appropriate if the class-conditional densities take the form

$$
c f(x) \exp \left\{-\frac{1}{2}(x-m)^{T} \Sigma^{-1}(x-m)\right\},
$$

where $c$ is a normalizing constant, $f$ is a density, $m$ is a vector, and $\Sigma$ is a positive definite matrix; $f$ and $\Sigma$ must be the same for the two densities, but $c$ and $m$ may be different. Unfortunately, obtaining the best values for $\alpha_{0}$ and $\alpha$ by maximum likelihood takes a serious computational effort. Had we tried to estimate $f, m$, and $\Sigma$ in the last example, we would have done more than what is needed, as both $f$ and $\Sigma$ drop out of the picture. In this respect, the regression format is both parsimonious and lightweight.

### 15.5 Classical Maximum Likelihood: Distribution Format

In a more classical approach, we assume that given $Y=1, X$ has a density $f_{1}$, and given $Y=0, X$ has a density $f_{0}$, where both $f_{0}$ and $f_{1}$ belong to a given family $\mathcal{F}$ of densities. A similar setup may be used for atomic distributions, but this will not add anything new here, and is rather routine. The likelihood product for the data $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ is

$$
\prod_{i=1}^{n}\left(p f_{1}\left(X_{i}\right)\right)^{Y_{i}}\left((1-p) f_{0}\left(X_{i}\right)\right)^{1-Y_{i}}
$$

where $p=\mathbf{P}\{Y=1\}$ is assumed to be unknown as well. The maximum likelihood choices for $p, f_{0}, f_{1}$ are given by

$$
\left(p^{*}, f_{0}^{*}, f_{1}^{*}\right)=\underset{p \in[0,1], f_{0}, f_{1} \in \mathcal{F}}{\arg \max } \mathcal{L}_{n}\left(p, f_{0}, f_{1}\right),
$$

where

$$
\mathcal{L}_{n}\left(p, f_{0}, f_{1}\right)=\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i} \log \left(p f_{1}\left(X_{i}\right)\right)+\left(1-Y_{i}\right) \log \left((1-p) f_{0}\left(X_{i}\right)\right)\right) .
$$

Having determined $p^{*}, f_{0}^{*}, f_{1}^{*}$ (recall that the solution is not necessarily unique), the maximum likelihood rule is

$$
g_{n}(x)= \begin{cases}0 & \text { if }\left(1-p^{*}\right) f_{0}^{*}(x) \geq p^{*} f_{1}^{*}(x) \\ 1 & \text { otherwise. }\end{cases}
$$

Generally speaking, the distribution format is more sensitive than the regression format. It may work better under the correct circumstances. However, we give
up our universality, as the nature of the distribution of $X$ must be known. For example, if $X$ is distributed on a lower dimensional nonlinear manifold of $\mathcal{R}^{d}$, the distribution format is particularly inconvenient. Consistency results and examples are provided in a few exercises.

## Problems and Exercises

Problem 15.1. Show that if $|\mathcal{F}|=k<\infty$ and $\eta \in \mathcal{F}$, then $L\left(\eta_{n}\right) \rightarrow L^{*}$ with probability one, where $\eta_{n}$ is obtained by the maximum likelihood method. Also, prove that convergence with probability one holds in Theorem 15.2.

Problem 15.2. Prove that if $\eta, \eta^{\prime}$ are $[0,1]$-valued regression functions, then the divergence

$$
D\left(\eta^{\prime}\right)=\mathbf{E}\left\{\eta(X) \log \frac{\eta^{\prime}(X)}{\eta(X)}+(1-\eta(X)) \log \frac{1-\eta^{\prime}(X)}{1-\eta(X)}\right\}
$$

is nonpositive, and that $D=0$ if and only if $\eta(X)=\eta^{\prime}(X)$ with probability one. In this sense, $D$ measures the distance between $\eta$ and $\eta^{\prime}$.

Problem 15.3. Let

$$
L^{*}=\mathbf{E}\{\min (\eta(X), 1-\eta(X))\}, \quad L\left(\eta^{\prime}\right)=\mathbf{P}\left\{I_{\left\{\eta^{\prime}(X)>1 / 2\right\}} \neq Y\right\},
$$

and

$$
D\left(\eta^{\prime}\right)=\mathbf{E}\left\{\eta(X) \log \frac{\eta^{\prime}(X)}{\eta(X)}+(1-\eta(X)) \log \frac{1-\eta^{\prime}(X)}{1-\eta(X)}\right\},
$$

where $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$, and $\eta^{\prime}: \mathcal{R}^{d} \rightarrow[0,1]$ is arbitrary. Prove that

$$
D\left(\eta^{\prime}\right) \leq-\frac{\mathbf{E}\left\{\left(\eta(X)-\eta^{\prime}(X)\right)^{2}\right\}}{2} \leq-\frac{1}{2}\left(L\left(\eta^{\prime}\right)-L^{*}\right)^{2}
$$

(see also Problem 3.22). Hint: First prove that for $p, q \in[0,1]$,

$$
p \log \frac{q}{p}+(1-p) \log \frac{1-q}{1-p} \leq-\frac{(p-q)^{2}}{2}
$$

(use Taylor's series with remainder term for $\mathcal{H}(\cdot)$ ).
Problem 15.4. Show that for each $\eta$ there exists an $\epsilon_{0}>0$ such that for all $\epsilon \in\left(0, \epsilon_{0}\right)$,

$$
\mathbf{P}\left\{\mathcal{L}_{n}(\eta)>\mathbf{E}\left\{\mathcal{L}_{n}(\eta)\right\}+\epsilon\right\} \leq e^{-n \epsilon^{2} / 16}
$$

Hint: Proceed by Chernoff's bounding technique.
Problem 15.5. Consider $\eta \in \mathcal{F}_{1}^{\prime}$ and let $\eta_{n}$ be a maximum likelihood estimate over $\mathcal{F}_{1}^{\prime}$. Let $p=\mathbf{P}\{Y=1\}$. Show that

$$
L^{*}=p \min \left(1, \frac{1-c}{c}\right) .
$$

Derive an upper bound for

$$
\mathbf{E}\left\{L\left(\eta_{n}\right)-L^{*}\right\},
$$

where in case of multiple choices for $\eta_{n}$, you take the smallest $\eta_{n}$ in the equivalence class. This is an important problem, as $\mathcal{F}_{1}^{\prime}$ picks a histogram cell in a data-based manner. $\mathcal{F}_{1}^{\prime}$ may be generalized to the automatic selection of the best $k$-cell histogram: let $\mathcal{F}$ be the collection of all $\eta$ 's that are constant on the $k$ intervals determined by breakpoints $-\infty<a_{1}<\cdots<$ $a_{k-1}<\infty$.

Problem 15.6. Show that for the class $\mathcal{F}_{2}$,

$$
\mathbf{E}\left\{L\left(\eta_{n}\right)\right\} \leq \frac{4 d}{n+1}
$$

when $\eta \in \mathcal{F}_{2}$ and $\eta_{n}$ is the maximum likelihood estimate.
Problem 15.7. Show that for the class $\mathcal{F}_{4}$,

$$
N(X, \epsilon) \leq 4^{[3 / \epsilon\rceil}
$$

holds for all $X$ and $\epsilon>0$, that is, the bracketing $\epsilon$-entropy of $\mathcal{F}_{4}$ is bounded by $\lceil 3 / \epsilon\rceil$. Hint: Cover the class $\mathcal{F}_{4}$ by a class of monotone decreasing, piecewise constant functions, whose values are multiples of $\epsilon / 3$, and whose breakpoints are at the $k \epsilon / 3$-quantiles of the distribution of $X(k=1,2, \ldots,\lceil 3 / \epsilon\rceil)$.

Problem 15.8. Discuss the maximum likelihood method for the class

$$
\mathcal{F}=\left\{\eta: \eta(x)=\left\{\begin{array}{ll}
c & \text { if } \alpha^{T} x>\beta \\
c^{\prime} & \text { otherwise } ;
\end{array} \quad c, c^{\prime} \in[0,1], \alpha \in \mathcal{R}^{d}, \beta \in \mathcal{R}\right\}\right.
$$

What do the discrimination rules look like? If $\eta \in \mathcal{F}$, is the rule consistent? Can you guarantee a certain rate of convergence for $\mathbf{E}\left\{L\left(\eta_{n}\right)\right\}$ ? If $\eta \notin \mathcal{F}$, can you prove that $L\left(\eta_{n}\right)$ does not converge in probability to $\inf _{\eta^{\prime} \in \mathcal{F}} L\left(\eta^{\prime}\right)$ for some distribution of $(X, Y)$ with $\eta(x)=$ $\mathbf{P}\{Y=1 \mid X=x\}$ ? How would you obtain the values of $c, c^{\prime}, \alpha, \beta$ for the maximum likelihood choice $\eta_{n}$ ?

Problem 15.9. Let $X_{1}, \ldots, X_{n}$ be i.i.d. uniform [0,1] random variables. Let $y_{1}, \ldots, y_{n}$ be arbitrary $\{0,1\}$-valued numbers. Show that with probability one,

$$
\limsup _{\theta \rightarrow \infty} \prod_{i: y_{i}=1} \sin ^{2}\left(\theta X_{i}\right) \times \prod_{i: y_{i}=0} \cos ^{2}\left(\theta X_{i}\right)=1
$$

while for any $t<\infty$, with probability one,

$$
\sup _{0 \leq \theta \leq t_{i: y_{i}}=1} \prod \sin ^{2}\left(\theta X_{i}\right) \times \prod_{i: y_{i}=0} \cos ^{2}\left(\theta X_{i}\right)<1
$$

## 16

## Parametric Classification

What do you do if you believe (or someone tells you) that the conditional distributions of $X$ given $Y=0$ and $Y=1$ are members of a given family of distributions, described by finitely many real-valued parameters? Of course, it does not make sense to say that there are, say, six parameters. By interleaving the bits of binary expansions, we can always make one parameter out of six, and by splitting binary expressions, we may make a countable number of parameters out of one parameter (by writing the bits down in triangular fashion as shown below).

| $b_{1}$ | $b_{2}$ | $b_{4}$ | $b_{7}$ |
| :---: | :---: | :---: | :---: |
| $b_{3}$ | $b_{5}$ | $b_{8}$ | $\cdots$ |
| $b_{6}$ | $b_{9}$ | $\cdots$ |  |
| $b_{10}$ | $\cdots$ |  |  |

Thus, we must proceed with care. The number of parameters of a family really is measured more by the sheer size or vastness of the family than by mere representation of numbers. If the family is relatively small, we will call it parametric but we will not give you a formal definition of "parametric." For now, we let $\Theta$, the set of all possible values of the parameter $\theta$, be a subset of a finite-dimensional Euclidean space. Formally, let

$$
\mathcal{P}_{\Theta}=\left\{P_{\theta}: \theta \in \Theta\right\},
$$

be a class of probability distributions on the Borel sets of $\mathcal{R}^{d}$. Typically, the family $\mathcal{P}_{\Theta}$ is parametrized in a smooth way. That is, two distributions, corresponding to two parameter vectors close to each other are in some sense close to each other,
as well. Assume that the class-conditional densities $f_{0}$ and $f_{1}$ exist, and that both belong to the class of densities

$$
\mathcal{F}_{\Theta}=\left\{f_{\theta}: \theta \in \Theta\right\} .
$$

Discrete examples may be handled similarly. Take for example all gaussian distributions on $\mathcal{R}^{d}$, in which

$$
f_{\theta}(x)=\frac{1}{\sqrt{(2 \pi)^{d} \operatorname{det}(\Sigma)}} e^{-\frac{1}{2}(x-m)^{T} \Sigma^{-1}(x-m)}
$$

where $m \in \mathcal{R}^{d}$ is the vector of means, and $\Sigma$ is the covariance matrix. Recall that $x^{T}$ denotes the transposition of the column vector $x$, and $\operatorname{det}(\Sigma)$ is the determinant of $\Sigma$. This class is conveniently parametrized by $\theta=(m, \Sigma)$, that is, a vector of $d+d(d+1) / 2$ real numbers.

Knowing that the class-conditional distributions are in $\mathcal{P}_{\Theta}$ makes discrimination so much easier-rates of convergence to $L^{*}$ are excellent. Take $\mathcal{F}_{\Theta}$ as the class of uniform densities on hyperrectangles of $\mathcal{R}^{d}$ : this has $2 d$ natural parameters, the coordinates of the lower left and upper right vertices.


Figure 16.1. The class-conditional densities are uniform on hyperrectangles.
class 0

Given $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, a child could not do things wrong-for class 1 , estimate the upper right vertex by

$$
\left(\max _{1 \leq i \leq n: Y_{i}=1} X_{i}^{(1)}, \max _{1 \leq i \leq n: Y_{i}=1} X_{i}^{(2)}, \ldots, \max _{1 \leq i \leq n: Y_{i}=1} X_{i}^{(d)}\right)
$$

and similarly for the upper right vertex of the class 0 density. Lower left vertices are estimated by considering minima. If $A_{0}, A_{1}$ are the two unknown hyperrectangles and $p=\mathbf{P}\{Y=1\}$, the Bayes rule is simply

$$
g^{*}(x)= \begin{cases}1 & \text { if } x \in A_{1}-A_{0} \\ 0 & \text { if } x \in A_{0}-A_{1} \\ 1 & \text { if } x \in A_{1} \cap A_{0}, \frac{p}{\lambda\left(A_{1}\right)}>\frac{1-p}{\lambda\left(A_{0}\right)} \\ 0 & \text { if } x \in A_{1} \cap A_{0}, \frac{p}{\lambda\left(A_{1}\right)} \leq \frac{1-p}{\lambda\left(A_{0}\right)} .\end{cases}
$$

In reality, replace $A_{0}, A_{1}$, and $p$ by the sample estimates $\widehat{A_{0}}, \widehat{A_{1}}$ (described above) and $\widehat{p}=(1 / n) \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}}$. This way of doing things works very well, and we will pick up the example a bit further on. However, it is a bit ad hoc. There are indeed a few main principles that may be used in the design of classifiers under the additional information given here. In no particular order, here are a few methodologies:
(A) As the Bayes classifiers belong to the class

$$
\mathcal{C}=\left\{\phi=I_{\left\{p f_{\theta_{1}}>(1-p) f_{\theta_{0}}\right\}}: p \in[0,1], \theta_{0}, \theta_{1} \in \Theta\right\},
$$

it suffices to consider classifiers in $\mathcal{C}$. For example, if $\mathcal{F}_{\Theta}$ is the normal family, then $\mathcal{C}$ coincides with indicators of functions in the set

$$
\left\{\left\{x: x^{T} A x+b^{T} x+c>0: A \text { is a } d \times d \text { matrix, } b \in \mathcal{R}^{d}, c \in \mathcal{R}\right\}\right\}
$$

that is, the family of quadratic decisions. In the hyperrectangular example above, every $\phi$ is of the form $I_{A_{1}-A_{2}}$ where $A_{1}$ and $A_{2}$ are hyperrectangles of $\mathcal{R}^{d}$. Finding the best classifier of the form $\phi=I_{A}$ where $A \in \mathcal{A}$ is something we can do in a variety of ways: one such way, empirical risk minimization, is dealt with in Chapter 12 for example.
(B) Plug-in rules estimate $\left(\theta_{0}, \theta_{1}\right)$ by $\left(\widehat{\theta_{0}}, \widehat{\theta_{1}}\right)$ and $p=\mathbf{P}\{Y=1\}$ by $\widehat{p}$ from the data, and form the rule

$$
g_{n}(x)= \begin{cases}1 & \text { if } \widehat{p} f_{\widehat{\theta}_{1}}(x)>(1-\widehat{p}) f_{\widehat{\theta_{0}}}(x) \\ 0 & \text { otherwise } .\end{cases}
$$

The rule here is within the class $\mathcal{C}$ described in the previous paragraph. We are hopeful that the performance with $g_{n}$ is close to the performance with the Bayes rule $g^{*}$ when ( $\widehat{p}, \widehat{\theta_{0}}, \widehat{\theta_{1}}$ ) is close to $\left(p, \theta_{0}, \theta_{1}\right)$. For this strategy to work it is absolutely essential that $L\left(p, \theta_{0}, \theta_{1}\right)$ (the probability of error when $p, \theta_{0}, \theta_{1}$ are the parameters) be continuous in ( $p, \theta_{0}, \theta_{1}$ ). Robustness is a key ingredient. If the continuity can be captured in an inequality, then we may get performance guarantees for $\mathbf{E}\left\{L_{n}\right\}-L^{*}$ in terms of the distance between ( $\left.\widehat{p}, \widehat{\theta_{0}}, \widehat{\theta_{1}}\right)$ and $\left(p, \theta_{0}, \theta_{1}\right)$. Methods of estimating the parameters include maximum likelihood. This methodology is dealt with in Chapter 15. This method is rather sensitive to incorrect hypotheses (what if we were wrong about our assumption that the class-conditional distributions were in $\mathcal{P}_{\Theta}$ ?). Another strategy, minimum distance estimation, picks that member from $\mathcal{P}_{\Theta}$ that is closest in some sense to the raw empirical measure that puts mass $1 / n$ at each if the $n$ data points. See Section 16.3 This approach does not care about continuity of $L\left(p, \theta_{0}, \theta_{1}\right)$, as it judges members of $\mathcal{P}_{\Theta}$ by closeness in some space under a metric that is directly related to the probability of error. Robustness will drop out naturally.

A general approach should not have to know whether $\Theta$ can be described by a finite number of parameters. For example, it should equally well handle descriptions
as $\mathcal{P}_{\Theta}$ is the class of all distributions of $(X, Y)$ in which $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$ is monotonically increasing in all the components of $x$. Universal paradigms such as maximum likelihood, minimum distance estimation, and empirical error minimization are all applicable here. This particular $\mathcal{P}_{\Theta}$ is dealt with in Chapter 15, just to show you that the description of the class does invalidate the underlying principles.

### 16.1 Example: Exponential Families

A class $\mathcal{P}_{\Theta}$ is exponential, if every class-conditional density $f_{\theta}$ can be written in the form

$$
f_{\theta}(x)=c \alpha(\theta) \beta(x) e^{\sum_{i=1}^{k} \pi_{i}(\theta) \psi_{i}(x)}
$$

where $\beta, \psi_{1}, \ldots, \psi_{k}: \mathcal{R}^{d} \rightarrow \mathcal{R}, \beta \geq 0, \alpha, \pi_{1}, \ldots, \pi_{k}: \Theta \rightarrow \mathcal{R}$ are fixed functions, and $c$ is a normalizing constant. Examples of exponential families include the gaussian, gamma, beta, Rayleigh, and Maxwell densities (see Problem 16.4). The Bayes-rule can be rewritten as

$$
g^{*}(x)= \begin{cases}1 & \text { if } \log \left(\frac{p f_{\theta_{1}^{*}}(x)}{(1-p) f_{\theta_{0}^{*}}(x)}\right)>0 \\ 0 & \text { otherwise }\end{cases}
$$

This is equivalent to

$$
g^{*}(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{k} \psi_{i}(x)\left(\pi_{i}\left(\theta_{0}^{*}\right)-\pi_{i}\left(\theta_{1}^{*}\right)\right)<\log \left(\frac{p \alpha\left(\theta_{1}^{*}\right)}{\left(1-\overline{p) \alpha\left(\theta_{0}^{*}\right)}\right)}\right. \\ 0 & \text { otherwise } .\end{cases}
$$

The Bayes-rule is a so-called generalized linear rule with $1, \psi_{1}, \ldots, \psi_{k}$ as basis functions. Such rules are easily dealt with by empirical risk minimization and related methods such as complexity regularization (Chapters $12,17,18,22$ ).

Another important point is that $g^{*}$ does not involve the function $\beta$. For all we know, $\beta$ may be some esoteric ill-behaved function that would make estimating $f_{\theta}(x)$ all but impossible if $\beta$ were unknown. Even if $\mathcal{P}_{\Theta}$ is the huge family in which $\beta \geq 0$ is left undetermined, but it is known to be identical for the two classconditional densities (and $\alpha(\theta)$ is just a normalization factor), we would still only have to look at the same small class of generalized linear discrimination rules! So, densities do not matter-ratios of densities do. Pattern recognition should be, and is, easier than density estimation. Those who first estimate $\left(f_{0}, f_{1}\right)$ by $\left(\widehat{f_{0}}, \widehat{f_{1}}\right)$ and then construct rules based on the sign of $\widehat{p} \widehat{f_{1}}(x)-(1-\widehat{p}) \widehat{f_{0}}(x)$ do themselves a disservice.

### 16.2 Standard Plug-In Rules

In standard plug-in methods, we construct estimates $\widehat{\theta_{0}}, \widehat{\theta_{1}}$ and $\widehat{p}$ from the data and use them to form a classifier

$$
g_{n}(x)= \begin{cases}1 & \text { if } \widehat{p} f_{\widehat{\theta_{1}}}(x)>(1-\widehat{p}) f_{\widehat{\theta_{0}}}(x) \\ 0 & \text { otherwise }\end{cases}
$$

It is generally not true that if $\widehat{p} \rightarrow p, \widehat{\theta_{0}} \rightarrow \theta_{0}$, and $\widehat{\theta_{1}} \rightarrow \theta_{1}$ in probability, then $L\left(g_{n}\right) \rightarrow L^{*}$ in probability, where $L\left(g_{n}\right)$ is the probability of error with $g_{n}$. Consider the following simple example:
Example. Let $\mathcal{F}_{\Theta}$ be the class of all uniform densities on $[-\theta, 0]$ if $\theta \neq 1$ and on $[0, \theta]$ if $\theta=1$. Let $\theta_{0}=1, \theta_{1}=2, p=1 / 2$. Then a reasonable estimate of $\theta_{i}$ would be $\widehat{\theta_{i}}=\max _{i: Y_{i}=1}\left|X_{i}\right|$. Clearly, $\widehat{\theta_{i}} \rightarrow \theta_{i}$ in probability. However, as $\widehat{\theta_{i}} \neq 1$ with probability one, we note that $g_{n}(x)=0$ for $x>0$ and thus, even though $L^{*}=0$ (as the supports of $f_{\theta_{0}}$ and $f_{\theta_{1}}$ are not overlapping), $L\left(g_{n}\right) \geq \mathbf{P}\{Y=1\}=p=1 / 2$. The problem with this is that there is no continuity with respect to $\theta$ in $\mathcal{F}_{\Theta}$.

Basic consistency based upon continuity considerations is indeed easy to establish. As ratios of densities matter, it helps to introduce

$$
\eta_{\theta}(x)=\frac{p f_{\theta_{1}}(x)}{(1-p) f_{\theta_{0}}(x)+p f_{\theta_{1}}(x)}=\mathbf{P}\{Y=1 \mid X=x\},
$$

where $\theta=\left(p, \theta_{0}, \theta_{1}\right)$ or $\left(\theta_{0}, \theta_{1}\right)$ as the case may be. We recall that if $g_{n}(x)=$ $I_{\left\{\eta_{\hat{\theta}}(x)>1 / 2\right\}}$ where $\widehat{\theta}$ is an estimate of $\theta$, then

$$
L\left(g_{n}\right)-L^{*} \leq 2 \mathbf{E}\left\{\left|\eta_{\widehat{\theta}}(X)-\eta_{\theta}(X)\right| \mid D_{n}\right\}
$$

where $D_{n}$ is the data sequence. Thus,

$$
\mathbf{E}\left\{L\left(g_{n}\right)-L^{*}\right\} \leq 2 \mathbf{E}\left\{\left|\eta_{\widehat{\theta}}(X)-\eta_{\theta}(X)\right|\right\}
$$

By the Lebesgue dominated convergence theorem, we have, without further ado:
Theorem 16.1. If $\eta_{\theta}$ is continuous $\operatorname{in} \theta$ in the $L_{1}(\mu)$ sense, where $\mu$ is the measure of $X$, and $\widehat{\theta} \rightarrow \theta$ in probability, then $\mathbf{E}\left\{L\left(g_{n}\right)\right\} \rightarrow L^{*}$ for the standard plug-in rule.

In some cases, we can do better and derive rates of convergence by examining the local behavior of $\eta_{\theta}(x)$. For example, if $\eta_{\theta}(x)=e^{-\theta|x|}, x \in \mathcal{R}$, then

$$
\left|\eta_{\theta}(x)-\eta_{\theta^{\prime}}(x)\right| \leq|x|\left|\theta-\theta^{\prime}\right|
$$

and

$$
\begin{aligned}
\mathbf{E}\left\{L\left(g_{n}\right)\right\}-L^{*} & \leq 2 \mathbf{E}\{|X||\theta-\widehat{\theta}|\} \\
& \leq 2 \sqrt{\mathbf{E}\left\{X^{2}\right\}} \sqrt{\mathbf{E}\left\{(\theta-\hat{\theta})^{2}\right\}}
\end{aligned}
$$

yielding an explicit bound. In general, $\theta$ is multivariate, consisting at least of the triple ( $p, \theta_{0}, \theta_{1}$ ), but the above example shows the way to happy analysis. For the simple example given here, $\mathbf{E}\left\{L\left(g_{n}\right)\right\} \rightarrow L^{*}$ if $\mathbf{E}\left\{(\theta-\widehat{\theta})^{2}\right\} \rightarrow 0$. In fact, this seems to suggest that for this family, $\widehat{\theta}$ should be found to minimize $\mathbf{E}\left\{(\theta-\widehat{\theta})^{2}\right\}$. This is false. One is always best off minimizing the probability of error. Other criteria may be relevant via continuity, but should be considered with care.

How certain parameters are estimated for given families of distributions is what mathematical statistics is all about. The maximum likelihood principle looms large: $\theta_{0}$ is estimated for densities $f_{\theta}$ by

$$
\widehat{\theta_{0}}=\underset{\theta}{\arg \max } \prod_{i: Y_{i}=0} f_{\theta}\left(X_{i}\right)
$$

for example. If you work out this (likelihood) product, you will often discover a simple form for the estimate of the data. In discrimination, only $\eta$ matters, not the class-conditional densities. Maximum likelihood in function of the $\eta$ 's was studied in Chapter 15. We saw there that this is often consistent, but that maximum likelihood behaves poorly when the true distribution is not in $\mathcal{P}_{\theta}$. We will work out two simple examples:

As an example of the maximum likelihood method in discrimination, we assume that

$$
\mathcal{F}_{\Theta}=\left\{f_{\theta}(x)=\frac{x^{\alpha-1} e^{-x / \beta}}{\beta^{\alpha} \Gamma(\alpha)} I_{\{x>0\}}: \alpha, \beta>0\right\}
$$

is the class of all gamma densities with $\theta=(\alpha, \beta)$. The likelihood product given $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, is, if $\theta_{0}, \theta_{1}$ are the unknown parameters,

$$
\prod_{i=1}^{n}\left(p f_{\theta_{1}}\left(X_{i}\right)\right)^{Y_{i}}\left((1-p) f_{\theta_{0}}\left(X_{i}\right)\right)^{1-Y_{i}}
$$

This is the probability of observing the data sequence if the $f_{\theta}$ 's were in fact discrete probabilities. This product is simply

$$
\begin{aligned}
& \left(\frac{p}{\beta_{1}^{\alpha_{1}} \Gamma\left(\alpha_{1}\right)}\right)^{\sum_{i=1}^{n} Y_{i}}\left(\frac{1-p}{\beta_{0}^{\alpha_{0}} \Gamma\left(\alpha_{0}\right)}\right)^{n-\sum_{i=1}^{n} Y_{i}} \\
& \quad \cdot e^{-\frac{\sum Y_{i} X_{i}}{\beta_{1}}-\frac{\sum\left(1-Y_{i}\right) X_{i}}{\beta_{0}}+\left(\alpha_{1}-1\right) \sum Y_{i} \log X_{i}+\left(\alpha_{0}-1\right) \sum\left(1-Y_{i}\right) \log X_{i}},
\end{aligned}
$$

where $\theta_{0}=\left(\alpha_{0}, \beta_{0}\right), \theta_{1}=\left(\alpha_{1}, \beta_{1}\right)$. The first thing we notice is that this expression depends only on certain functions of the data, notably $\sum Y_{i} X_{i}, \sum\left(1-Y_{i}\right) X_{i}$, $\sum Y_{i} \log X_{i}, \sum\left(1-Y_{i}\right) \log X_{i}$, and $\sum Y_{i}$. These are called the sufficient statistics for the problem at hand. We may in fact throw away the data and just store the sufficient statistics. The likelihood product has to be maximized. Even in this rather simple univariate example, this is a nontrivial task. Luckily, we immediately note that $p$ occurs in the factor $p^{N}(1-p)^{n-N}$, where $N=\sum Y_{i}$. This is maximal at $\widehat{p}=N / n$, a well-known result. For fixed $\alpha_{0}, \alpha_{1}$, we can also get $\beta_{0}, \beta_{1}$; but for
variable $\alpha_{0}, \alpha_{1}, \beta_{0}, \beta_{1}$, the optimization is difficult. In $d$-dimensional cases, one has nearly always to resort to specialized algorithms.

As a last example, we return to $\mathcal{F}_{\Theta}=\left\{\right.$ uniform densities on rectangles of $\left.\mathcal{R}^{d}\right\}$. Here the likelihood product once again has $p^{N}(1-p)^{n-N}$ as a factor, leading to $\widehat{p}=$ $N / n$. The other factor (if $\left(a_{1}, b_{1}\right),\left(a_{0}, b_{0}\right)$ are the lower left, upper right vertices of the rectangles for $\left.\theta_{1}, \theta_{0}\right)$ is zero if for some $i, Y_{i}=1, X_{i} \notin \operatorname{rectangle}\left(a_{1}, b_{1}\right)$, or $Y_{i}=0, X_{i} \notin \operatorname{rectangle}\left(a_{0}, b_{0}\right)$. Otherwise, it is

$$
\frac{1}{\left\|b_{1}-a_{1}\right\|^{N}} \cdot \frac{1}{\left\|b_{0}-a_{0}\right\|^{n-N}}
$$

where $\left\|b_{1}-a_{1}\right\|=\prod_{j=1}^{d}\left(b_{1}^{(j)}-a_{1}^{(j)}\right)$ denotes the volume of the rectangle $\left(a_{1}, b_{1}\right)$, and similarly for $\left\|b_{0}-a_{0}\right\|$. This is maximal if $\left\|b_{1}-a_{1}\right\|$ and $\left\|b_{0}-a_{0}\right\|$ are minimal. Thus, the maximum likelihood estimates are

$$
\begin{aligned}
& \widehat{a}_{i}^{(k)}=\min _{j: Y_{j}=i} X_{i}^{(k)}, \quad 1 \leq k \leq d, i=0,1 \\
& \widehat{b}_{i}^{(k)}=\max _{j: Y_{j}=i} X_{i}^{(k)}, \quad 1 \leq k \leq d, i=0,1
\end{aligned}
$$

where $\widehat{a}_{i}=\left(\widehat{a}_{i}^{(1)}, \ldots, \widehat{a}_{i}^{(d)}\right), \widehat{b}_{i}=\left(\widehat{b}_{i}^{(1)}, \ldots, \widehat{b}_{i}^{(d)}\right)$, and $X_{i}=\left(X_{i}^{(1)}, \ldots, X_{i}^{(d)}\right)$.
Rates of convergence may be obtained via some of the (in)equalities of Chapter 6 , such as

$$
\begin{align*}
\mathbf{E}\left\{L\left(g_{n}\right)\right\}-L^{*} \leq & 2 \mathbf{E}\left\{\left|\eta_{n}(X)-\eta(X)\right|\right\}  \tag{1}\\
& \left(\text { where } \eta_{n}=\widehat{p} f_{\widehat{\theta_{1}}} /\left(\widehat{p} f_{\widehat{\theta_{1}}}+(1-\widehat{p}) f_{\widehat{\theta_{0}}}\right)\right) \\
\mathbf{E}\left\{L\left(g_{n}\right)\right\}-L^{*} \leq & 2 \sqrt{\mathbf{E}\left\{\left(\eta_{n}(X)-\eta(X)\right)^{2}\right\}}  \tag{2}\\
\mathbf{E}\left\{L\left(g_{n}\right)\right\}-L^{*}= & 2 \mathbf{E}\left\{\left|\eta(X)-\frac{1}{2}\right| I_{\left\{g_{n}(X) \neq g^{*}(X)\right\}}\right\} \tag{3}
\end{align*}
$$

The rate with which $\widehat{\theta}$ approaches $\theta$ in $\Theta$ (measured with some metric) may be very different from that with which $L\left(g_{n}\right)$ approaches $L^{*}$. As shown in Theorem 6.5 , the inequality (2) is always loose, yet it is this inequality that is often used to derive rates of convergence by authors and researchers. Let us take a simple example to illustrate this point. Assume $\mathcal{F}_{\Theta}$ is the family of normal densities on the real line. If $\theta_{0}=\left(m_{0}, \sigma_{0}\right), \theta_{1}=\left(m_{1}, \sigma_{1}\right)$ are the unknown means and standard deviations of the class-conditional densities, and $p=\mathbf{P}\{Y=1\}$ is also unknown, then we may estimate $p$ by $\widehat{p}=\sum_{j=1}^{n} Y_{j} / n$, and $m_{i}$ and $\sigma_{i}$ by

$$
\widehat{m}_{i}=\frac{\sum_{j=1}^{n} X_{j} I_{\left\{Y_{j}=i\right\}}}{\sum_{j=1}^{n} I_{\left\{Y_{j}=i\right\}}} \quad \text { and } \quad \widehat{\sigma}_{i}^{2}=\frac{\sum_{j=1}^{n}\left(X_{j}-\widehat{m}_{i}\right)^{2} I_{\left\{Y_{j}=i\right\}}}{\sum_{j=1}^{n} I_{\left\{Y_{j}=i\right\}}}, i=0,1,
$$

respectively, when denominators are positive. If a denominator is 0 , set $\widehat{m}_{i}=0$, $\widehat{\sigma}_{i}^{2}=1$. From Chebyshev's inequality, we can verify that for fixed $p \in(0,1)$,
$\mathbf{E}\{|p-\widehat{p}|\}=O(1 / \sqrt{n}), \mathbf{E}\left\{\left|m_{i}-\widehat{m}_{i}\right|\right\}=O(1 / \sqrt{n}), \mathbf{E}\left\{\left|\sigma_{i}-\widehat{\sigma}_{i}\right|\right\}=O(1 / \sqrt{n})$ (Problem 16.3).
If we compute $\mathbf{E}\left\{\left|\eta_{n}(X)-\eta(X)\right|\right\}$, we will discover that $\mathbf{E}\left\{L\left(g_{n}\right)\right\}-L^{*}=$ $O(1 / \sqrt{n})$. However, if we compute $\mathbf{E}\left\{\left|\eta(X)-\frac{1}{2}\right| I_{\left\{g_{n}(X) \neq \xi^{*}(X)\right\}}\right\}$, we will find that $\mathbf{E}\left\{L\left(g_{n}\right)\right\}-L^{*}=O(1 / n)$. Thus, while the parameters converge at a rate $O(1 / \sqrt{n})$ dictated by the central limit theorem, and while $\eta_{n}$ converges to $\eta$ in $L_{1}(\mu)$ with the same rate, the error rate in discrimination is much smaller. See Problems 16.7 to 16.9 for some practice in this respect.

Bibliographic remarks. McLachlan (1992) has a comprehensive treatment on parametric classification. Duda and Hart (1973) have many good introductory examples and a nice discussion on sufficient statistics, a topic we do not deal with in this text. For maximum likelihood estimation, see Hjort (1986a; 1986b).

### 16.3 Minimum Distance Estimates

Here we describe a general parameter estimation principle that appears to be more suitable for plug-in classification rules than the maximum likelihood method. The estimated parameter is obtained by the projection of the empirical measure on the parametric family.

The principle of minimum distance estimation may be described as follows. Let $\mathcal{P}_{\Theta}=\left\{P_{\theta}: \theta \in \Theta\right\}$ be a parametric family of distributions, and assume that $P_{\theta^{*}}$ is the unknown distribution of the i.i.d. observations $Z_{1}, \ldots, Z_{n}$. Denote by $v_{n}$ the empirical measure

$$
v_{n}(A)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{Z_{i} \in A\right\}}, \quad A \subset \mathcal{R}^{d} .
$$

Let $D(\cdot, \cdot)$ be a metric on the set of all probability distributions on $\mathcal{R}^{d}$. The minimum distance estimate of $\theta^{*}$ is defined as

$$
\theta_{n}=\underset{\theta \in \Theta}{\arg \min } D\left(v_{n}, P_{\theta}\right),
$$

if it exists and is unique. If it is not unique, select one candidate for which the minimum is attained.

Consider for example the Kolmogorov-Smirnov distance

$$
D_{\mathrm{KS}}(P, Q)=\sup _{z \in \mathcal{R}^{d}}|F(z)-G(z)|,
$$



Figure 16.2. The member of $\mathcal{P}_{\Theta}$ closest to the empirical measure $\nu_{n}$ is chosen by the minimum distance estimate.
where $F$ and $G$ are the distribution functions of the measures $P$ and $Q$ on $\mathcal{R}^{d}$, respectively. It is easy to see that $D_{\mathrm{KS}}$ is a metric. Note that

$$
\sup _{z \in \mathcal{R}^{d}}|F(z)-G(z)|=\sup _{A \in \mathcal{A}}|P(A)-Q(A)|,
$$

where $\mathcal{A}$ is the class of sets of the form $\left(-\infty, z^{(1)}\right) \times \cdots \times\left(-\infty, z^{(d)}\right)$, for $z=\left(z^{(1)}, \ldots, z^{(d)}\right) \in \mathcal{R}^{d}$. For the Kolmogorov-Smirnov distance between the estimated and the true distributions, by the triangle inequality, we have

$$
\begin{aligned}
D_{\mathrm{KS}}\left(P_{\theta_{n}}, P_{\theta^{*}}\right) & \leq D_{\mathrm{KS}}\left(P_{\theta_{n}}, v_{n}\right)+D_{\mathrm{KS}}\left(v_{n}, P_{\theta^{*}}\right) \\
& \leq 2 D_{\mathrm{KS}}\left(v_{n}, P_{\theta^{*}}\right),
\end{aligned}
$$

where in the second inequality we used the definition of $\theta_{n}$. Now notice that the upper bound is just twice the Kolmogorov-Smirnov distance between an empirical distribution and the true distribution. By a straightforward application of Theorem 12.5 , for every $n$ and $\epsilon>0$,

$$
\mathbf{P}\left\{D_{\mathrm{KS}}\left(v_{n}, P_{\theta^{*}}\right)>\epsilon\right\}=\mathbf{P}\left\{\sup _{A \in \mathcal{A}}\left|v_{n}(A)-P_{\theta^{*}}(A)\right|>\epsilon\right\} \leq 8\left(\frac{n e}{d}\right)^{d} e^{-n \epsilon^{2} / 32}
$$

since the $n$-th shatter coefficient $s(\mathcal{A}, n)$ of the class of sets

$$
\mathcal{A}=\left\{\left(-\infty, z^{(1)}\right) \times \cdots \times\left(-\infty, z^{(d)}\right), z=\left(z^{(1)}, \ldots, z^{(d)}\right) \in \mathcal{R}^{d}\right\}
$$

cannot exceed $(n e / d)^{d}$. This can easily be seen by an argument similar to that in the proof of Theorem 13.8. From the inequalities above, we see that the KolmogorovSmirnov distance between the estimated and the true distributions is always $O(\sqrt{\log n / n})$. The only condition we require is that $\theta_{n}$ be well defined.

Of course, rather than the Kolmogorov-Smirnov distance, it is the error probability of the plug-in classification rule in which we are primarily interested. In order to make the connection, we adapt the notions of minimum distance estimation and Kolmogorov-Smirnov distance to better suit the classification problem we are after. Every parametric family of distributions defines a class of sets in $\mathcal{R}^{d}$ as the collection $\mathcal{A}$ of sets of the form $\{x: \phi(x)=1\}$, where the classifiers $\phi$ are the possible plug-in rules defined by the parametric family. The idea here is to perform minimum distance estimation with the generalized Kolmogorov-Smirnov distance with respect to a class of sets closely related to $\mathcal{A}$.

Assume that both class-conditional distributions $P_{\theta_{0}}, P_{\theta_{1}}$ belong to a parametric family $\mathcal{P}_{\Theta}$, and the class-conditional densities $f_{\theta_{0}}, f_{\theta_{1}}$ exist. Then the Bayes rule may be written as

$$
g^{*}(x)= \begin{cases}1 & \text { if } \alpha_{\theta}(x)>0 \\ 0 & \text { otherwise }\end{cases}
$$

where $\alpha_{\theta}(x)=p f_{\theta_{1}}(x)-(1-p) f_{\theta_{0}}(x)$ and $p=\mathbf{P}\{Y=1\}$. We use the short notation $\theta=\left(p, \theta_{0}, \theta_{1}\right)$. The function $\alpha_{\theta}$ may be thought of as the Radon-Nikodym derivative of the signed measure $Q_{\theta}=p P_{\theta_{1}}-(1-p) P_{\theta_{0}}$. In other words, to each Borel set $A \subset \mathcal{R}^{d}, Q_{\theta}$ assigns the real number $Q_{\theta}(A)=p P_{\theta_{1}}(A)-(1-p) P_{\theta_{0}}(A)$. Given the data $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$, we define the empirical counterpart of $Q_{\theta}$ by

$$
v_{n}(A)=\frac{1}{n} \sum_{i=1}^{n}\left(2 Y_{i}-1\right) I_{\left\{X_{i} \in A\right\}} .
$$

The minimum distance classification rule we propose projects the empirical signed measure $v_{n}$ on the set of measures $Q_{\theta}$. The metric we use is also specifically fitted to the given pattern recognition problem: define the class of sets

$$
\mathcal{A}=\left\{\left\{x \in \mathcal{R}^{d}: \alpha_{\theta}(x)>0\right\}: p \in[0,1], \theta_{0}, \theta_{1} \in \Theta\right\}
$$

$\mathcal{A}$ is just the class of sets $A \subset \mathcal{R}^{d}$ such that $I_{\{x \in A\}}$ is the Bayes rule for some $\theta=\left(p, \theta_{0}, \theta_{1}\right)$. Also introduce

$$
\mathcal{B}=\left\{A \cap B^{c}: A, B \in \mathcal{A}\right\}
$$

Given two signed measures $Q, Q^{\prime}$, we define their generalized KolmogorovSmirnov distance by

$$
D_{\mathcal{B}}\left(Q, Q^{\prime}\right)=\sup _{A \in \mathcal{B}}\left|Q(A)-Q^{\prime}(A)\right|
$$

that is, instead of the class of half-infinite intervals as in the definition of the ordinary Kolmogorov-Smirnov distance, here we take the supremum over $\mathcal{B}$, a class tailored to our discrimination problem. Now, we are ready to define our minimum distance estimate

$$
\widehat{\theta}=\underset{\theta}{\arg \min } D_{\mathcal{B}}\left(Q_{\theta}, v_{n}\right)
$$

where the minimum is taken over all triples $\theta=\left(p, \theta_{0}, \theta_{1}\right)$ with $p \in[0,1]$, $\theta_{0}, \theta_{1} \in \Theta$. The corresponding classification rule is

$$
g_{\widehat{\theta}}(x)= \begin{cases}1 & \text { if } \alpha_{\widehat{\theta}}(x)>0 \\ 0 & \text { otherwise }\end{cases}
$$

The next theorem shows that if the parametric assumption is valid, then $g_{\widehat{\theta}}$ performs extremely well. The theorem shows that if $\mathcal{A}$ has finite vC dimension $V_{\mathcal{A}}$, then without any additional conditions imposed on the parametric class, the corresponding error probability $L\left(g_{\widehat{\theta}}\right)$ is not large: $L\left(g_{\widehat{\theta}}\right)-L^{*}=O\left(\sqrt{V_{\mathcal{A}} \log n / n}\right)$.

Theorem 16.2. Assume that both conditional densities $f_{\theta_{0}}, f_{\theta_{1}}$ are in the parametric class $\mathcal{F}_{\Theta}$. Then for the classification rule defined above, we have for every $n$ and $\epsilon>0$ that

$$
\mathbf{P}\left\{L\left(g_{\widehat{\theta}}\right)-L^{*}>\epsilon\right\} \leq 8 s(\mathcal{B}, n) e^{-n \epsilon^{2} / 512}
$$

where $s(\mathcal{B}, n)$ is the $n$-th shatter coefficient of $\mathcal{B}$. Furthermore,

$$
\mathbf{E}\left\{L\left(g_{\widehat{\theta}}\right)-L^{*}\right\} \leq 32 \sqrt{\frac{\log (8 e s(\mathcal{B}, n))}{2 n}}
$$

Remark. Recalling from Chapter 13 that $s(\mathcal{B}, n) \leq s^{2}(\mathcal{A}, n)$ and that $s(\mathcal{A}, n) \leq$ $(n+1)^{V_{\mathcal{A}}}$, where $V_{\mathcal{A}}$ is the vC dimension of $\mathcal{A}$, we obtain the bound

$$
\mathbf{P}\left\{L\left(g_{\widehat{\theta}}\right)-L^{*}>\epsilon\right\} \leq 8(n+1)^{2 V_{\mathcal{A}}} e^{-n \epsilon^{2} / 512}
$$

The proof of the theorem is based upon two key observations. The first lemma provides a bound on $L\left(g_{\widehat{\theta}}\right)-L^{*}$ in terms of the generalized Kolmogorov-Smirnov distance between the estimated and the true parameters. The second lemma is a straightforward extension of the Vapnik-Chervonenkis inequality to signed measures.

## Lemma 16.1.

$$
L\left(g_{\widehat{\theta}}\right)-L^{*} \leq 2 D_{\mathcal{B}}\left(Q_{\theta}, Q_{\widehat{\theta}}\right)
$$

Proof. Denote

$$
A_{\theta}=\left\{x: \alpha_{\theta}(x)>0\right\} \quad \text { and } \quad A_{\widehat{\theta}}=\left\{x: \alpha_{\widehat{\theta}}(x)>0\right\}
$$

that is, $g_{\widehat{\theta}}(x)=I_{\left\{\alpha_{\widehat{\theta}}(x)>0\right\}}, g^{*}(x)=I_{\left\{\alpha_{\theta}(x)>0\right\}}$, and $A_{\theta}, A_{\widehat{\theta}} \in \mathcal{A}$. At the first crucial step, we use the equality of Theorem 2.2:

$$
\begin{aligned}
L\left(g_{\widehat{\theta}}\right)-L^{*} & =\int I_{\left\{g_{\overparen{\theta}}(x) \neq \neq g_{\theta}(x)\right\}}\left|\alpha_{\theta}(x)\right| d x \\
& =-\int_{A_{\overparen{\theta} \cap A_{\theta}^{c}}} \alpha_{\theta}(x) d x+\int_{A_{\overparen{\theta}}^{c} \cap A_{\theta}} \alpha_{\theta}(x) d x
\end{aligned}
$$

$$
\begin{aligned}
\leq & \int_{A_{\widehat{\theta}} \cap A_{\theta}^{c}} \alpha_{\widehat{\theta}}(x) d x-\int_{A_{\widehat{\theta}} \cap A_{\theta}^{c}} \alpha_{\theta}(x) d x \\
& +\int_{A_{\overparen{\theta}}^{c} \cap A_{\theta}} \alpha_{\theta}(x) d x-\int_{A_{\widehat{\theta}}^{c} \cap A_{\theta}} \alpha_{\widehat{\theta}}(x) d x \\
= & Q_{\widehat{\theta}}\left(A_{\widehat{\theta}} \cap A_{\theta}^{c}\right)-Q_{\theta}\left(A_{\widehat{\theta}} \cap A_{\theta}^{c}\right)+Q_{\theta}\left(A_{\widehat{\theta}}^{c} \cap A_{\theta}\right)-Q_{\widehat{\theta}}\left(A_{\widehat{\theta}}^{c} \cap A_{\theta}\right) \\
\leq & 2 D_{\mathcal{B}}\left(Q_{\theta}, Q_{\widehat{\theta}}\right)
\end{aligned}
$$

since both $A_{\widehat{\theta}} \cap A_{\theta}^{c}$ and $A_{\widehat{\theta}}^{c} \cap A_{\theta}$ are members of $\mathcal{B}_{\Theta}$.
The next lemma is an extension of the Vapnik-Chervonenkis inequality. The proof is left to the reader (Problem 16.11).

Lemma 16.2. For every $n$ and $\epsilon>0$,

$$
\mathbf{P}\left\{D_{\mathcal{B}}\left(Q_{\theta}, v_{n}\right)>\epsilon\right\} \leq 8 s(\mathcal{B}, n) e^{-n \epsilon^{2} / 32}
$$

The rest of the proof of Theorem 16.2 shows that the generalized KolmogorovSmirnov distance with respect to $\mathcal{B}$ between the estimated and true distributions is small with large probability. This can be done as we proved for the KolmogorovSmirnov distance at the beginning of the section.
Proof of Theorem 16.2. By Lemma 16.1,

$$
\begin{aligned}
L\left(g_{\widehat{\theta}}\right)-L^{*} & \leq 2 D_{\mathcal{B}}\left(Q_{\theta}, Q_{\widehat{\theta}}\right) \\
& \leq 2 D_{\mathcal{B}}\left(Q_{\theta}, v_{n}\right)+2 D_{\mathcal{B}}\left(Q_{\widehat{\theta}}, v_{n}\right) \quad \text { (by the triangle inequality) } \\
& \left.\leq 4 D_{\mathcal{B}}\left(Q_{\theta}, v_{n}\right) \quad \text { (by the definition of } \widehat{\theta}\right)
\end{aligned}
$$

The theorem now follows by Lemma 16.2.

Finally, we examine robustness of the minimum distance rule against modeling errors, that is, what happens if the distributions are not in $\mathcal{P}_{\Theta}$. A good rule should still work reasonably well if the distributions are in some sense close to the modeled parametric class $\mathcal{P}_{\Theta}$. Observe that if for some $\theta=\left(p, \theta_{0}, \theta_{1}\right)$ the Bayes rule can be written as

$$
g^{*}(x)= \begin{cases}1 & \text { if } \alpha_{\theta}(x)>0 \\ 0 & \text { otherwise }\end{cases}
$$

then Lemma 16.1 remains valid even when the class-conditional distributions are not in $\mathcal{P}_{\Theta}$. Denote the true class-conditional distributions by $P_{0}^{*}, P_{1}^{*}$, let $p^{*}=$ $\mathbf{P}\{Y=1\}$, and introduce $Q^{*}=p^{*} P_{1}^{*}-(1-p) P_{0}^{*}$. Thus,

$$
\begin{aligned}
& L\left(g_{\widehat{\theta}}\right)-L^{*} \\
& \quad \leq 2 D_{\mathcal{B}}\left(Q^{*}, Q_{\widehat{\theta}}\right) \quad \text { (by Lemma 16.1) } \\
& \quad \leq 2 D_{\mathcal{B}}\left(Q^{*}, v_{n}\right)+2 D_{\mathcal{B}}\left(v_{n}, Q_{\widehat{\theta}}\right) \quad \text { (by the triangle inequality) } \\
& \quad \leq 2 D_{\mathcal{B}}\left(Q^{*}, v_{n}\right)+2 D_{\mathcal{B}}\left(v_{n}, Q_{\theta}\right) \quad \text { (by the definition of } Q_{\widehat{\theta}} \text { ) } \\
& \quad \leq 4 D_{\mathcal{B}}\left(Q^{*}, v_{n}\right)+2 D_{\mathcal{B}}\left(Q^{*}, Q_{\theta}\right) \quad \text { (again by the triangle inequality). }
\end{aligned}
$$

Lemma 16.2 now applies to the first term on the right-hand side. Thus, we conclude that if the Bayes rule is in $\mathcal{A}$, then for all $\epsilon \geq 4 \inf _{Q_{\theta}} D_{\mathcal{B}}\left(Q^{*}, Q_{\theta}\right)$,

$$
\mathbf{P}\left\{L\left(g_{\widehat{\theta}}\right)-L^{*}>\epsilon\right\} \leq 8(n+1)^{2 V_{\mathcal{A}}} e^{-n \epsilon^{2} / 2^{11}}
$$

The constant in the exponent may be improved significantly by more careful analysis. In other words, if the Bayes rule is in $\mathcal{A}$ and the true distribution is close to the parametric family in the generalized Kolmogorov-Smirnov distance specified above, then the minimum distance rule still performs close to the Bayes error. Unfortunately, we cannot say the same if $\mathcal{A}$ does not contain the Bayes rule. Empirical error minimization, discussed in the next section, is however very robust in all situations.

### 16.4 Empirical Error Minimization

In this section we explore the connection between parametric classification and rule selection by minimizing the empirical error, studied in Chapter 12.

Consider the class $\mathcal{C}$ of classifiers of the form

$$
\phi(x)= \begin{cases}0 & \text { if }(1-p) f_{\theta_{0}}(x) \geq p f_{\theta_{1}}(x) \\ 1 & \text { otherwise }\end{cases}
$$

where $p \in[0,1]$, and $\theta_{0}, \theta_{1} \in \Theta$. The parametric assumption means that the Bayes rule is contained in $\mathcal{C}$. Then it is a very natural approach to minimize the empirical error probability

$$
\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi\left(X_{i}\right) \neq F Y_{i}\right\}},
$$

measured on the training data $D_{n}$ over classifiers $\phi$ in the class $\mathcal{C}$. Denote the empirically selected rule (i.e., the one minimizing $\left.\widehat{L}_{n}(\phi)\right)$ by $\phi_{n}^{*}$. For most typical parametric classes $\Theta$, the VC dimension $V_{\mathcal{C}}$ is finite. Therefore, as a straightforward consequence of Theorem 12.6, we have

Corollary 16.1. If both conditional distributions are contained in the parametric family $\mathcal{P}_{\Theta}$, then for the error probability $L\left(\phi_{n}^{*}\right)=\mathbf{P}\left\{\phi_{n}^{*}(X) \neq Y \mid D_{n}\right\}$ of the empirically optimal rule $\phi_{n}^{*}$, we have for every $n$ and $\epsilon>0$

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-L^{*}>\epsilon\right\} \leq 8 \mathcal{S}(\mathcal{C}, n) e^{-n \epsilon^{2} / 128} .
$$

The result above means that $O(\sqrt{\log n / n})$ rate of convergence to the Bayes rule is guaranteed for the empirically optimal rule, whenever the VC dimension $V_{\mathcal{C}}$ is finite. This is the case, for example, for exponential families. If $\mathcal{P}_{\Theta}$ is an exponential family, with densities of the form

$$
f_{\theta}(x)=c \alpha(\theta) \beta(x) \exp \left\{\sum_{i=1}^{k} \pi_{i}(\theta) \psi_{i}(x)\right\}
$$

then by results from Chapter $13, \mathcal{S}(\mathcal{C}, n) \leq(n+1)^{k+1}$. Observe that in this approach, nothing but properties of the class $\mathcal{C}$ are used to derive the $O(\sqrt{\log n / n})$ rate of convergence.

REmARK. Robustness. The method of empirical minimization is clearly extremely robust against errors in the parametric model. Obviously, (see Theorem 12.6) if the true conditional distributions are not contained in the class $\mathcal{P}_{\Theta}$, then $L^{*}$ can be replaced in Corollary 16.1 by $\inf _{\phi \in \mathcal{C}} \mathbf{P}\{\phi(X) \neq Y\}$. If the model is fairly good, then this number should be very close to the Bayes risk $L^{*}$.

REmARK. LOWER BOUNDS. The results of Chapter 14 imply that for some distributions the error probability of the selected rule is about $O(1 / \sqrt{n})$ away from the Bayes risk. In the parametric case, however, since the class of distributions is restricted by the parametric model, this is not necessarily true. In some cases, a much faster rate of convergence is possible than the $O(\sqrt{\log n / n})$ rate guaranteed by Corollary 16.1. See, for example, Problem 16.3, where an example is given in which the error rate is $O(1 / n)$.

## Problems and Exercises

Problem 16.1. Show that if both conditional distributions of $X$, given $Y=0$ and $Y=1$, are gaussian, then the Bayes decision is quadratic, that is, it can be written as

$$
g^{*}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{d+d(d+1) / 2} a_{i} \psi_{i}(x)+a_{0} \geq 0 \\ 1 & \text { otherwise }\end{cases}
$$

where the functions $\psi_{i}(x)$ are either of the form $x^{(i)}$ (the $i$-th component of the vector $x$ ), or $x^{(i)} x^{(j)}$, and $a_{0}, \ldots, a_{d+d(d+1) / 2} \in \mathcal{R}$.

Problem 16.2. Let $f$ be the normal density on the real line with mean $m$ and standard deviation $\sigma^{2}$, and we draw an i.i.d. sample $X_{1}, \ldots, X_{n}$ from $f$, and set

$$
\widehat{m}=\frac{1}{n} \sum_{i=1}^{n} X_{i} \quad \text { and } \quad \widehat{\sigma}^{2}=\frac{1}{n} \sum_{i=1}^{n}\left(X_{i}-\widehat{m}\right)^{2}
$$

Show that $\mathbf{E}\{|m-\widehat{m}|\}=O(1 / \sqrt{n})$ and $\mathbf{E}\{|\sigma-\widehat{\sigma}|\}=O(1 / \sqrt{n})$ by using Chebyshev's inequality. Show that this rate is in fact tight. Prove also that the result remains true if $n$ is replaced by $N$, a binomial $(n, p)$ random variable independent of $X_{1}, \ldots, X_{n}$, where $p \in(0,1)$. That is, $\widehat{m}$ becomes $(1 / N) \sum_{i=1}^{N} X_{i}$ if $N>0$ and 0 otherwise, and similarly for $\widehat{\sigma}$.

Problem 16.3. Assume that $p=1 / 2$, and that both class-conditional densities $f_{0}$ and $f_{1}$ are gaussian on $\mathcal{R}$ with unit variance, but different means. We use the maximum likelihood estimates $\widehat{m}_{0}, \widehat{m}_{1}$ of the conditional means $m_{0}=\mathbf{E}\{X \mid Y=0\}$ and $m_{1}=\mathbf{E}\{X \mid Y=1\}$ to obtain the plug-in classifier $g_{\widehat{\theta}}$. Show that $\mathbf{E}\left\{\left(\widehat{m}_{0}-m_{0}\right)^{2}\right\}=O(1 / n)$. Then go on to show that $\mathbf{E}\left\{L\left(g_{\widehat{\theta}}\right)\right\}-L^{*} \leq O(1 / n)$.

Problem 16.4. Show that the following classes of densities on $\mathcal{R}$ constitute exponential families:
(1) gaussian family:

$$
\left\{\frac{e^{-\frac{(x-m)^{2}}{2 \sigma^{2}}}}{\sqrt{2 \pi} \sigma}: m \in \mathcal{R}, \sigma>0\right\}
$$

(2) gamma family:

$$
\left\{\frac{x^{\alpha-1} e^{-x / \beta}}{\beta^{\alpha} \Gamma(\alpha)} I_{\{x>0\}}: \alpha, \beta>0\right\}
$$

(3) beta family:

$$
\left\{\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha) \Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1} I_{\{x \in 0,1\}}: \alpha, \beta>0\right\}
$$

(4) Rayleigh family:

$$
\left\{\theta^{\lambda} \beta^{1-\lambda} e^{-\beta^{2} /(2 \theta)} x^{\lambda} e^{-\theta x^{2} / 2} \sum_{j=0}^{\infty} \frac{1}{j!\Gamma(j+\lambda)}\left(\frac{x}{2}\right)^{2 j+\lambda-1}: \theta, \beta>0, \lambda \geq 0\right\}
$$

Problem 16.5. This exercise shows that one-parameter classes may be incredibly rich. Let $\mathcal{C}$ be the class of rules of the form

$$
g_{\theta}(x)= \begin{cases}1 & \text { if } x \in A+\theta \\ 0 & \text { if } x \notin A+\theta\end{cases}
$$

where $x \in \mathcal{R}, \theta \in \mathcal{R}$ is a parameter, and $A$ is a union of intervals on the real line. Equivalently, $g_{\theta}=I_{A+\theta}$. Let

$$
A=\bigcup_{i: b_{i}=1}\left[i-\frac{1}{2}, i+\frac{1}{2}\right)
$$

where $b_{1}, b_{2}, \ldots$ are bits, obtained as follows: first list all sequences of length 1 , then those of length 2 , et cetera, so that $\left(b_{1}, b_{2}, \ldots\right)$ is a concatenation of $(1,0),(1,1,1,0,0,1,0,0)$, et cetera.
(1) Show that for all $n$, there exists a set $\left\{x_{1}, \ldots, x_{n}\right\} \subset \mathcal{R}$ that can be shattered by a set from $\{A+\theta: \theta \in \mathcal{R}\}$. Conclude that the vc dimension of $\mathcal{C}$ is infinite.
(2) If we use $\mathcal{C}$ for empirical error minimization and $L^{*}=0$, what can you say about $\mathbf{E}\left\{L_{n}\right\}$, the probability of error of the selected rule?

Problem 16.6. Continuation. Let $X$ be uniform on $\theta+[i-1 / 2, i+1 / 2)$ with probability $1 / 2^{i}, i \geq 1$. Set $Y=b_{i}$ if $X \in \theta+[i-1 / 2, i+1 / 2)$, so that $L^{*}=0$.
(1) Derive the class of Bayes rules.
(2) Work out the details of the generalized Kolmogorov-Smirnov distance minimizer based on the class of (1).
(3) Provide the best upper bound on $\mathbf{E}\left\{L_{n}\right\}$ you can get with any method.
(4) Regardless of discrimination, how would you estimate $\theta$ ? Derive the asymptotic behavior of $\mathbf{E}\left\{L_{n}\right\}$ for the plug-in rule based on your estimate of $\theta$.

Problem 16.7. If $\mathcal{F}_{\Theta}=\left\{\right.$ all uniform densities on rectangles of $\left.\mathcal{R}^{d}\right\}$ and if we use the maximum likelihood estimates of $p, \theta_{0}, \theta_{1}$ derived in the text, show that $\mathbf{E}\left\{L_{n}\right\}-L^{*}=$ $O(1 / n)$.

Problem 16.8. Continued. Assume that the true class-conditional densities are $f_{0}, f_{1}$, and $f_{0}, f_{1} \notin \mathcal{F}_{\theta}$. With the same maximum likelihood estimates given above, find $f_{0}, f_{1}$ for which $\mathbf{E}\left\{L_{n}\right\} \rightarrow 1$, yet $L^{*}=0$.

Problem 16.9. Continued. Show that the $O(1 / n)$ rate above can, in general, not be improved.

Problem 16.10. Show that

$$
\mathbf{P}\left\{D_{\mathrm{KS}}\left(v_{n}, P_{\theta^{*}}\right)>\epsilon\right\} \leq 4 n^{d} e^{-2(n-d) \epsilon^{2}}
$$

by noting that the supremum in the definition of $D_{\mathrm{KS}}$ may be replaced by a maximum over $n^{d}$ carefully selected points of $\mathcal{R}^{d}$. Hint: Use the idea of fingering from Chapters 4 and 12.

Problem 16.11. Prove Lemma 16.2 by following the line of the proof of the VapnikChervonenkis inequality (Theorem 12.5). Hint: In the second symmetrization step observe that

$$
\sup _{A \in \mathcal{B}}\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i}\left(2 Y_{i}-1\right) I_{\left\{X_{i} \in A\right\}}\right|
$$

has the same distribution as that of

$$
\sup _{A \in \mathcal{B}}\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} I_{\left\{X_{i} \in A\right\}}\right|
$$

Problem 16.12. Minimum distance density estimation. Let $\mathcal{F}_{\Theta}=\left\{f_{\theta}: \theta \in \Theta\right\}$ be a parametric class of densities on $\mathcal{R}^{d}$, and assume that the i.i.d. sample $X_{1}, \ldots, X_{n}$ is drawn from the density $f_{\theta} \in \mathcal{F}_{\Theta}$. Define the class of sets

$$
\mathcal{A}=\left\{\left\{x \in \mathcal{R}^{d}: f_{\theta_{1}}(x)>f_{\theta_{2}}(x)\right\}: \theta_{1}, \theta_{2} \in \Theta\right\}
$$

and define the minimum distance estimate of $\theta$ by

$$
\widehat{\theta}=\underset{\theta^{\prime} \in \Theta}{\arg \min } D_{\mathcal{A}}\left(P_{\theta^{\prime}}, \mu_{n}\right),
$$

where $P_{\theta^{\prime}}$ is the distribution belonging to the density $f_{\theta^{\prime}}, \mu_{n}$ is the empirical distribution defined by $X_{1}, \ldots, X_{n}$, and $D_{\mathcal{A}}$ is the generalized Kolmogorov-Smirnov distance defined by

$$
D_{\mathcal{A}}(P, Q)=\sup _{A \in \mathcal{A}}|P(A)-Q(A)| .
$$

Prove that if $\mathcal{A}$ has finite vc dimension, then

$$
\mathbf{E}\left\{\int\left|f_{\widehat{\theta}}(x)-f_{\theta}(x)\right| d x\right\}=O(1 / \sqrt{n})
$$

For more information on minimum distance density estimation, see Yatracos (1985) and Devroye (1987). Hint: Follow the steps listed below:
(1) $\int\left|f_{\widehat{\theta}}-f_{\theta}\right| \leq 2 D_{\mathcal{A}}\left(P_{\widehat{\theta}}, P_{\theta}\right)$ (use Scheffé's theorem).
(2) $D_{\mathcal{A}}\left(P_{\widehat{\theta}}, P_{\theta}\right) \leq 2 D_{\mathcal{A}}\left(P_{\theta}, \mu_{n}\right)$ (proceed as in the text).
(3) Finish the proof by applying Alexander's improvement of the Vapnik-Chervonenkis inequality (Theorem 12.10).

## 17

## Generalized Linear Discrimination

The classifiers we study here have their roots in the Fourier series estimate or other series estimates of an unknown density, potential function methods (see Chapter 10), and generalized linear classifiers. All these estimators can be put into the following form: classify $x$ as belonging to class 0 if

$$
\sum_{j=1}^{k} a_{n, j} \psi_{j}(x) \leq 0
$$

where the $\psi_{j}$ 's are fixed functions, forming a base for the series estimate, $a_{n, j}$ is a fixed function of the training data, and $k$ controls the amount of smoothing. When the $\psi_{j}$ 's are the usual trigonometric basis, then this leads to the Fourier series classifier studied by Greblicki and Pawlak (1981; 1982). When the $\psi_{j}$ 's form an orthonormal system based upon Hermite polynomials, we obtain the classifiers studied by Greblicki (1981), and Greblicki and Pawlak (1983; 1985). When $\left\{\psi_{j}(x)\right\}$ is the collection of all products of components of $x$ (such as $1,\left(x^{(i)}\right)^{k},\left(x^{(i)}\right)^{k}\left(x^{(j)}\right)^{l}$, etcetera), we obtain the polynomial method of Specht (1971).

We start with classification based on Fourier series expansion, which has its origins in Fourier series density estimation, which, in turn, was developed by Cencov (1962), Schwartz (1967), Kronmal and Tarter (1968), Tarter and Kronmal (1970), and Specht (1971). Its use in classification was considered by Greblicki (1981), Specht (1971), Greblicki and Pawlak (1981; 1982; 1983), and others.

### 17.1 Fourier Series Classification

Let $f$ be the density of $X$, and assume that $f \in L_{2}(\lambda)$, that is, $f \geq 0, \int f(x) d x=1$ and $\int f^{2}(x) d x<\infty$, and recall that $\lambda$ denotes the Lebesgue measure on $\mathcal{R}^{d}$. Let $\psi_{1}, \psi_{2}, \ldots$ be a complete orthonormal set of bounded functions in $L_{2}(\lambda)$ with $\sup _{j, x}\left|\psi_{j}(x)\right| \leq B<\infty$. An orthonormal set of functions $\left\{\psi_{j}\right\}$ in $L_{2}(\mu)$ is such that $\int \psi_{j} \psi_{i}=I_{\{i=j\}}$ for all $i, j$. It is complete if $\int \alpha \psi_{i}=0$ for all $i$ for some function $\alpha \in L_{2}(\lambda)$ implies that $\alpha \equiv 0$ almost everywhere (with respect to $\lambda$ ). If $f \in L_{2}(\lambda)$, then the class-conditional densities $f_{0}$ and $f_{1}$ exist, and are in $L_{2}(\lambda)$. Then the function

$$
\alpha(x)=p f_{1}(x)-(1-p) f_{0}(x)=(2 \eta(x)-1) f(x)
$$

is in $L_{2}(\lambda)$ as well. Recall that the Bayes decision is given by

$$
g^{*}(x)= \begin{cases}0 & \text { if } \alpha(x) \leq 0 \\ 1 & \text { otherwise }\end{cases}
$$

Let the bounded functions $\psi_{1}, \psi_{2}, \ldots$ form a complete orthonormal basis, and let the Fourier representation of $\alpha$ be given by

$$
\alpha=\sum_{j=1}^{\infty} \alpha_{j} \psi_{j}
$$

The above convergence is understood in $L_{2}(\lambda)$, that is, the infinite sum means that

$$
\lim _{k \rightarrow \infty} \int\left(\alpha(x)-\sum_{j=1}^{k} \alpha_{j} \psi_{j}(x)\right)^{2} d x=0
$$

The coefficients are given by

$$
\alpha_{j}=\int \psi_{j}(x) \alpha(x) d x
$$

We approximate $\alpha(x)$ by a truncation of its Fourier representation to finitely many terms, and use the data $D_{n}$ to estimate the coefficients appearing in this sum. Formally, consider the classification rule

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{j=1}^{k_{n}} \alpha_{n, j} \psi_{j}(x) \leq 0  \tag{17.1}\\ 1 & \text { otherwise }\end{cases}
$$

where the estimates $\alpha_{n, j}$ of $\alpha_{j}$ are very easy to compute:

$$
\alpha_{n, j}=\frac{1}{n} \sum_{i=1}^{n}\left(2 Y_{i}-1\right) \psi_{j}\left(X_{i}\right)
$$

Before discussing the properties of these rules, we list a few examples of complete orthonormal systems on the real line. Some of these systems contain functions
on a bounded interval. These, of course, can only be used if the distribution of $X$ is concentrated on an interval. The completeness of these systems may typically be checked via the Stone-Weierstrass theorem (Theorem A.9). A general way of producing complete orthonormal systems on $\mathcal{R}^{d}$ is taking products of onedimensional functions of the $d$ components, as sketched in Problem 17.1. For more information on orthogonal series, we refer to Sansone (1969), Szegő (1959), and Zygmund (1959).
(1) The standard trigonometric basis on the interval $[-\pi, \pi]$ is formed by the functions

$$
\psi_{0}(x)=\frac{1}{\sqrt{2 \pi}}, \quad \psi_{2 i-1}(x)=\frac{\cos (i x)}{\sqrt{\pi}}, \quad \psi_{2 i}(x)=\frac{\sin (i x)}{\sqrt{\pi}}, \quad i=1,2, \ldots
$$

This is a complete orthonormal system in $L_{2}([-\pi, \pi])$.
(2) The Legendre polynomials form a complete orthonormal basis over the interval $[-1,1]$ :

$$
\psi_{i}(x)=\sqrt{\frac{2 i+1}{2}} \frac{1}{2^{i} i!} \frac{d^{i}}{d x^{i}}\left(\left(x^{2}-1\right)^{i}\right), \quad i=0,1,2, \ldots
$$

(3) The set of Hermite functions is complete and orthonormal over the whole real line:

$$
\psi_{i}(x)=\frac{e^{x^{2} / 2}}{\sqrt{2^{i} i!\sqrt{\pi}}} \frac{d^{i}}{d x^{i}}\left(e^{-x^{2}}\right), \quad i=0,1,2, \ldots
$$

(4) Functions of the Laguerre basis are defined on [0, $\infty$ ) by

$$
\psi_{i}(x)=\left(\frac{\Gamma(i+1)}{\Gamma(i+\alpha+1)} x^{-\alpha} e^{x}\right)^{1 / 2} \frac{1}{i!} \frac{d^{i}}{d x^{i}}\left(x^{i+\alpha} e^{-x}\right), \quad i=0,1,2, \ldots,
$$

where $\alpha>-1$ is a parameter of the system. A complete orthonormal system over the whole real line may be obtained by

$$
\psi_{2 i}^{\prime}(x)=\psi_{i}(x) I_{\{x \geq 0\}}, \quad \psi_{2 i+1}^{\prime}(x)=\psi_{i}(-x) I_{\{x<0\}}, \quad i=0,1,2, \ldots
$$

(5) The Haar basis contains functions on [0, 1] that take three different values. It is orthonormal and complete. Functions with finitely many values have computational advantages in pattern recognition. Write the integer $i \geq 0$ as $i=2^{k}+j$, where $k=\left\lfloor\log _{2} i\right\rfloor$ (i.e., $0 \leq j<2^{k}$ ). Then

$$
\psi_{i}(x)= \begin{cases}2^{k / 2} & \text { if } x \in\left((j-1) / 2^{k},(j-1 / 2) / 2^{k}\right) \\ 2^{-k / 2} & \text { if } x \in\left((j-1 / 2) / 2^{k}, j / 2^{k}\right) \\ 0 & \text { otherwise }\end{cases}
$$

(6) Functions on [0, 1] of the Rademacher basis take only two values, -1 and 1. It is an orthonormal system, but it is not complete:

$$
\psi_{0}(x) \equiv 1, \quad \psi_{i}(x)=(-1)^{\left\lfloor 2^{i} x\right\rfloor}, \quad i=1,2, \ldots
$$

The basis may be completed as follows: write $i=\sum_{j=1}^{k} 2^{i_{j}}$ such that $0 \leq$ $i_{1}<i_{2}<\cdots<i_{k}$. This form is unique. Define

$$
\psi^{\prime}(x)=\psi_{i_{1}}(x) \cdots \psi_{i_{k}}(x)
$$

where the $\psi_{j}$ 's are the Rademacher functions. The resulting basis $\psi_{0}^{\prime}, \psi_{1}^{\prime}, \ldots$ is orthonormal and complete, and is called the Walsh basis.

There is no system of basis functions that is better than another for all distributions. In selecting the basis of a Fourier series rule, the designer must use a mixture of intuition, error estimation, and computational concerns. We have the following consistency theorem:

Theorem 17.1. Let $\left\{\psi_{0}, \psi_{1}, \ldots\right\}$ be a complete orthonormal basis on $\mathcal{R}^{d}$ such that for some $B<\infty,\left|\psi_{i}(x)\right| \leq B$ for each $i$ and $x$. Assume that the class-conditional densities $f_{0}$ and $f_{1}$ exist and are in $L_{2}(\lambda)$. If

$$
k_{n} \rightarrow \infty \text { and } \frac{k_{n}}{n} \rightarrow 0 \text { as } n \rightarrow \infty
$$

then the Fourier series classification rule defined in (17.1) is consistent:

$$
\lim _{n \rightarrow \infty} \mathbf{E} L\left(g_{n}\right)=L^{*}
$$

If

$$
k_{n} \rightarrow \infty \text { and } \frac{k_{n} \log n}{n} \rightarrow 0 \text { as } n \rightarrow \infty
$$

then

$$
\lim _{n \rightarrow \infty} L\left(g_{n}\right)=L^{*}
$$

with probability one, that is, the rule is strongly consistent.

Proof. First observe that the $\alpha_{n, j}$ 's are unbiased estimates of the $\alpha_{j}$ 's:

$$
\begin{aligned}
\mathbf{E}\left\{\alpha_{n, j}\right\} & =\mathbf{E}\left\{(2 Y-1) \psi_{j}(X)\right\}=\mathbf{E}\left\{\mathbf{E}\left\{(2 Y-1) \psi_{j}(X) \mid X\right\}\right\} \\
& =\mathbf{E}\left\{\psi_{j}(X) \mathbf{E}\{(2 Y-1) \mid X\}\right\}=\mathbf{E}\left\{(2 \eta(X)-1) \psi_{j}(X)\right\} \\
& =\int(2 \eta(x)-1) \psi_{j}(x) f(x) d x=\int \psi_{j}(x) \alpha(x) d x \\
& =\alpha_{j}
\end{aligned}
$$

and that their variance can be bounded from above as follows:

$$
\begin{aligned}
\operatorname{Var}\left\{\alpha_{n, j}\right\} & =\frac{\operatorname{Var}\left\{\psi_{j}\left(X_{1}\right)\left(2 \eta\left(X_{1}\right)-1\right)\right\}}{n} \\
& =\frac{\int \psi_{j}^{2}(x)(2 \eta(x)-1)^{2} f(x) d x-\alpha_{j}^{2}}{n} \\
& \leq \frac{B^{2}-\alpha_{j}^{2}}{n} \leq \frac{B^{2}}{n}
\end{aligned}
$$

where we used the boundedness of the $\psi_{j}$ 's. By Parseval's identity,

$$
\int \alpha^{2}(x) d x=\sum_{j=1}^{\infty} \alpha_{j}^{2}
$$

Therefore, exploiting orthonormality of the $\psi_{j}$ 's, we have

$$
\begin{aligned}
& \int\left(\alpha(x)-\sum_{j=1}^{k_{n}} \alpha_{n, j} \psi_{j}(x)\right)^{2} d x \\
&= \int \alpha^{2}(x) d x+\int\left(\sum_{j=1}^{k_{n}} \alpha_{n, j} \psi_{j}(x)\right)^{2} d x \\
&-2 \int \alpha(x) \sum_{j=1}^{k_{n}} \alpha_{n, j} \psi_{j}(x) d x \\
&= \sum_{j=1}^{\infty} \alpha_{j}^{2}+\sum_{j=1}^{k_{n}} \alpha_{n, j}^{2}-2 \sum_{j=1}^{k_{n}} \alpha_{j} \alpha_{n, j} \\
&= \sum_{j=1}^{k_{n}}\left(\alpha_{n, j}-\alpha_{j}\right)^{2}+\sum_{j=k_{n}+1}^{\infty} \alpha_{j}^{2}
\end{aligned}
$$

Thus, the expected $L_{2}$-error is bounded as follows:

$$
\begin{aligned}
\mathbf{E} & \left\{\int\left(\alpha(x)-\sum_{j=1}^{k_{n}} \alpha_{n, j} \psi_{j}(x)\right)^{2} d x\right\} \\
& =\mathbf{E}\left\{\sum_{j=1}^{k_{n}}\left(\alpha_{n, j}-\alpha_{j}\right)^{2}\right\}+\sum_{j=k_{n}+1}^{\infty} \alpha_{j}^{2} \\
& =\sum_{j=1}^{k_{n}} \operatorname{Var}\left\{\alpha_{n, j}\right\}+\sum_{j=k_{n}+1}^{\infty} \alpha_{j}^{2} \\
& \leq \frac{k_{n} B^{2}}{n}+\sum_{j=k_{n}+1}^{\infty} \alpha_{j}^{2} .
\end{aligned}
$$

Since $\sum_{j=1}^{\infty} \alpha_{j}^{2}<\infty$, the second term tends to zero if $k_{n} \rightarrow \infty$. If at the same time $k_{n} / n \rightarrow 0$, then the expected $L_{2}$-error converges to zero, that is, the estimate is consistent in $L_{2}$. Now, convergence of the error probability follows from Problem 2.11 .

To prove strong consistency (i.e., convergence with probability one), fix $\epsilon>0$, and assume that $n$ is so large that

$$
\sum_{j=k_{n}+1}^{\infty} \alpha_{j}^{2}<\epsilon / 2
$$

Then

$$
\begin{aligned}
& \mathbf{P}\left\{\int\left(\alpha(x)-\sum_{j=1}^{k_{n}} \alpha_{n, j} \psi_{j}(x)\right)^{2} d x>\epsilon\right\} \\
& =\mathbf{P}\left\{\sum_{j=1}^{k_{n}}\left(\alpha_{n, j}-\alpha_{j}\right)^{2}+\sum_{j=k_{n}+1}^{\infty} \alpha_{j}^{2}>\epsilon\right\} \\
& \quad \leq \mathbf{P}\left\{\sum_{j=1}^{k_{n}}\left(\alpha_{n, j}-\alpha_{j}\right)^{2}>\epsilon / 2\right\} \\
& \quad \leq \sum_{j=1}^{k_{n}} \mathbf{P}\left\{\left(\alpha_{n, j}-\alpha_{j}\right)^{2}>\epsilon /\left(2 k_{n}\right)\right\}
\end{aligned}
$$

(by the union bound)

$$
\begin{aligned}
& =\sum_{j=1}^{k_{n}} \mathbf{P}\left\{\left|\frac{1}{n} \sum_{i=1}^{n}\left(\psi_{j}\left(X_{i}\right)\left(2 Y_{i}-1\right)-\mathbf{E}\left\{\psi_{j}(X)\left(2 Y_{i}-1\right)\right\}\right)\right|>\sqrt{\frac{\epsilon}{2 k_{n}}}\right\} \\
& \leq 2 k_{n} e^{-n \epsilon /\left(4 k_{n} B^{2}\right)},
\end{aligned}
$$

where we used Hoeffding's inequality (Theorem 8.1). Because $k_{n} \log n / n \rightarrow 0$, the upper bound is eventually smaller than $e^{-2 \log n}=n^{-2}$, which is summable. The Borel-Cantelli lemma yields strong $L_{2}$ consistency. Strong consistency of the classifier then follows from Problem 2.11.

REMARK. It is clear from the inequality

$$
\mathbf{E}\left\{\int\left(\alpha(x)-\sum_{j=1}^{k_{n}} \alpha_{n, j} \psi_{j}(x)\right)^{2} d x\right\} \leq \frac{k_{n} B^{2}}{n}+\sum_{j=k_{n}+1}^{\infty} \alpha_{j}^{2}
$$

that the rate of convergence is determined by the choice of $k_{n}$. If $k_{n}$ is small, then the first term, which corresponds to the estimation error, is small, but the approximation error, expressed by the second term, is larger. For large $k_{n}$, the
situation is reversed. The optimal choice of $k_{n}$ depends on how fast the second term goes to zero as $k_{n}$ grows. This depends on other properties of $f$, such as the size of its tail and its smoothness.

REmark. Consistency in Theorem 17.1 is not universal, since we needed to assume the existence of square integrable conditional densities. This, however, is not a restrictive assumption in some practical situations. For example, if the observations are corrupted by additive gaussian noise, then the conditions of the theorem hold (Problem 17.2). However, if $\mu$ does not have a density, the method may be inconsistent (see Problem 17.4).

Fourier series classifiers have two rather unattractive features in general:
(i) They are not invariant under translations of the coordinate space.
(ii) Most series classifiers are not local in nature-points at arbitrary distances from $x$ affect the decision at $x$. In kernel and nearest neighbor rules, the locality is easily controlled.

### 17.2 Generalized Linear Classification

When $X$ is purely atomic or singular continuous, Theorem 17.1 is not applicable. A theme of this book is that pattern recognition can be developed in a distributionfree manner since, after all, the distribution of $(X, Y)$ is not known. Besides, even if we had an i.i.d. sample $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ at our disposal, we do not know of any test for verifying whether $X$ has a square integrable density. So, we proceed a bit differently to develop universally consistent rules. To generalize Fourier series classifiers, let $\psi_{1}, \psi_{2}, \ldots$ be bounded functions on $\mathcal{R}^{d}$. These functions do not necessarily form an orthonormal basis of $L_{2}$. Consider the classifier

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{j=1}^{k_{n}} \alpha_{n, j} \psi_{j}(x) \leq 0 \\ 1 & \text { otherwise }\end{cases}
$$

where the coefficients $\alpha_{n, j}$ are some functions of the data $D_{n}$. This may be viewed as a generalization of Fourier series rules, where the coefficients were unbiased estimates of the Fourier coefficients of $\alpha(x)=(2 \eta(x)-1) f(x)$. Here we will consider some other choices of the coefficients. Observe that $g_{n}$ is a generalized linear classifier, as defined in Chapter 13. An intuitively appealing way to determine the coefficients is to pick them to minimize the empirical error committed on $D_{n}$, as in empirical risk minimization. The critical parameter here is $k_{n}$, the number of basis functions used in the linear combination. If we keep $k_{n}$ fixed, then as we saw in Chapter 13, the error probability of the selected rule converges quickly to that of the best classifier of the above form. However, for some distributions, this infimum may be far from the Bayes risk, so it is useful to let $k_{n}$ grow as $n$ becomes
larger. However, choosing $k_{n}$ too large may result in overfitting the data. Using the terminology introduced in Chapter 12 , let $\mathcal{C}^{\left(k_{n}\right)}$ be the class of classifiers of the form

$$
\phi(x)= \begin{cases}0 & \text { if } \sum_{j=1}^{k_{n}} a_{j} \psi_{j}(x) \leq 0 \\ 1 & \text { otherwise }\end{cases}
$$

where the $a_{j}$ 's range through $\mathcal{R}$. Choose the coefficients to minimize the empirical error

$$
\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi\left(X_{i}\right) \neq Y_{i}\right\}} .
$$

Let $g_{n}=\phi_{n}^{*}$ be the corresponding classifier. We recall from Chapter 13 that the VC dimension of $\mathcal{C}^{\left(k_{n}\right)}$ is not more than $k_{n}$. Therefore, by Theorem 13.12, for every $n>2 k_{n}+1$,

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}^{\left(k k_{n}\right)}} L(\phi)>\epsilon\right\} \leq 8 e^{n \mathcal{H}\left(\frac{k_{n}}{n}\right)} e^{-n \epsilon^{2} / 128}
$$

where $\mathcal{H}$ is the binary entropy function. The right-hand side is $e^{-n \epsilon^{2} / 128+o(n)}$ if $k_{n}=$ $o(n)$. However, to obtain consistency, we need to know how close $\inf _{\phi \in \mathcal{C}^{\left(k_{n}\right)}} L(\phi)$ is to $L^{*}$. This obviously depends on the choice of the $\psi_{j}$ 's, as well as on the distribution. If $k_{n}$ is not allowed to grow with $n$, and is bounded by a number $K$, then universal consistency eludes us, as for some distribution $\inf _{\phi \in \mathcal{C}^{(K)}} L(\phi)-L^{*}>0$. It follows from Theorem 2.2 that for every $B>0$

$$
\begin{aligned}
& \inf _{\phi \in \mathcal{C}\left(k_{n}\right)} L(\phi)-L^{*} \\
& \quad \leq \inf _{a_{1}, \ldots, a_{k_{n}}} \int_{\|x\| \leq B}\left|(2 \eta(x)-1)-\sum_{j=1}^{k_{n}} a_{j} \psi_{j}(x)\right| \mu(d x)+\int_{\|x\|>B} \mu(d x) .
\end{aligned}
$$

The limit supremum of the right-hand side can be arbitrarily close to zero for all distributions if $k_{n} \rightarrow \infty$ as $n \rightarrow \infty$ and the set of functions

$$
\bigcup_{k=1}^{\infty}\left\{\sum_{j=1}^{k} a_{j} \psi_{j}(x) ; a_{1}, a_{2}, \ldots \in \mathcal{R}\right\}
$$

is dense in $L_{1}(\mu)$ on balls of the form $\{x:\|x\| \leq B\}$ for all $\mu$. This means that given any probability measure $\mu$, and function $f$ with $\int|f| d \mu<\infty$, for every $\epsilon>0$, there exists an integer $k$ and coefficients $a_{1}, \ldots, a_{k}$ such that

$$
\int\left|f(x)-\sum_{j=1}^{k} a_{j} \psi_{j}(x)\right| \mu(d x)<\epsilon
$$

Thus, we have obtained the following consistency result:

Theorem 17.2. Let $\psi_{1}, \psi_{2}, \ldots$ be a sequence of functions such that the set of all finite linear combinations of the $\psi_{j}$ 's is dense in $L_{1}(\mu)$ on balls of the form $\{x:\|x\| \leq B\}$ for any probability measure $\mu$. Then then $g_{n}$ is strongly universally consistent when

$$
k_{n} \rightarrow \infty \text { and } \frac{k_{n}}{n} \rightarrow 0 \text { as } n \rightarrow \infty
$$

Remark. To see that the statement of Theorem 17.2 is not vacuous, note that denseness in $L_{1}(\mu)$ on balls follows from denseness with respect to the supremum norm on balls. Denseness in the $L_{\infty}$ sense may be checked by the Stone-Weierstrass theorem (Theorem A.9). For example, the class of all polynomial classifiers satisfies the conditions of the theorem. This class can be obtained by choosing the functions $\psi_{1}, \psi_{2}, \ldots$ as the simple polynomials $x^{(1)}, \ldots, x^{(d)}, x^{(1)} x^{(2)}, \ldots$. Similarly, the functions $\psi_{j}$ can be chosen as bases of a trigonometric series.

Remark. The histogram rule. It is clear that Theorem 17.2 can be modified in the following way: let $\psi_{j, k}, j=1, \ldots, k ; k=1,2, \ldots$, be functions such that for every $f \in L_{1}(\mu)$ (with $\mu$ concentrated on a bounded ball) and $\epsilon>0$ there is a $k_{0}$ such that for all $k>k_{0}$ there is a function of the form $\sum_{j=1}^{k} a_{j} \psi_{j, k}$ whose distance from $f$ in $L_{1}(\mu)$ is smaller than $\epsilon$. Let $\mathcal{C}^{\left(k_{n}\right)}$ be the class of generalized linear classifiers based on the functions $\psi_{1, k_{n}}, \ldots, \psi_{k_{n}, k_{n}}$. If $k_{n} \rightarrow \infty$ and $k_{n} / n \rightarrow 0$, then the classifier $g_{n}$ that minimizes the empirical error over $\mathcal{C}^{\left(k_{n}\right)}$ is strongly universally consistent. This modification has an interesting implication. Consider for example functions $\psi_{1, k_{n}}, \ldots, \psi_{k_{n}, k_{n}}$ that are indicators of sets of a partition of $\mathcal{R}^{d}$. Then it is easy to see that the classifier that minimizes the empirical error is just the histogram classifier based on the same partition. The denseness assumption requires that the partition becomes infinitesimally fine as $n \rightarrow \infty$. In fact, we have obtained an alternative proof for the strong universal consistency of the regular histogram rule (Theorem 9.4). The details are left as an exercise (Problem 17.3).

## Problems and Exercises

Problem 17.1. Let $\psi_{1}, \psi_{2}, \ldots$ be a sequence of real-valued functions on a bounded interval $[a, b]$ such that $\int_{a}^{b} \psi_{i}(x) \psi_{j}(x) d x=I_{\{i=j\}}$, and the set of finite linear combinations $\sum_{i=1}^{k} a_{i} \psi_{i}(x)$ is dense in the class of continuous functions on [ $\left.a, b\right]$ with respect to the supremum norm. Define the functions $\Psi_{i_{1}, \ldots, i_{d}}:[a, b]^{d} \rightarrow \mathcal{R}$ by

$$
\Psi_{i_{1}, \ldots, i_{d}}(x)=\psi_{i_{1}}\left(x^{(1)}\right) \cdots \psi_{i_{d}}\left(x^{(d)}\right), \quad i_{1}, \ldots, i_{d} \in\{1,2, \ldots\} .
$$

Show that these functions form a complete orthonormal set of functions on $[a, b]^{d}$.
Problem 17.2. Let $Z$ be an arbitrary random variable on $\mathcal{R}$, and $V$ be a real random variable, independent of $Z$, that has a density $h \in L_{2}(\lambda)$. Show that the density $f$ of the
random variable $X=Z+V$ exists, and $\int f^{2}(x) d x<\infty$. Hint: Use Jensen's inequality to prove that $\int f^{2}(x) d x \leq \int h^{2}(x) d x$.

Problem 17.3. Derive the strong universal consistency of the regular histogram rule from Theorem 17.2, as indicated in the remark following it.

Problem 17.4. Let $\left\{\psi_{0}, \psi_{1}, \psi_{2}, \ldots\right\}$ be the standard trigonometric basis, and consider the classification rule defined in (17.1). Show that the rule is not consistent for the following distribution: given $Y=1$, let $X$ be 0 , and given $Y=0$, let $X$ be uniformly distributed on $[-\pi, \pi]$. Assume furthermore that $\mathbf{P}\{Y=1\}=1 / 2$.

Problem 17.5. The Haar basis is not bounded. Determine whether or not the Laguerre, Hermite, and Legendre bases are bounded.

## 18

## Complexity Regularization

This chapter offers key theoretical results that confirm the existence of certain "good" rules. Although the proofs are constructive-we do tell you how you may design such rules-the computational requirements are often prohibitive. Many of these rules are thus not likely to filter down to the software packages and pattern recognition implementations. An attempt at reducing the computational complexity somewhat is described in the section entitled "Simple empirical covering." Nevertheless, we feel that much more serious work on discovering practical algorithms for empirical risk minimization is sorely needed.

In Chapter 12, the empirical error probability was minimized over a class $\mathcal{C}$ of decision rules $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$. We saw that if the vc dimension $V_{\mathcal{C}}$ of the class is finite, then the error probability of the selected rule is within constant times $\sqrt{V_{\mathcal{C}} \log n / n}$ of that of the best rule in $\mathcal{C}$. Unfortunately, classes with finite vc dimension are nearly always too small, and thus the error probability of the best rule in $\mathcal{C}$ is typically far from the Bayes risk $L^{*}$ for some distribution (see Theorem 18.4). In Chapter 17 we investigated the special classes of generalized linear rules. Theorem 17.2, for example, shows that if we increase the size of the class in a controlled fashion as the sample size $n$ increases, the error probability of the selected rule approaches $L^{*}$ for any distribution. Thus, $V_{\mathcal{C}}$ increases with $n$ !

Theorem 17.2 may be generalized straightforwardly for other types of classifiers. Consider a sequence of classes $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$, containing classifiers of the form $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$. The training data $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$ are used to select a classifier $\phi_{n}^{*}$ by minimizing the empirical error probability $\widehat{L}_{n}(\phi)$ over the class $\mathcal{C}^{\left(k_{n}\right)}$, where the integer $k_{n}$ depends on $n$ in a specified way. The following generalization is based on the proof of Theorem 17.2 (see Problem 18.1):

Theorem 18.1. Assume that $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$ is a sequence of classes of decision rules such that for any distribution of $(X, Y)$,

$$
\lim _{i \rightarrow \infty} \inf _{\phi \in \mathcal{C}^{(i)}} L(\phi)=L^{*}
$$

and that the vC dimensions $V_{\mathcal{C}^{(1)}}, V_{\mathcal{C}^{(2)}}, \ldots$ are all finite. If

$$
k_{n} \rightarrow \infty \quad \text { and } \quad \frac{V_{\mathcal{C}^{\left(k_{n}\right)}} \log n}{n} \rightarrow 0 \quad \text { as } n \rightarrow \infty
$$

then the classifier $\phi_{n}^{*}$ that minimizes the empirical error over the class $\mathcal{C}^{\left(k_{n}\right)}$ is strongly universally consistent.

Theorem 18.1 is the missing link-we are now ready to apply the rich theory of Vapnik-Chervonenkis classes in the construction of universally consistent rules. The theorem does not help us, however, with the choice of the classes $\mathcal{C}^{(i)}$, or the choice of the sequence $\left\{k_{n}\right\}$. Examples for sequences of classes satisfying the condition of Theorem 18.1 include generalized linear decisions with properly chosen basis functions (Chapter 17), or neural networks (Chapter 30).

A word of warning here. The universally consistent rules obtained via Theorem 18.1 may come at a tremendous computational price. As we will see further on, we will often have exponential complexities in $n$ instead of polynomial time that would be obtained if we just minimized the empirical risk over a fixed vc class. The computational complexity of these rules are often incomparably larger than that of some simple universally consistent rules as $k$-nearest neighbor, kernel, or partitioning rules.

### 18.1 Structural Risk Minimization

Let $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$ be a sequence of classes of classifiers, from which we wish to select a sequence of classifiers with the help of the training data $D_{n}$. Previously, we picked $\phi_{n}^{*}$ from the $k_{n}$-th class $\mathcal{C}^{\left(k_{n}\right)}$, where the integer $k_{n}$ is some prespecified function of the sample size $n$ only. The integer $k_{n}$ basically determines the complexity of the class from which the decision rule is selected. Theorem 18.1 shows that under mild conditions on the sequence of classes, it is possible to find sequences $\left\{k_{n}\right\}$ such that the rule is universally consistent. Typically, $k_{n}$ should grow with $n$ in order to assure convergence of the approximation $\operatorname{error}^{\inf }{ }_{\phi \in \mathcal{C}^{(k n)}} L(\phi)-L^{*}$, but it cannot grow too rapidly, for otherwise the estimation error $L\left(\phi_{n}^{*}\right)-\inf _{\phi \in \mathcal{C}^{\left(k_{n}\right)}} L(\phi)$ might fail to converge to zero. Ideally, to get best performance, the two types of error should be about the same order of magnitude. Clearly, a prespecified choice of the complexity $k_{n}$ cannot balance the two sides of the trade-off for all distributions. Therefore, it is important to find methods such that the classifier is selected from a class whose index is automatically determined by the data $D_{n}$. This section deals with such methods.

The most obvious method would be based on selecting a candidate decision rule $\tilde{\phi}_{n, j}$ from every class $\mathcal{C}^{(j)}$ (for example, by minimizing the empirical error over $\mathcal{C}^{(j)}$ ), and then minimizing the empirical error over these rules. However, typically, the VC dimension of

$$
\mathcal{C}^{*}=\bigcup_{j=1}^{\infty} \mathcal{C}^{(j)}
$$

equals infinity, which, in view of results in Chapter 14, implies that this approach will not work. In fact, in order to guarantee

$$
\inf _{j \geq 1} \inf _{\phi \in \mathcal{C}^{(j)}} L(\phi)=L^{*}
$$

it is necessary that

$$
V_{\mathcal{C}^{*}}=\infty
$$

(see Theorems 18.4 and 18.5).
A possible solution for the problem is a method proposed by Vapnik and Chervonenkis (1974c) and Vapnik (1982), called structural risk minimization. First we select a classifier $\tilde{\phi}_{n, j}$ from every class $\mathcal{C}^{(j)}$ which minimizes the empirical error over the class. Then we know from Theorem 12.5 that for every $j$, with very large probability,

$$
L\left(\tilde{\phi}_{n, j}\right) \leq \widehat{L}_{n}\left(\tilde{\phi}_{n, j}\right)+O\left(\sqrt{\frac{V_{\mathcal{C}^{(j)}} \log n}{n}}\right)
$$

Now, pick a classifier that minimizes the upper bound over $j \geq 1$. To make things more precise, for every $n$ and $j$, we introduce the complexity penalty

$$
r(j, n)=\sqrt{\frac{32}{n} V_{\mathcal{C}^{(j)}} \log (e n)}
$$

Let $\tilde{\phi}_{n, 1}, \tilde{\phi}_{n, 2}, \ldots$ be classifiers minimizing the empirical error $\widehat{L}_{n}(\phi)$ over the classes $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$, respectively. For $\phi \in \mathcal{C}^{(j)}$, define the complexity-penalized error estimate

$$
\tilde{L}_{n}(\phi)=\widehat{L}_{n}(\phi)+r(j, n)
$$

Finally, select the classifier $\phi_{n}^{*}$ minimizing the complexity penalized error estimate $\tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right)$ over $j \geq 1$. The next theorem states that this method avoids overfitting the data. The only condition is that each class in the sequence has finite vc dimension.
Theorem 18.2. Let $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$ be a sequence of classes of classifiers such that for any distribution of $(X, Y)$,

$$
\lim _{j \rightarrow \infty} \inf _{\phi \in \mathcal{C}^{(j)}} L(\phi)=L^{*}
$$

Assume also that the VC dimensions $V_{\mathcal{C}^{(1)}}, V_{\mathcal{C}^{(2)}}, \ldots$ are finite and satisfy

$$
\Delta=\sum_{j=1}^{\infty} e^{-V_{c(j)}}<\infty
$$

Then the classification rule $\phi_{n}^{*}$ based on structural risk minimization, as defined above, is strongly universally consistent.

REmARK. Note that the condition on the vc dimensions is satisfied if we insist that $V_{\mathcal{C}^{(j+1)}}>V_{\mathcal{C}^{(j)}}$ for all $j$, a natural assumption. See also Problem 18.3.

Instead of minimizing the empirical error $\widehat{L}_{n}(\phi)$ over the set of candidates $\mathcal{C}^{*}$, the method of structural risk minimization minimizes $\tilde{L}_{n}(\phi)$, the sum of the empirical error, and a term $\sqrt{\frac{32}{n} V_{\mathcal{C}^{(j)}} \log (e n)}$, which increases as the vc dimension of the class $\mathcal{C}^{(j)}$ containing $\phi$ increases. Since classes with larger vc dimension can be considered as more complex than those with smaller vc dimension, the term added to the empirical error may be considered as a penalty for complexity. The idea of minimizing the sum of the empirical error and a term penalizing the complexity has been investigated in various statistical problems by, for example, Rissanen (1983), Akaike (1974), Barron (1985; 1991), Barron and Cover (1991), and Lugosi and Zeger (1996). Barron (1991) minimizes the penalized empirical risk over a suitably chosen countably infinite list of candidates. This approach is close in spirit to the skeleton estimates discussed in Chapter 28 . He makes the connection between structural risk minimization and the minimum description length principle, and obtains results similar to those discussed in this section. The theorems presented here were essentially developed in Lugosi and Zeger (1996).
Proof of Theorem 18.2. We show that both terms on the right-hand side of the following decomposition converge to zero with probability one:

$$
L\left(\phi_{n}^{*}\right)-L^{*}=\left(L\left(\phi_{n}^{*}\right)-\inf _{j \geq 1} \tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right)\right)+\left(\inf _{j \geq 1} \tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right)-L^{*}\right)
$$

For the first term we have

$$
\begin{aligned}
\mathbf{P} & \left\{L\left(\phi_{n}^{*}\right)-\inf _{j \geq 1} \tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right)>\epsilon\right\} \\
& =\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-\tilde{L}_{n}\left(\phi_{n}^{*}\right)>\epsilon\right\} \\
& \leq \mathbf{P}\left\{\sup _{j \geq 1}\left(L\left(\tilde{\phi}_{n, j}\right)-\tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right)\right)>\epsilon\right\} \\
& =\mathbf{P}\left\{\sup _{j \geq 1}\left(L\left(\tilde{\phi}_{n, j}\right)-\widehat{L}_{n}\left(\tilde{\phi}_{n, j}\right)-r(j, n)\right)>\epsilon\right\} \\
& \leq \sum_{j=1}^{\infty} \mathbf{P}\left\{\left|L\left(\tilde{\phi}_{n, j}\right)-\widehat{L}_{n}\left(\tilde{\phi}_{n, j}\right)\right|>\epsilon+r(j, n)\right\} \\
& \leq \sum_{j=1}^{\infty} 8 n^{V_{c(j)}} e^{-n(\epsilon+r(j, n))^{2} / 32} \quad \text { (by Theorem 12.5) }
\end{aligned}
$$

$$
\begin{aligned}
& \leq \sum_{j=1}^{\infty} 8 n^{V_{C}(j)} e^{-n r^{2}(j, n) / 32} e^{-n \epsilon^{2} / 32} \\
& =8 e^{-n \epsilon^{2} / 32} \sum_{j=1}^{\infty} e^{-V_{C(j)}}=\Delta e^{-n \epsilon^{2} / 32}
\end{aligned}
$$

using the defining expression for $r(j, n)$, where $\Delta=8 \sum_{j=1}^{\infty} e^{-V_{C(j)}}<\infty$, by assumption. Thus, it follows from the Borel-Cantelli lemma that

$$
L\left(\phi_{n}^{*}\right)-\inf _{j \geq 1} \tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right) \rightarrow 0
$$

with probability one as $n \rightarrow \infty$. Now, we can turn to investigating the second term $\inf _{j \geq 1} \tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right)-L^{*}$. By our assumptions, for every $\epsilon>0$, there exists an integer $k$ such that

$$
\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)-L^{*} \leq \epsilon
$$

Fix such a $k$. Thus, it suffices to prove that

$$
\limsup _{n \rightarrow \infty} \inf _{j \geq 1} \tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right)-\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi) \leq 0 \quad \text { with probability one. }
$$

Clearly, for any fixed $k$, if $n$ is large enough, then

$$
r(k, n)=\sqrt{\frac{32}{n} V_{\mathcal{C}^{(k)}} \log (e n)} \leq \frac{\epsilon}{2}
$$

Thus, for such large $n$,

$$
\begin{aligned}
\mathbf{P} & \left\{\inf _{j \geq 1} \tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right)-\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)>\epsilon\right\} \\
& \leq \mathbf{P}\left\{\tilde{L}_{n}\left(\tilde{\phi}_{n, k}\right)-\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)>\epsilon\right\} \\
& =\mathbf{P}\left\{\widehat{L}_{n}\left(\tilde{\phi}_{n, k}\right)+r(k, n)-\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)>\epsilon\right\} \\
& \leq \mathbf{P}\left\{\widehat{L}_{n}\left(\tilde{\phi}_{n, k}\right)-\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)>\frac{\epsilon}{2}\right\} \\
& \leq \mathbf{P}\left\{\sup _{\phi \in \mathcal{C}^{(k)}}\left|\widehat{L}_{n}(\phi)-L(\phi)\right|>\frac{\epsilon}{2}\right\} \\
& \leq 8 n^{V_{C^{(k)}} e^{-n \epsilon^{2} / 128}}
\end{aligned}
$$

by Theorem 12.5 . Therefore, since $V_{\mathcal{C}^{(k)}}<\infty$, the proof is completed.

Theorem 18.2 shows that the method of structural risk minimization is universally consistent under very mild conditions on the sequence of classes $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$

This property, however, is shared with the minimizer of the empirical error over the class $\mathcal{C}^{\left(k_{n}\right)}$, where $k_{n}$ is a properly chosen function of the sample size $n$ (Theorem 18.1). The real strength, then, of structural risk minimization is seen from the next result.

Theorem 18.3. Let $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$ be a sequence of classes of classifiers such that the VC dimensions $V_{\mathcal{C}^{(1)}}, V_{\mathcal{C}^{(2)}}, \ldots$ are all finite. Assume further that the Bayes rule

$$
g^{*} \in \mathcal{C}^{*}=\bigcup_{j=1}^{\infty} \mathcal{C}^{(j)}
$$

that is, a Bayes rule is contained in one of the classes. Let $k$ be the smallest integer such that $g^{*} \in C^{(k)}$. Then for every $n$ and $\epsilon>0$ satisfying

$$
V_{\mathcal{C}^{(k)}} \log (e n) \leq \frac{n \epsilon^{2}}{512}
$$

the error probability of the classification rule based on structural risk minimization $\phi_{n}^{*}$ satisfies

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-L^{*}>\epsilon\right\} \leq \Delta e^{-n \epsilon^{2} / 128}+8 n^{V_{c}(k)} e^{-n \epsilon^{2} / 512}
$$

where $\Delta=\sum_{j=1}^{\infty} e^{-V_{\mathcal{C}}(j)}$.
Proof. Theorem 18.3 follows by examining the proof of Theorem 18.2. The only difference is that by assumption, there is an integer $k$ such that $\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)=L^{*}$. Therefore, for this $k$,

$$
L\left(\phi_{n}^{*}\right)-L^{*}=\left(L\left(\phi_{n}^{*}\right)-\inf _{j \geq 1} \tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right)\right)+\left(\inf _{j \geq 1} \tilde{L}_{n}\left(\tilde{\phi}_{n, j}\right)-\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)\right)
$$

and the two terms on the right-hand side may be bounded as in the proof of Theorem 18.2.

Theorem 18.3 implies that if $g^{*} \in \mathcal{C}^{*}$, there is a universal constant $c_{0}$ and a finite $k$ such that

$$
\mathbf{E} L\left(\phi_{n}^{*}\right)-L^{*} \leq c_{0} \sqrt{\frac{V_{\mathcal{C}^{(k)}} \log n}{n}}
$$

that is, the rate of convergence is always of the order of $\sqrt{\log n / n}$, and the constant factor $V_{\mathcal{C}^{(k)}}$ depends on the distribution. The number $V_{\mathcal{C}^{(k)}}$ may be viewed as the inherent complexity of the Bayes rule for the distribution. The intuition is that the simplest rules are contained in $\mathcal{C}^{(1)}$, and more and more complex rules are added to the class as the index of the class increases. The size of the error is about the same as if we had known $k$ beforehand, and minimized the empirical error over $\mathcal{C}^{(k)}$. In view of the results of Chapter 14 it is clear that the classifier described in Theorem 18.1 does not share this property, since if $L^{*}>0$, then the error probability of the
rule selected from $C^{\left(k_{n}\right)}$ is larger than a constant times $\sqrt{V_{\mathcal{C}^{\left(k_{n}\right)}} \log n / n}$ for some distribution-even if $g^{*} \in \mathcal{C}^{(k)}$ for some fixed $k$. Just as in designs based upon minimum description length, automatic model selection, and other complexity regularization methods (Rissanen (1983), Akaike (1974), Barron (1985; 1991), and Barron and Cover (1991)), structural risk minimization automatically finds where to look for the optimal classifier. The constants appearing in Theorem 18.3 may be improved by using refined versions of the Vapnik-Chervonenkis inequality.
The condition $g^{*} \in \mathcal{C}^{*}$ in Theorem 18.3 is not very severe, as $\mathcal{C}^{*}$ can be a large class with infinite vC dimension. The only requirement is that it should be written as a countable union of classes of finite vc dimension. Note however that the class of all decision rules can not be decomposed as such (see Theorem 18.6). We also emphasize that in order to achieve the $O(\sqrt{\log n / n})$ rate of convergence, we do not have to assume that the distribution is a member of a known finite-dimensional parametric family (see Chapter 16). The condition is imposed solely on the form of the Bayes classifier $g^{*}$.
By inspecting the proof of Theorem 18.2 we see that for every $k$,

$$
\mathbf{E} L\left(\phi_{n}^{*}\right)-L^{*} \leq c_{0} \sqrt{\frac{V_{\mathcal{C}^{(k)}} \log n}{n}}+\left(\inf _{\phi \in \mathcal{C}(k)} L(\phi)-L^{*}\right) .
$$

In fact, Theorem 18.3 is the consequence of this inequality. The first term on the right-hand side, which may be called estimation error, usually increases with growing $k$, while the second, the approximation error, usually decreases with it. Importantly, the above inequality is true for every $k$, so that

$$
\mathbf{E} L\left(\phi_{n}^{*}\right)-L^{*} \leq \inf _{k}\left(c_{0} \sqrt{\frac{V_{\mathcal{C}^{(k)}} \log n}{n}}+\left(\inf _{\phi \in \mathcal{C}(k)} L(\phi)-L^{*}\right)\right) .
$$

Thus structural risk minimization finds a nearly optimal balance between the two terms. See also Problem 18.6.

Remark. It is worthwhile mentioning that under the conditions of Theorem 18.3, an even faster, $O(\sqrt{1 / n})$, rate of convergence is achievable at the expense of magnifying the constant factor. More precisely, it is possible to define the complexity penalties $r(j, n)$ such that the resulting classifier satisfies

$$
\mathbf{E} L\left(\phi_{n}^{*}\right)-L^{*} \leq \frac{c_{1}}{\sqrt{n}},
$$

where the constant $c_{1}$ depends on the distribution. These penalties may be defined by exploiting Alexander's inequality (Theorem 12.10), and the inequality above can be proved by using the bound in Problem 12.10, see Problem 18.6.

Remark. We see from the proof of Theorem 18.2 that

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)>\tilde{L}_{n}\left(\phi_{n}^{*}\right)+\epsilon\right\} \leq \Delta e^{-n \epsilon^{2} / 32} .
$$

This means that $\tilde{L}_{n}\left(\phi_{n}^{*}\right)$ does not underestimate the true error probability of $\phi_{n}^{*}$ by much. This is a very attractive feature of $\tilde{L}_{n}$ as an error estimate, as the designer can be confident about the performance of the rule $\phi_{n}^{*}$.

The initial excitement over a consistent rule with a guaranteed $O(\sqrt{\log n / n})$ rate of convergence to the Bayes error is tempered by a few sobering facts:
(1) The user needs to choose the sequence $\mathcal{C}^{(i)}$ and has to know $V_{\mathcal{C}^{(j)}}$ (see, however, Problems 18.3, 18.4, and the method of simple empirical covering below).
(2) It is difficult to find the structural risk minimizer $\phi_{n}^{*}$ efficiently. After all, the minimization is done over an infinite sequence of infinite sets.

The second concern above deserves more attention. Consider the following simple example: let $d=1$, and let $\mathcal{C}^{(j)}$ be the class of classifiers $\phi$ for which

$$
\{x: \phi(x)=1\}=\bigcup_{i=1}^{j} A_{i}
$$

where each $A_{i}$ is an interval of $\mathcal{R}$. Then, from Theorem 13.7, $V_{\mathcal{C}^{(j)}}=2 j$, and we may take

$$
r(j, n)=\sqrt{\frac{64 j}{n} \log (e n)}
$$

In structural risk minimization, we find those $j$ (possibly empty) intervals that minimize

$$
\widehat{L}_{n}(\phi)+r(j, n)
$$

and call the corresponding classifier $\tilde{\phi}_{n, j}$. For $j=1$ we have $r(j, n)=\sqrt{\frac{64}{n} \log (e n)}$. As $r(n, n)>1+r(1, n)$, and $r(j, n)$ is monotone in $j$, it is easily seen that to pick the best $j$ as well, we need only consider $1 \leq j \leq n$. Still, this is a formidable exercise. For fixed $j$, the best $j$ intervals may be found by considering all possible insertions of $2 j$ interval boundaries among the $n X_{i}$ 's. This brute force method takes computation time bounded from below by $\binom{n+2 j}{2 j}$. If we let $j$ run up to $n$, then we have as a lower bound

$$
\sum_{j=1}^{n}\binom{n+2 j}{2 j}
$$

Just the last term alone, $\binom{3 n}{2 n}$, grows as $\sqrt{\frac{3}{4 \pi n}}\left(\frac{27}{4}\right)^{n}$, and is prohibitively large for $n \geq 20$. Fortunately, in this particular case, there is an algorithm which finds a classifier minimizing $\widehat{L}_{n}(\phi)+r(j, n)$ over $\mathcal{C}^{*}=\bigcup_{j=1}^{\infty} \mathcal{C}^{(j)}$ in computational time polynomial in $n$, see Problem 18.5. Another example when the structural risk minimizer $\phi_{n}^{*}$ is easy to find is described in Problem 18.6. However, such fast algorithms are not always available, and exponential running time prohibits the use of structural risk minimization even for relatively small values of $n$.

### 18.2 Poor Approximation Properties of VC Classes

We pause here for a moment to summarize some interesting by-products that readily follow from Theorem 18.3 and the slow convergence results of Chapter 7. Theorem 18.3 states that for a large class of distributions an $O(\sqrt{\log n / n})$ rate of convergence to the Bayes error $L^{*}$ is achievable. On the other hand, Theorem 7.2 asserts that no universal rates of convergence to $L^{*}$ exist. Therefore, the class of distributions for which Theorem 18.3 is valid cannot be that large, after all. The combination of these facts results in the following three theorems, which say that vC classes-classes of subsets of $\mathcal{R}^{d}$ with finite vC dimension-have necessarily very poor approximation properties. The proofs are left to the reader as easy exercises (see Problems 18.7 to 18.9). For direct proofs, see, for example, Benedek and Itai (1994).

Theorem 18.4. Let $\mathcal{C}$ be any class of classifiers of the form $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$, with vc dimension $V_{\mathcal{C}}<\infty$. Then for every $\epsilon>0$ there exists a distribution such that

$$
\inf _{\phi \in \mathcal{C}} L(\phi)-L^{*}>\frac{1}{2}-\epsilon
$$

Theorem 18.5. Let $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$ be a sequence of classifiers such that the vC dimensions $V_{\mathcal{C}^{(1)}}, V_{\mathcal{C}^{(2)}}, \ldots$ are all finite. Then for any sequence $\left\{a_{k}\right\}$ of positive numbers converging to zero arbitrarily slowly, there exists a distribution such that for every $k$ large enough,

$$
\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)-L^{*}>a_{k}
$$

Theorem 18.6. The class $\mathcal{C}^{*}$ of all (Borel measurable) decision rules of form $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$ cannot be written as

$$
\mathcal{C}^{*}=\bigcup_{j=1}^{\infty} \mathcal{C}^{(j)},
$$

where the VC dimension of each class $\mathcal{C}^{(j)}$ is finite. In other words, the class $\mathcal{B}$ of all Borel subsets of $\mathcal{R}^{d}$ cannot be written as a union of countably many vc classes. In fact, the same is true for the class of all subsets of the set of positive integers.

### 18.3 Simple Empirical Covering

As Theorem 18.2 shows, the method of structural risk minimization provides automatic protection against the danger of overfitting the data, by penalizing complex candidate classifiers. One of the disadvantages of the method is that the penalty terms $r(j, n)$ require knowledge of the vc dimensions of the classes $\mathcal{C}^{(j)}$ or upper bounds of these dimensions. Next we discuss a method proposed by

Buescher and Kumar (1996b) which is applicable even if the vc dimensions of the classes are completely unknown. The method, called simple empirical covering, is closely related to the method based on empirical covering studied in Problem 12.14. In simple empirical covering, we first split the data sequence $D_{n}$ into two parts. The first part is $D_{m}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{m}, Y_{m}\right)\right)$ and the second part is $T_{l}=\left(\left(X_{m+1}, Y_{m+1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$. The positive integers $m$ and $l=n-m$ will be specified later. The first part $D_{m}$ is used to cover $\mathcal{C}^{*}=\bigcup_{j=1}^{\infty} \mathcal{C}^{(j)}$ as follows. For every $\phi \in \mathcal{C}^{*}$, define the binary $m$-vector $b_{m}(\phi)$ by

$$
b_{m}(\phi)=\left(\phi\left(X_{1}\right), \ldots, \phi\left(X_{m}\right)\right) .
$$

There are $N \leq 2^{m}$ different values of $b_{m}(\phi)$. Usually, as $V_{\mathcal{C}^{*}}=\infty, N=2^{m}$, that is, all possible values of $b_{m}(\phi)$ occur as $\phi$ is varied over $\mathcal{C}^{*}$, but of course, $N$ depends on the values of $X_{1}, \ldots, X_{m}$. We call a classifier $\phi$ simpler than another classifier $\phi^{\prime}$, if the smallest index $i$ such that $\phi \in \mathcal{C}^{(i)}$ is smaller than or equal to the smallest index $j$ such that $\phi^{\prime} \in \mathcal{C}^{(j)}$. For every binary $m$-vector $b \in\{0,1\}^{m}$ that can be written as $b=b_{m}(\phi)$ for some $\phi \in \mathcal{C}^{*}$, we pick a candidate classifier $\tilde{\phi}_{m, k}$ with $k \in\{1, \ldots, N\}$ such that $b_{m}\left(\tilde{\phi}_{m, k}\right)=b$, and it is the simplest such classifier, that is, there is no $\phi \in \mathcal{C}^{*}$ such that simultaneously $b_{m}\left(\tilde{\phi}_{m, k}\right)=b(\phi)$ and $\phi$ is simpler than $\tilde{\phi}_{m, k}$. This yields $N \leq 2^{m}$ candidates $\tilde{\phi}_{m, 1}, \ldots, \tilde{\phi}_{m, N}$. Among these, we select one that minimizes the empirical error

$$
\widehat{L}_{l}\left(\tilde{\phi}_{m, k}\right)=\frac{1}{l} \sum_{i=1}^{I} I_{\left\{\tilde{\phi}_{m, k}\left(X_{m+i}\right) \neq Y_{m+i}\right\}}
$$

measured on the independent testing sequence $T_{l}$. Denote the selected classifier by $\phi_{n}^{*}$. The next theorem asserts that the method works under circumstances similar to structural risk minimization.

Theorem 18.7. (BUESCHER AND KUMAR (1996B)). Let $\mathcal{C}^{(1)} \subseteq \mathcal{C}^{(2)} \subseteq \ldots$ be a nested sequence of classes of classifiers such that for any distribution of $(X, Y)$,

$$
\lim _{j \rightarrow \infty} \inf _{\phi \in \mathcal{C}^{(j)}} L(\phi)=L^{*}
$$

Assume also that the vC dimensions $V_{\mathcal{C}^{(1)}}, V_{\mathcal{C}^{(2)}}, \ldots$ are all finite. If $m / \log n \rightarrow \infty$ and $m / n \rightarrow 0$, then the classification rule $\phi_{n}^{*}$ based on simple empirical covering as defined above is strongly universally consistent.

Proof. We decompose the difference between the error probability of the selected rule $\phi_{n}^{*}$ and the Bayes risk as follows:

$$
L\left(\phi_{n}^{*}\right)-L^{*}=\left(L\left(\phi_{n}^{*}\right)-\inf _{1 \leq k \leq N} L\left(\tilde{\phi}_{m, k}\right)\right)+\left(\inf _{1 \leq k \leq N} L\left(\tilde{\phi}_{m, k}\right)-L^{*}\right)
$$

The first term can be handled by Lemma 8.2 and Hoeffding's inequality:

$$
\begin{aligned}
& \mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-\inf _{1 \leq k \leq N} L\left(\tilde{\phi}_{m, k}\right)>\epsilon\right\} \\
& \\
& \quad \leq \mathbf{P}\left\{2 \sup _{1 \leq k \leq N}\left|L\left(\tilde{\phi}_{m, k}\right)-\widehat{L}_{l}\left(\tilde{\phi}_{m, k}\right)\right|>\epsilon\right\} \\
& \\
& =\mathbf{E}\left\{\mathbf{P}\left\{2 \sup _{1 \leq k \leq N}\left|L\left(\tilde{\phi}_{m, k}\right)-\widehat{L}_{l}\left(\tilde{\phi}_{m, k}\right)\right|>\epsilon \mid D_{m}\right\}\right\} \\
& \\
& \leq \mathbf{E}\left\{2 N e^{-l \epsilon^{2} / 2}\right\} \\
& \leq 2^{m+1} e^{-l \epsilon^{2} / 2}=2^{m+1} e^{-(n-m) \epsilon^{2} / 2} .
\end{aligned}
$$

Because $m=o(n)$, the latter expression converges to zero exponentially rapidly. Thus, it remains to show that

$$
\inf _{1 \leq k \leq N} L\left(\tilde{\phi}_{m, k}\right)-L^{*} \rightarrow 0
$$

with probability one. By our assumptions, for every $\epsilon>0$, there is an integer $k$ such that

$$
\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)-L^{*}<\epsilon
$$

Then there exists a classifier $\phi^{(\epsilon)} \in \mathcal{C}^{(k)}$ with $L\left(\phi^{(\epsilon)}\right)-L^{*} \leq \epsilon$. Therefore, we are done if we prove that

$$
\limsup _{n \rightarrow \infty} \inf _{1 \leq k \leq N} L\left(\tilde{\phi}_{m, k}\right)-L\left(\phi^{(\epsilon)}\right) \leq 0 \quad \text { with probability one. }
$$

Clearly, there is a classifier $\tilde{\phi}_{m, j}$ among the candidates $\tilde{\phi}_{m, 1}, \ldots, \tilde{\phi}_{m, N}$, such that

$$
b_{m}\left(\tilde{\phi}_{m, j}\right)=b_{m}\left(\phi^{(\epsilon)}\right)
$$

Since by definition, $\tilde{\phi}_{m, j}$ is simpler than $\phi^{(\epsilon)}$, and the classes $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$ are nested, it follows that $\tilde{\phi}_{m, j} \in \mathcal{C}^{(k)}$. Therefore,

$$
\begin{aligned}
\inf _{1 \leq k \leq N} L\left(\tilde{\phi}_{m, k}\right)-L\left(\phi^{(\epsilon)}\right) & \leq L\left(\tilde{\phi}_{m, j}\right)-L\left(\phi^{(\epsilon)}\right) \\
& \leq \sup _{\phi, \phi^{\prime} \in \mathcal{C}^{(k)}: b_{m}(\phi)=b_{m}\left(\phi^{\prime}\right)}\left|L(\phi)-L\left(\phi^{\prime}\right)\right|
\end{aligned}
$$

where the last supremum is taken over all pairs of classifiers such that their corresponding binary vectors $b_{m}(\phi)$ and $b_{m}\left(\phi^{\prime}\right)$ are equal. But from Problem 12.14,

$$
\mathbf{P}\left\{\sup _{\phi, \phi^{\prime} \in \mathcal{C}^{(k)}: b_{m}(\phi)=b_{m}\left(\phi^{\prime}\right)}\left|L(\phi)-L\left(\phi^{\prime}\right)\right|>\epsilon\right\} \leq 2 \mathcal{S}^{4}\left(\mathcal{C}^{(k)}, 2 m\right) e^{-m \epsilon \log 2 / 4},
$$

which is summable if $m / \log n \rightarrow \infty$. $\square$

REmARK. As in Theorem 18.3, we may assume again that there is an integer $k$ such that $\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)=L^{*}$. Then from the proof of the theorem above we see that the error probability of the classifier $\phi_{n}^{*}$, obtained by the method of simple empirical covering, satisfies

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-L^{*}>\epsilon\right\} \leq 2^{m+1} e^{-(n-m) \epsilon^{2} / 8}+2 \mathcal{S}^{4}\left(\mathcal{C}^{(k)}, 2 m\right) e^{-m \epsilon \log 2 / 8}
$$

Unfortunately, for any choice of $m$, this bound is much larger than the analogous bound obtained for structural risk minimization. In particular, it does not yield an $O(\sqrt{\log n / n})$ rate of convergence. Thus, it appears that the price paid for the advantages of simple empirical covering-no knowledge of the vc dimensions is required, and the implementation may require significantly less computational time in general-is a slower rate of convergence. See Problem 18.10.

## Problems and Exercises

Problem 18.1. Prove Theorem 18.1.
Problem 18.2. Define the complexity penalties $r(j, n)$ so that under the conditions of Theorem 18.3, the classification rule $\phi_{n}^{*}$ based upon structural risk minimization satisfies

$$
\mathbf{E} L\left(\phi_{n}^{*}\right)-L^{*} \leq \frac{c_{1}}{\sqrt{n}},
$$

where the constant $c_{1}$ depends on the distribution. Hint: Use Alexander's bound (Theorem 12.10 ), and the inequality of Problem 12.10.

Problem 18.3. Let $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$ be a sequence of classes of decision rules with finite vC dimensions. Assume that only upper bounds $\alpha_{j} \geq V_{\mathcal{C}^{(j)}}$ on the vC dimensions are known. Define the complexity penalties by

$$
r(j, n)=\sqrt{\frac{32}{n} \alpha_{j} \log (e n)} .
$$

Show that if $\sum_{j=1}^{\infty} e^{-\alpha_{j}}<\infty$, then Theorems 18.2 and 18.3 carry over to the classifier based on structural risk minimization defined by these penalties. This points out that knowledge of relatively rough estimates of the VC dimensions suffice.

Problem 18.4. Let $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \ldots$ be a sequence of classes of classifiers such that the vc dimensions $V_{\mathcal{C}^{(1)}}, V_{\mathcal{C}^{(2)}}, \ldots$ are all finite. Assume furthermore that the Bayes rule is contained in one of the classes and that $L^{*}=0$. Let $\phi_{n}^{*}$ be the rule obtained by structural risk minimization using the positive penalties $r(j, n)$, satisfying
(1) For each $n, r(j, n)$ is strictly monotone increasing in $j$.
(2) For each $j, \lim _{n \rightarrow \infty} r(j, n)=0$.

Show that $\mathbf{E}\left\{L\left(\phi_{n}^{*}\right)\right\}=O(\log n / n)$ (Lugosi and Zeger (1996)). For related work, see Benedek and Itai (1994).

Problem 18.5. Let $\mathcal{C}^{(j)}$ be the class of classifiers $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}$ satisfying

$$
\{x: \phi(x)=1\}=\bigcup_{i=1}^{j} A_{i}
$$

where the $A_{i}$ 's are bounded intervals in $\mathcal{R}$. The purpose of this exercise is to point out that there is a fast algorithm to find the structural risk minimizer $\phi_{n}^{*} \operatorname{over} \mathcal{C}^{*}=\bigcup_{j=1}^{\infty} \mathcal{C}^{(j)}$, that is, which minimizes $\tilde{L}_{n}=\widehat{L}_{n}+r(j, n)$ over $\mathcal{C}^{*}$, where the penalties $r(j, n)$ are defined as in Theorems 18.2 and 18.3. The property below was pointed out to us by Miklós Csürös and Réka Szabó.
(1) Let $A_{1, j}^{*}, \ldots, A_{j, j}^{*}$ be the $j$ intervals defining the classifier $\tilde{\phi}_{n, j}$ minimizing $\widehat{L}_{n}$ over $\mathcal{C}^{(j)}$. Show that the optimal intervals for $C^{(j+1)}, A_{1, j+1}^{*}, \ldots, A_{j+1, j+1}^{*}$ satisfy the following property: either $j$ of the intervals coincide with $A_{1, j}^{*}, \ldots, A_{j, j}^{*}$, or $j-1$ of them are among $A_{1, j}^{*}, \ldots, A_{j, j}^{*}$, and the remaining two intervals are subsets of the $j$-th interval.
(2) Use property (1) to define an algorithm that finds $\phi_{n}^{*}$ in running time polynomial in the sample size $n$.

Problem 18.6. Assume that the distribution of $X$ is concentrated on the unit cube, that is, $\mathbf{P}\left\{X \in[0,1]^{d}\right\}=1$. Let $\mathcal{P}_{j}$ be a partition of $[0,1]^{d}$ into cubes of size $1 / j$, that is, $\mathcal{P}_{j}$ contains $j^{d}$ cubic cells. Let $\mathcal{C}^{(j)}$ be the class of all histogram classifiers $\phi:[0,1]^{d} \rightarrow\{0,1\}$ based on $\mathcal{P}_{j}$. In other words, $\mathcal{P}_{j}$ contains all $2^{j^{d}}$ classifiers which assign the same label to points falling in the same cell of $\mathcal{P}_{j}$. What is the vc dimension $V_{\mathcal{C}^{(j)}}$ of $\mathcal{C}^{(j)}$ ? Point out that the classifier minimizing $\widehat{L}_{n}$ over $\mathcal{C}^{(j)}$ is just the regular histogram rule based on $\mathcal{P}_{j}$. (See Chapter 9.) Thus, we have another example in which the empirically optimal classifier is computationally inexpensive. The structural risk minimizer $\phi_{n}^{*}$ based on $\mathcal{C}^{*}=\bigcup_{j=1}^{\infty} \mathcal{C}^{(j)}$ is also easy to find. Assume that the a posteriori probability $\eta(x)$ is uniformly Lipschitz, that is, for any $x, y \in[0,1]^{d}$,

$$
|\eta(x)-\eta(y)| \leq c\|x-y\|
$$

where $c$ is some constant. Find upper bounds for the rate of convergence of $\mathbf{E} L\left(\phi_{n}^{*}\right)$ to $L^{*}$.
Problem 18.7. Prove Theorem 18.4. Hint: Use Theorem 7.1.
Problem 18.8. Prove Theorem 18.5. Hint: Use Theorem 7.2.
Problem 18.9. Prove Theorem 18.6. Hint: Use Theorems 7.2 and 18.3.
Problem 18.10. Assume that the Bayes rule $g^{*}$ is contained in $\mathcal{C}^{*}=\bigcup_{j=1}^{\infty} \mathcal{C}^{(j)}$. Let $\phi_{n}^{*}$ be the classifier obtained by simple empirical covering. Determine the value of the design parameter $m$ that minimizes the bounds obtained in the proof of Theorem 18.7. Obtain a tight upper bound for $\mathbf{E} L\left(\phi_{n}^{*}\right)-L^{*}$. Compare your results with Theorem 18.3.

## 19

## Condensed and Edited Nearest Neighbor Rules

### 19.1 Condensed Nearest Neighbor Rules

Condensing is the process by which we eliminate data points, yet keep the same behavior. For example, in the nearest neighbor rule, by condensing we might mean the reduction of $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ to $\left(X_{1}^{\prime}, Y_{1}^{\prime}\right), \ldots,\left(X_{m}^{\prime}, Y_{m}^{\prime}\right)$ such that for all $x \in \mathcal{R}^{d}$, the $1-\mathrm{NN}$ rule is identical based on the two samples. This will be called pure condensing. This operation has no effect on $L_{n}$, and therefore is recommended whenever space is at a premium. The space savings should be substantial whenever the classes are separated. Unfortunately, pure condensing is computationally expensive, and offers no hope of improving upon the performance of the ordinary 1 -NN rule.


Figure 19.1. Pure condensing: Eliminating the marked points does not change the decision.

Hart (1968) has the first simple algorithm for condensing. He picks a subset that correctly classifies the remaining data by the $1-$ nN rule. Finding a minimal such subset is computationally difficult, but heuristics may do the job. Hart's heuristic is also discussed in Devijver and Kittler (1982, p.120). For probabilistic analysis, we take a more abstract setting. Let $\left(X_{1}^{\prime}, Y_{1}^{\prime}\right), \ldots,\left(X_{m}^{\prime}, Y_{m}^{\prime}\right)$ be a sequence that depends in an arbitrary fashion on the data $D_{n}$, and let $g_{n}$ be the 1-nearest neighbor rule with $\left(X_{1}^{\prime}, Y_{1}^{\prime}\right), \ldots,\left(X_{m}^{\prime}, Y_{m}^{\prime}\right)$, where for simplicity, $m$ is fixed beforehand. The data might, for example, be obtained by finding the subset of the data of size $m$ for which the error with the $1-\mathrm{NN}$ rule committed on the remaining $n-m$ data is minimal (this will be called Hart's rule). Regardless, if $\widehat{L}_{n}=(1 / n) \sum_{i=1}^{n} I_{\left\{g_{n}\left(X_{i}\right) \neq Y_{i}\right\}}$ and $L_{n}=\mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}$, then we have the following:

Theorem 19.1. (Devroye and Wagner (1979c)). For all $\epsilon>0$ and all distributions,

$$
\mathbf{P}\left\{\left|L_{n}-\widehat{L}_{n}\right| \geq \epsilon\right\} \leq 8\left(\frac{n e}{d+1}\right)^{(d+1) m(m-1)} e^{-n \epsilon^{2} / 32}
$$

REMARK. The estimate $\widehat{L}_{n}$ is called the resubstitution estimate of the error probability. It is thoroughly studied in Chapter 23, where several results of the aforementioned kind are stated.

Proof. Observe that

$$
\widehat{L}_{n}=\frac{1}{n} \sum_{j=1}^{n} I_{\left\{\left(X_{j}, Y_{j}\right) \notin \bigcup_{i=1}^{m} B_{i} \times\left\{Y_{i}^{\prime}\right\}\right\}},
$$

where $B_{i}$ is the Voronoi cell of $X_{i}^{\prime}$ in the Voronoi partition corresponding to $X_{1}^{\prime}, \ldots, X_{m}^{\prime}$, that is, $B_{i}$ is the set of points of $\mathcal{R}^{d}$ closer to $X_{i}^{\prime}$ than to any other $X_{J}^{\prime}$ (with appropriate distance-tie breaking). Similarly,

$$
L_{n}=\mathbf{P}\left\{(X, Y) \notin \bigcup_{i=1}^{m} B_{i} \times\left\{Y_{i}^{\prime}\right\} \mid D_{n}\right\}
$$

We use the simple upper bound

$$
\left|L_{n}-\widehat{L}_{n}\right| \leq \sup _{A \in \mathcal{A}_{m}}\left|v_{n}(A)-v(A)\right|,
$$

where $v$ denotes the measure of $(X, Y), v_{n}$ is the corresponding empirical measure, and $\mathcal{A}_{m}$ is the family of all subsets of $\mathcal{R}^{d} \times\{0,1\}$ of the form $\bigcup_{i=1}^{m} B_{i} \times\left\{y_{i}\right\}$, where $B_{1}, \ldots, B_{m}$ are Voronoi cells corresponding to $x_{1}, \ldots, x_{m}, x_{i} \in \mathcal{R}^{d}, y_{i} \in\{0,1\}$. We use the Vapnik-Chervonenkis inequality to bound the above supremum. By Theorem 13.5 (iv),

$$
s\left(\mathcal{A}_{m}, n\right) \leq s(\mathcal{A}, n)^{m}
$$

where $\mathcal{A}$ is the class of sets $B_{1} \times\left\{y_{1}\right\}$. But each set in $\mathcal{A}$ is an intersection of at most $m-1$ hyperplanes. Therefore, by Theorems 13.5 (iii), 13.9, and 13.3,

$$
s(\mathcal{A}, n) \leq \sup _{n_{0}, n_{1}: n_{0}+n_{1}=n}\left(\prod_{j=0}^{1}\left(\frac{n_{j} e^{\prime}}{d+1}\right)^{(d+1)(k-1)}\right) \leq\left(\frac{n e}{d+1}\right)^{(d+1)(k-1)}
$$

where $n_{j}$ denotes the number of points in $\mathcal{R}^{d} \times\{j\}$. The result now follows from Theorem 12.5.

Remark. With Hart's rule, at least $m$ data points are correctly classified by the 1 -NN rule (if we handle distance ties satisfactorily). Therefore, $\widehat{L}_{n} \leq 1-m / n$.

The following is a special case: Let $m<n$ be fixed and let $D_{m}$ be an arbitrary (possibly random) subset of $m$ pairs from $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, used in $g_{n}$. Let the remaining $n-m$ points be denoted by $T_{n-m}$. We write $L_{n}\left(D_{m}\right)$ for the probability of error with the 1-NN based upon $D_{m}$, and we define

$$
\widehat{L}_{n, m}\left(D_{m}, T_{n-m}\right)=\frac{1}{n-m} \sum_{\left(X_{i}, Y_{i}\right) \in T_{n-m}} I_{\left\{g_{n}\left(X_{i}\right) \neq Y_{i}\right\}}
$$

In Hart's rule, $\widehat{L}_{n, m}$ would be zero, for example. Then we have:
Theorem 19.2. For all $\epsilon>0$,

$$
\begin{aligned}
\mathbf{P}\left\{\left|L_{n}-\widehat{L}_{n, m}\right|>\epsilon\right\} & \leq 2\binom{n}{m} e^{-2(n-m) \epsilon^{2}} \\
& \leq 2 n^{m} e^{-2(n-m) \epsilon^{2}}
\end{aligned}
$$

where $L_{n}$ is the probability of error with $g_{n}$ (note that $D_{m}$ depends in an arbitrary fashion upon $D_{n}$ ), and $\widehat{L}_{n, m}$ is $\widehat{L}_{n, m}\left(D_{m}, T_{n-m}\right)$ with the data set $D_{m}$.

Proof. List the $m$-element subsets $\left\{i_{1}, \ldots, i_{m}\right\}$ of $\{1,2, \ldots, n\}$, and define $D_{m}^{(i)}$ as the sequence of $m$ pairs from $D_{n}$ indexed by $i=\left\{i_{1}, \ldots, i_{m}\right\}, 1 \leq i \leq\binom{ n}{m}$. Accordingly, denote $T_{n-m}^{(i)}=D_{n}-D_{m}^{(i)}$. Then

$$
\begin{aligned}
\mathbf{P} & \left\{\left|L_{n}-\widehat{L}_{n, m}\right|>\epsilon\right\} \\
& \leq \sum_{i=1}^{\binom{n}{m}} \mathbf{P}\left\{\left|L_{n}\left(D_{m}^{(i)}\right)-\widehat{L}_{n, m}\left(D_{m}^{(i)}, T_{n-m}^{(i)}\right)\right|>\epsilon\right\} \\
& =\mathbf{E}\left\{\sum_{i=1}^{\binom{n}{m}} \mathbf{P}\left\{\left|L_{n}\left(D_{m}^{(i)}\right)-\widehat{L}_{n, m}\left(D_{m}^{(i)}, T_{n-m}^{(i)}\right)\right|>\epsilon \mid D_{m}^{(i)}\right\}\right\}
\end{aligned}
$$

$\leq \mathbf{E}\left\{\sum_{i=1}^{\substack{n \\ m \\ m}} 2 e^{-2(n-m) \epsilon^{2}}\right\}$
(by Hoeffding's inequality, because given $D_{m}^{(i)}$,
$(n-m) \widehat{L}_{n, m}\left(D_{n}^{(i)}, T_{n-m}^{(i)}\right)$ is binomial $\left.\left((n-m), L_{n}\left(D_{m}^{(i)}\right)\right)\right)$
$=2\binom{n}{m} e^{-2(n-m) \epsilon^{2}}$.

By checking the proof, we also see that if $D_{m}$ is selected to minimize the error estimate $\widehat{L}_{n, m}\left(D_{m}, T_{n-m}\right)$, then the error probability $L_{n}$ of the obtained rule satisfies

$$
\begin{align*}
& \mathbf{P}\left\{\left|L_{n}-\inf _{D_{m}} L_{n}\left(D_{m}\right)\right|>\epsilon\right\} \\
& \\
& \leq \mathbf{P}\left\{2 \sup _{D_{m}}\left|L_{n}\left(D_{m}\right)-\widehat{L}_{n, m}\left(D_{m}, T_{n-m}\right)\right|>\epsilon\right\} \quad \text { (by Theorem 8.4) }  \tag{19.1}\\
& \quad \leq 2\binom{n}{m} e^{-(n-m) \epsilon^{2} / 2} \quad \text { (as in the proof of Theorem 19.2). }
\end{align*}
$$

Thus, for the particular rule that mimics Hart's rule (with the exception that $m$ is fixed), if $m$ is not too large--it must be much smaller than $n / \log n-L_{n}$ is likely to be close to the best possible we can hope for with a $1-\mathrm{NN}$ rule based upon a subsample of size $m$. With some work (see Problem 12.1), we see that

$$
\mathbf{E}\left\{\left|L_{n}-\inf _{D_{m}} L_{n}\left(D_{m}\right)\right|\right\}=O\left(\sqrt{\frac{m \log m}{n}}\right)
$$

By Theorem 5.1,

$$
\mathbf{E}\left\{\inf _{D_{m}} L_{n}\left(D_{m}\right)\right\} \leq \mathbf{E}\left\{L_{m}\right\} \rightarrow L_{\mathrm{NN}} \quad \text { as } \quad m \rightarrow \infty
$$

where $L_{m}$ is the probability of error with the 1-NN rule based upon a sample of $m$ data pairs. Hence, if $m=o(n / \log n), m \rightarrow \infty$,

$$
\limsup _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\} \leq L_{N N}
$$

for the 1-NN rule based on $m$ data pairs $D_{m}$ selected to minimize the error estimate $\widehat{L}_{n, m}\left(D_{m}, T_{n-m}\right)$. However, this is very pessimistic indeed. It reassures us that with only a small fraction of the original data, we obtain at least as good a performance as with the full data set-so, this method of condensing is worthwhile. This is not very surprising. Interestingly, however, the following much stronger result is true.

Theorem 19.3. If $m=o(n / \log n)$ and $m \rightarrow \infty$, and if $L_{n}$ is the probability of error for the condensed nearest neighbor rule in which $\widehat{L}_{n, m}\left(D_{m}, T_{n-m}\right)$ is minimized over all data sets $D_{m}$, then

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}=L^{*}
$$

Proof. By (19.1), it suffices to establish that

$$
\mathbf{E}\left\{\inf _{D_{m}} L_{n}\left(D_{m}\right)\right\} \rightarrow L^{*}
$$

as $m \rightarrow \infty$ such that $m=o(n)$, where $L_{n}\left(D_{m}\right)$ is the probability of error of the 1-NN rule with $D_{m}$. As this is one of the fundamental properties of nearest neighbor rules not previously found in texts, we offer a thorough analysis and proof of this result in the remainder of this section, culminating in Theorem 19.4.

The distribution-free result of Theorem 19.3 sets the stage for many consistency proofs for rules that use condensing (or editing, or prototyping, as defined in the next two sections). It states that inherently, partitions of the space by $1-\mathrm{NN}$ rules are rich.

Historical remark. Other condensed nearest neighbor rules are presented by Gates (1972), Ullmann (1974), Ritter, Woodruff, Lowry, and Isenhour (1975), Tomek (1976b), Swonger (1972), Gowda and Krishna (1979), and Fukunaga and Mantock (1984).

Define $Z=I_{\{\eta(X)>1 / 2\}}$. Let $\left(X_{i}^{\prime}, Y_{i}, Z_{i}\right), i=1,2, \ldots, n$, be i.i.d. tuples, independent of $(X, Y, Z)$, where $X_{i}^{\prime}$ may have a distribution different from $X$, but the support set of the distribution $\mu^{\prime}$ of $X_{i}^{\prime}$ is identical to that of $\mu$, the distribution of $X$. Furthermore, $\mathbf{P}\left\{Y_{i}=1 \mid X_{i}^{\prime}=x\right\}=\eta(x)$, and $Z_{i}=I_{\left\{\eta\left(X_{i}^{\prime}\right)>1 / 2\right\}}$ is the Bayes decision at $X_{i}^{\prime}$.

Lemma 19.1. Let $L_{n}=\mathbf{P}\left\{Z_{(1)}^{\prime}(X) \neq Z \mid X_{1}^{\prime}, Z_{1}, \ldots, X_{n}^{\prime}, Z_{n}\right\}$ be the probability of error for the 1 -NN rule based on $\left(X_{i}^{\prime}, Z_{i}\right), 1 \leq i \leq n$, that is, $Z_{(1)}^{\prime}(X)=Z_{i}$ if $X_{i}^{\prime}$ is the nearest neighbor of $X$ among $X_{1}^{\prime}, \ldots, X_{n}^{\prime}$. Then

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}=0
$$

Proof. Denote by $X_{(1)}^{\prime}(X)$ the nearest neighbor of $X$ among $X_{1}^{\prime}, \ldots, X_{n}^{\prime}$. Notice that the proof of Lemma 5.1 may be extended in a straightforward way to show that $\left\|X_{(1)}^{\prime}(X)-X\right\| \rightarrow 0$ with probability one. Since this is the only property of the nearest neighbor of $X$ that we used in deriving the asymptotic formula for the ordinary 1-NN error, $\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}$ equals $L_{\mathrm{NN}}$ corresponding to the pair ( $X, Z$ ). But we have $\mathbf{P}\{Z=1 \mid X=x\}=I_{\{\eta(x)>1 / 2\}}$. Thus, the Bayes probability of error $L^{*}$
for the pattern recognition problem with $(X, Z)$ is zero. Hence, for this distribution, the $1-\mathrm{NN}$ rule is consistent as $L_{\mathrm{NN}}=L^{*}=0$.

Lemma 19.2. Let $Z_{(1)}^{\prime}(X)$ be as in the previous lemma. Let

$$
L_{n}=\mathbf{P}\left\{Z_{(1)}^{\prime}(X) \neq Y \mid X_{1}^{\prime}, Y_{1}, \ldots, X_{n}^{\prime}, Y_{n}\right\}
$$

be the probability of error for the discrimination problem for $(X, Y)$ (not $(X, Z)$ ). Then

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}=L^{*}
$$

where $L^{*}$ is the Bayes error corresponding to $(X, Y)$.

Proof.

$$
\begin{aligned}
\mathbf{E}\left\{L_{n}\right\} & =\mathbf{P}\left\{Z_{(1)}^{\prime}(X) \neq Y\right\} \\
& \leq \mathbf{P}\left\{Z_{(1)}^{\prime}(X) \neq Z\right\}+\mathbf{P}\{Y \neq Z\} \\
& =o(1)+L^{*}
\end{aligned}
$$

by Lemma 19.1.

Theorem 19.4. Let $D_{m}$ be a subset of size $m$ drawn from $D_{n}$. If $m \rightarrow \infty$ and $m / n \rightarrow 0$ as $n \rightarrow \infty$, then

$$
\lim _{n \rightarrow \infty} \mathbf{P}\left\{\inf _{D_{m}} L_{n}\left(D_{m}\right)>L^{*}+\epsilon\right\}=0
$$

for all $\epsilon>0$, where $L_{n}\left(D_{m}\right)$ denotes the conditional probability of error of the nearest neighbor rule with $D_{m}$, and the infimum ranges over all $\binom{n}{m}$ subsets.

Proof. Let $D$ be the subset of $D_{n}$ consisting of those pairs ( $X_{i}, Y_{i}$ ) for which $Y_{i}=I_{\left\{\eta\left(X_{i}\right)>1 / 2\right\}}=Z_{i}$. If $|D| \geq m$, let $D^{*}$ be the first $m$ pairs of $D$, and if $|D|<m$, let $D^{*}=\left\{\left(X_{1}, Y_{1}\right), \ldots,\left(X_{m}, Y_{m}\right)\right\}$. Then

$$
\inf _{D_{m}} L_{n}\left(D_{m}\right) \leq L_{n}\left(D^{*}\right)
$$

If $N=|D| \geq m$, then we know that the pairs in $D^{*}$ are i.i.d. and drawn from the distribution of ( $X^{\prime}, Z$ ), where $X^{\prime}$ has the same support set as $X$; see Problem 19.2 for properties of $X^{\prime}$. In particular,

$$
\begin{aligned}
& \mathbf{P}\left\{\inf _{D_{m}} L_{n}\left(D_{m}\right)>L^{*}+\epsilon\right\} \\
& \quad \leq \mathbf{P}\{N<m\}+\mathbf{P}\left\{N \geq m, L_{n}\left(D^{*}\right)>L^{*}+\epsilon\right\}
\end{aligned}
$$

$$
\begin{aligned}
\leq & \mathbf{P}\{N<m\}+\mathbf{P}\left\{L_{n}\left(D^{*}\right)>L^{*}+\epsilon\right\} \\
\leq & \mathbf{P}\{\operatorname{Binomial}(n, p)<m\}+\frac{\mathbf{E}\left\{L_{n}\left(D^{*}\right)\right\}-L^{*}}{\epsilon} \\
& \left(\text { where } p=\mathbf{P}\{Y=Z\}=1-L^{*} \geq 1 / 2,\right. \text { and by Markov's inequality) } \\
= & o(1)
\end{aligned}
$$

by the law of large numbers (here we use $m / n \rightarrow 0$ ), and by Lemma 19.2 (here we use $m \rightarrow \infty$ ).

### 19.2 Edited Nearest Neighbor Rules

Edited nearest neighbor rules are 1-nN rules that are based upon carefully selected subsets $\left(X_{1}^{\prime}, Y_{1}^{\prime}\right), \ldots,\left(X_{m}^{\prime}, Y_{m}^{\prime}\right)$. This situation is partially dealt with in the previous section, as the frontier between condensed and edited nearest neighbor rules is illdefined. The idea of editing based upon the $k$-nN rule was first suggested by Wilson (1972) and later studied by Wagner (1973) and Penrod and Wagner (1977). Wilson suggests the following scheme: compute $\left(X_{i}, Y_{i}, Z_{i}\right)$, where $Z_{i}$ is the $k$-NN decision at $X_{i}$ based on the full data set with $\left(X_{i}, Y_{i}\right)$ deleted. Then eliminate all data pairs for which $Y_{i} \neq Z_{i}$. The remaining data pairs are used with the 1 -nN rule (not the $k$-NN rule). Another rule, based upon data splitting is dealt with by Devijver and Kittler (1982). A survey is given by Dasarathy (1991), Devijver (1980), and Devijver and Kittler (1980). Repeated editing was investigated by Tomek (1976a). Devijver and Kittler (1982) propose a modification of Wilson's leave-one-out method of editing based upon data splitting.

### 19.3 Sieves and Prototypes

Let $g_{n}$ be a rule that uses the 1-NN classification based upon prototype data pairs $\left(X_{1}^{\prime}, Y_{1}^{\prime}\right), \ldots,\left(X_{m}^{\prime}, Y_{m}^{\prime}\right)$ that depend in some fashion on the original data. If the pairs form a subset of the data pairs (and thus, $m \leq n$ ), we have edited or condensed nearest neighbor rules. However, the ( $X_{i}^{\prime}, Y_{i}^{\prime}$ ) pairs may be strategically picked outside the original data set. For example, in relabeling (see Section 11.7), $m=n$, $X_{i}^{\prime}=X_{i}$ and $Y_{i}^{\prime}=g_{n}^{\prime}\left(X_{i}\right)$, where $g_{n}^{\prime}\left(X_{i}\right)$ is the $k$-nN decision at $X_{i}$. Under some conditions, the relabeling rule is consistent (see Theorem 11.2). The true objective of prototyping is to extract information from the data by insisting that $m$ be much smaller than $n$.


Chang (1974) describes a rule in which we iterate the following step until a given stopping rule is satisfied: merge the two closest nearest neighbors of the same class and replace both pairs by a new average prototype pair. Kohonen (1988; 1990) recognizes the advantages of such prototyping in general as a device for partitioning $\mathcal{R}^{d}$-he calls this learning vector quantization. This theme was picked up again by Geva and Sitte (1991), who pick $X_{1}^{\prime}, \ldots, X_{m}^{\prime}$ as a random subset of $X_{1}, \ldots, X_{n}$ and allow $Y_{i}^{\prime}$ to be different from $Y_{i}$. Diverging a bit from Geva and Sitte, we might minimize the empirical error with the prototyped 1-NN rule over all $Y_{1}^{\prime}, \ldots, Y_{m}^{\prime}$, where the empirical error is that committed on the remaining data. We show that this simple strategy leads to a Bayes-risk consistent rule whenever $m \rightarrow \infty$ and $m / n \rightarrow 0$. Note, in particular, that we may take $\left(X_{1}^{\prime}, \ldots, X_{m}^{\prime}\right)=\left(X_{1}, \ldots, X_{m}\right)$, and that we "throw away" $Y_{1}, \ldots, Y_{m}$, as these are not used. We may, in fact, use additional data with missing $Y_{i}$-values for this purpose-the unclassified data are thus efficiently used to partition the space.

Let

$$
\widehat{L}_{n}\left(g_{n}\right)=\frac{1}{n-m} \sum_{i=m+1}^{n} I_{\left\{g_{n}\left(X_{i}\right) \neq Y_{i}\right\}}
$$

be the empirical risk on the remaining data, where $g_{n}$ is the 1-NN rule based upon $\left(X_{1}, Y_{1}^{\prime}\right), \ldots,\left(X_{m}, Y_{m}^{\prime}\right)$. Let $g_{n}^{*}$ be the 1-NN rule with the choice of $Y_{1}^{\prime}, \ldots, Y_{m}^{\prime}$ that minimizes $\widehat{L}_{n}\left(g_{n}\right)$. Let $L\left(g_{n}^{*}\right)$ denote its probability of error.

Theorem 19.5. $L\left(g_{n}^{*}\right) \rightarrow L^{*}$ in probability for all distributions whenever $m \rightarrow$ $\infty$ and $m / n \rightarrow 0$.

Proof. There are $2^{m}$ different possible functions $g_{n}$. Thus,

$$
\begin{aligned}
& \mathbf{P}\left\{\sup _{g_{n}}\left\{\widehat{L}_{n}\left(g_{n}\right)-L\left(g_{n}\right)|>\epsilon| X_{1}, \ldots, X_{n}\right\}\right. \\
& \\
& \leq 2^{m} \sup _{g_{n}} \mathbf{P}\left\{\left|\widehat{L}_{n}\left(g_{n}\right)-L\left(g_{n}\right)\right|>\epsilon \mid X_{1}, \ldots, X_{m}\right\} \\
& \\
& \leq 2^{m+1} e^{-2(n-m) \epsilon^{2}}
\end{aligned}
$$

by Hoeffding's inequality. Also,

$$
\begin{aligned}
& \mathbf{P}\left\{L\left(g_{n}^{*}\right)>L^{*}+3 \epsilon\right\} \\
& \leq \mathbf{P}\left\{L\left(g_{n}^{*}\right)-\widehat{L}_{n}\left(g_{n}^{*}\right)>\epsilon\right\} \\
&+\mathbf{P}\left\{\widehat{L}_{n}\left(\tilde{g}_{n}\right)-L\left(\tilde{g}_{n}\right)>\epsilon\right\} \quad\left(\tilde{g}_{n} \text { minimizes } L\left(g_{n}\right)\right) \\
&\left.+\mathbf{P}\left\{L\left(\tilde{g}_{n}\right)>L^{*}+\epsilon\right\} \quad \text { (here we used } \widehat{L}_{n}\left(\tilde{g}_{n}\right) \geq \widehat{L}_{n}\left(g_{n}^{*}\right)\right) \\
& \leq 2 \mathbf{E}\left\{\mathbf{P}\left\{\sup _{g_{n}}\left|\widehat{L}_{n}\left(g_{n}\right)-L\left(g_{n}\right)\right|>\epsilon \mid X_{1}, \ldots, X_{n}\right\}\right\} \\
&+\mathbf{P}\left\{L\left(g_{n}^{+}\right)>L^{*}+\epsilon\right\} \\
&\left(\text { where } g_{n}^{+} \text {is the } 1-\text { NN rule based on }\left(X_{1}, Z_{1}\right), \ldots,\left(X_{m}, Z_{m}\right)\right. \\
&\text { with } \left.Z_{i}=I_{\left\{n\left(X_{i}\right)>1 / 2\right\}} \text { as in Lemma } 19.1\right) \\
& \leq 2^{m+2} e^{-(n-m) \epsilon^{2}}+o(1) \\
&(\text { if } m \rightarrow \infty, \text { by Lemma 19.1) } \\
&= o(1)
\end{aligned}
$$

if $m \rightarrow \infty$ and $m / n \rightarrow 0 . \square$

If we let $X_{1}^{\prime}, \ldots, X_{m}^{\prime}$ have arbitrary values-not only among those taken by $X_{1}, \ldots, X_{n}$-then we get a much larger, more flexible class of classifiers. For example, every linear discriminant is nothing but a prototype $1-$ nN rule with $m=2$ just take $\left(X_{1}^{\prime}, 1\right),\left(X_{2}^{\prime}, 0\right)$ and place $X_{1}^{\prime}$ and $X_{2}^{\prime}$ in the right places. In this sense, prototype 1 -nN rules generalize a vast class of rules. The most promising strategy of choosing prototypes is to minimize the empirical error committed on the training sequence $D_{n}$. Finding this optimum may be computationally very expensive. Nevertheless, the theoretical properties provided in the next result may provide useful guidance.

Theorem 19.6. Let $\mathcal{C}$ be the class of nearest neighbor rules based on prototype pairs $\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right), m \geq 3$, where the $\left(x_{i}, y_{i}\right)$ 's range through $\mathcal{R}^{d} \times\{0,1\}$. Given the training data $D_{n}=\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, let $g_{n}$ be the nearest neighbor rule from $\mathcal{C}$ minimizing the error estimate

$$
\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi\left(X_{i}\right) \neq Y_{i}\right\}}, \quad \phi \in \mathcal{C} .
$$

Then for each $\epsilon>0$,

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq 8 \cdot 2^{m} n^{(d+1) m(m-1) / 2} e^{-n \epsilon^{2} / 128}
$$

The rule is consistent if $m \rightarrow \infty$ such that $m^{2} \log m=o(n)$. For $d=1$ and $d=2$
the probability bound may be improved significantly. For $d=1$,

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq 8 \cdot(2 n)^{m} e^{-n \epsilon^{2} / 128}
$$

and for $d=2$,

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq 8 \cdot 2^{m} n^{9 m-18} e^{-n \epsilon^{2} / 128}
$$

In both cases, the rule is consistent if $m \rightarrow \infty$ and $m \log m=o(n)$.

Proof. It follows from Theorem 12.6 that

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)>\epsilon\right\} \leq 8 \mathcal{S}(\mathcal{C}, n) e^{-n \epsilon^{2} / 128}
$$

where $\mathcal{S}(\mathcal{C}, n)$ is the $n$-th shatter coefficient of the class of sets $\{x: \phi(x)=1\}$, $\phi \in \mathcal{C}$. All we need is to find suitable upper bounds for $\mathcal{S}(\mathcal{C}, n)$. Each classifier $\phi$ is a partitioning rule based on the $m$ Voronoi cells defined by $x_{1}, \ldots, x_{m}$. Therefore, $\mathcal{S}(\mathcal{C}, n)$ is not more than $2^{m}$ times the number of different ways $n$ points in $\mathcal{R}^{d}$ can be partitioned by Voronoi partitions defined by $m$ points. In each partition, there are at most $m(m-1) / 2$ cell boundaries that are subsets of $d-1$-dimensional hyperplanes. Thus, the sought number is not greater than the number of different ways $m(m-1) / 2$ hyperplanes can partition $n$ points. By results of Chapter 13, this is at most $n^{(d+1) m(m-1) / 2}$, proving the first inequality.

The other two inequalities follow by sharper bounds on the number of cell boundaries. For $d=1$, this is clearly at most $m$. To prove the third inequality, for each Voronoi partition construct a graph whose vertices are $x_{1}, \ldots, x_{m}$, and two vertices are connected with an edge if and only if their corresponding Voronoi cells are connected. It is easy to see that this graph is planar. But the number of edges of a planar graph with $m$ vertices cannot exceed $3 m-6$ (see Nishizeki and Chiba (1988, p.10)) which proves the inequality.

The consistency results follow from the stated inequalities and from the fact that $\inf _{\phi \in \mathcal{C}} L(\phi)$ tends to $L^{*}$ as $m \rightarrow \infty$ (check the proof of Theorem 5.15 again).

## Problems and Exercises

Problem 19.1. Let $N \leq n$ be the size of the data set after pure condensing. Show that $\lim _{\inf _{n \rightarrow \infty}} \mathbf{E}\{N\} / n>0$ whenever $L^{*}>0$. True or false: if $L^{*}=0$, then $\mathbf{E}\{N\}=o(n)$. Hint: Consider the real line, and note that all points whose right and left neighbors are of the same class are eliminated.

Problem 19.2. Let $\left(X_{1}, Y_{1}\right),\left(X_{2}, Y_{2}\right), \ldots$ be an i.i.d. sequence of pairs of random variables in $\mathcal{R}^{d} \times\{0,1\}$ with $\mathbf{P}\left\{Y_{1}=1 \mid X_{1}=x\right\}=\eta(x)$. Let $\left(X^{\prime}, Z\right)$ be the first pair $\left(X_{i}, Y_{i}\right)$ in
the sequence such that $Y_{i}=I_{\left\{\eta\left(X_{i}\right)>1 / 2\right\}}$. Show that the distribution $\mu^{\prime}$ of $X^{\prime}$ is absolutely continuous with respect to the common distribution $\mu$ of the $X_{i}$ 's, with density (i.e., RadonNikodym derivative)

$$
\frac{d \mu^{\prime}}{d \mu}(x)=\frac{1-\min (\eta(x), 1-\eta(x))}{1-L^{*}}
$$

where $L^{*}$ is the Bayes error corresponding to $\left(X_{1}, Y_{1}\right)$. Let $Y$ be a $\{0,1\}$-valued random variable with $\mathbf{P}\left\{Y=1 \mid X^{\prime}=x\right\}=\eta(x)$. If $L^{\prime}$ denotes the Bayes error corresponding to ( $X^{\prime}, Y$ ), then show that $L^{\prime} \leq L^{*}$.
Problem 19.3. Consider the following edited nn rule. The pair ( $X_{i}, Y_{i}$ ) is eliminated from the training sequence if the $k$-NN rule (based on the remaining $n-1$ pairs) incorrectly classifies $X_{i}$. The 1 -NN rule is used with the edited data. Show that this rule is consistent whenever $X$ has a density if $k \rightarrow \infty$ and $k / n \rightarrow 0$. Related papers: Wilson (1972), Wagner (1973), Penrod and Wagner (1977), and Devijver and Kittler (1980).

## 20

## Tree Classifiers

Classification trees partition $\mathcal{R}^{d}$ into regions, often hyperrectangles parallel to the axes. Among these, the most important are the binary classification trees, since they have just two children per node and are thus easiest to manipulate and update. We recall the simple terminology of books on data structures. The top of a binary tree is called the root. Each node has either no child (in that case it is called a terminal node or leaf), a left child, a right child, or a left child and a right child. Each node is the root of a tree itself. The trees rooted at the children of a node are called the left and right subtrees of that node. The depth of a node is the length of the path from the node to the root. The height of a tree is the maximal depth of any node.
Trees with more than two children per node can be reduced to binary trees by


Figure 20.1. A binary tree.
a simple device-just associate a left child with each node by selecting the oldest child in the list of children. Call the right child of a node its next sibling (see Figures 20.2 and 20.3). The new binary tree is called the oldest-child/next-sibling binary tree (see, e.g., Cormen, Leiserson, and Rivest (1990) for a general introduction). We only mention this particular mapping because it enables us to only consider binary trees for simplicity.


Figure 20.2. Ordered tree: The children are ordered from oldest to


Figure 20.3. The corresponding binary tree. youngest.

In a classification tree, each node represents a set in the space $\mathcal{R}^{d}$. Also, each node has exactly two or zero children. If a node $u$ represents the set $A$ and its children $u^{\prime}, u^{\prime \prime}$ represent $A^{\prime}$ and $A^{\prime \prime}$, then we require that $A=A^{\prime} \cup A^{\prime \prime}$ and $A^{\prime} \cap A^{\prime \prime}=$ $\emptyset$. The root represents $\mathcal{R}^{d}$, and the leaves, taken together, form a partition of $\mathcal{R}^{d}$.

Assume that we know $x \in A$. Then the question "is $x \in A$ ?" should be answered in a computationally simple manner so as to conserve time. Therefore, if $x=$ ( $x^{(1)}, \ldots, x^{(d)}$ ), we may just limit ourselves to questions of the following forms:
(i) Is $x^{(i)} \leq \alpha$ ? This leads to ordinary binary classification trees with partitions into hyperrectangles.
(ii) Is $a_{1} x^{(1)}+\cdots+a_{d} x^{(d)} \leq \alpha$ ? This leads to BSP trees (binary space partition trees). Each decision is more time consuming, but the space is more flexibly cut up into convex polyhedral cells.
(iii) Is $\|x-z\| \leq \alpha$ ? (Here $z$ is a point of $\mathcal{R}^{d}$, to be picked for each node.) This induces a partition into pieces of spheres. Such trees are called sphere trees.
(iv) Is $\psi(x) \geq 0$ ? Here, $\psi$ is a nonlinear function, different for each node. Every classifier can be thought of as being described in this format-decide class one if $\psi(x) \geq 0$. However, this misses the point, as tree classifiers should really be built up from fundamental atomic operations and queries such as those listed in (i)-(iii). We will not consider such trees any further.


Figure 20.4. Partition induced by an ordinary binary tree.


Figure 20.6. Partition induced by a BSP tree.


Figure 20.5. Corresponding tree.


Figure 20.7. Partition induced by a sphere tree.

We associate a class in some manner with each leaf in a classification tree. The tree structure is usually data dependent, as well, and indeed, it is in the construction itself where methods differ. If a leaf represents region $A$, then we say that the classifier $g_{n}$ is natural if

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i: X_{i} \in A} Y_{i}>\sum_{i: X_{i} \in A}\left(1-Y_{i}\right), x \in A \\ 0 & \text { otherwise. }\end{cases}
$$

That is, in every leaf region, we take a majority vote over all ( $X_{i}, Y_{i}$ )'s with $X_{i}$ in the same region. Ties are broken, as usual, in favor of class 0 . In this set-up, natural tree classifiers are but special cases of data-dependent partitioning rules. The latter are further described in detail in Chapter 21.


Figure 20.8. A natural classifier based on an ordinary binary tree. The decision is 1 in regions where points with label 1 form a majority. These areas are shaded.

Regular histograms can also be thought of as natural binary tree classifiers-the construction and relationship is obvious. However, as $n \rightarrow \infty$, histograms change size, and usually, histogram partitions are not nested as $n$ grows. Trees offer the exciting perspective of fully dynamic classification-as data are added, we may update the tree slightly, say, by splitting a leaf or so, to obtain an updated classifier.

The most compelling reason for using binary tree classifiers is to explain complicated data and to have a classifier that is easy to analyze and understand. In fact, expert system design is based nearly exclusively upon decisions obtained by going down a binary classification tree. Some argue that binary classification trees are preferable over BSP trees for this simple reason. As argued in Breiman, Friedman, Olshen, and Stone (1984), trees allow mixing component variables that are heterogeneous-some components may be of a nonnumerical nature, others may represent integers, and still others may be real numbers.

### 20.1 Invariance

Nearly all rules in this chapter and in Chapters 21 and 30 show some sort of invariance with respect to certain transformations of the input. This is often a major asset in pattern recognition methods. We say a rule $g_{n}$ is invariant under transformation $T$ if

$$
g_{n}\left(T(x) ; T\left(X_{1}\right), Y_{1}, \ldots, T\left(X_{n}\right), Y_{n}\right)=g_{n}\left(x ; X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right)
$$

for all values of the arguments. In this sense, we may require translation invariance, rotation invariance, linear translation invariance, and monotone transformation invariance $(T(\cdot)$ maps each coordinate separately by a strictly increasing but possibly nonlinear function).

Monotone transformation invariance frees us from worries about the kind of measuring unit. For example, it would not matter whether earthquakes were measured on a logarithmic (Richter) scale or a linear scale. Rotation invariance matters of course in situations in which input data have no natural coordinate axis system. In many cases, data are of the ordinal form-colors and names spring to mind-and
ordinal values may be translated into numeric values by creating bit vectors. Here, distance loses its physical meaning, and any rule that uses ordinal data perhaps mixed in with numerical data should be monotone transformation invariant.
Tree methods that are based upon perpendicular splits are usually (but not always) monotone transformation invariant and translation invariant. Tree methods based upon linear hyperplane splits are sometimes linear transformation invariant.
The partitions of space cause some problems if the data points can line up along hyperplanes. This is just a matter of housekeeping, of course, but the fact that some projections of $X$ to a line have atoms or some components of $X$ have atoms will make the proofs heavier to digest. For this reason only, we assume throughout this chapter that $X$ has a density $f$. As typically no conditions are put on $f$ in our consistency theorems, it will be relatively easy to generalize them to all distributions. The density assumption affords us the luxury of being able to say that, with probability one, no $d+1$ points fall in a hyperplane, no $d$ points fall in a hyperplane perpendicular to one axis, no $d-1$ points fall in a hyperplane perpendicular to two axes, etcetera. If a rule is monotone transformation invariant, we can without harm transform all the data as follows for the purpose of analysis only. Let $f_{1}, \ldots, f_{d}$ be the marginal densities of $X$ (see Problem 20.1), with corresponding distribution functions $F_{1}, \ldots, F_{d}$. Then replace in the data each $X_{i}$ by $T\left(X_{i}\right)$, where

$$
T(x)=\left(F_{1}\left(x^{(1)}\right), \ldots, F_{d}\left(x^{(d)}\right)\right) .
$$

Each component of $T\left(X_{i}\right)$ is now uniformly distributed on $[0,1]$. Of course, as we do not know $T$ beforehand, this device could only be used in the analysis. The transformation $T$ will be called the uniform marginal transformation. Observe that the original density is now transformed into another density.

### 20.2 Trees with the $X$-Property

It is possible to prove the convergence of many tree classifiers all at once. What is needed, clearly, is a partition into small regions, yet, most majority votes should be over sufficiently large sample. In many of the cases considered, the form of the tree is determined by the $X_{i}$ 's only, that is, the labels $Y_{i}$ do not play a role in constructing the partition, but they are used in voting. This is of course rather simplistic, but as a start, it is very convenient. We will say that the classification tree has the $X$-property, for lack of a better mnemonic. Let the leaf regions be $\left\{A_{1}, \ldots, A_{N}\right\}$ (with $N$ possibly random). Define $N_{j}$ as the number of $X_{i}$ 's falling in $A_{j}$. As the leaf regions form a partition, we have $\sum_{j=1}^{N} N_{j}=n$. $\operatorname{By} \operatorname{diam}\left(A_{j}\right)$ we mean the diameter of the cell $A_{j}$, that is, the maximal distance between two points of $A_{j}$. Finally, decisions are taken by majority vote, so for $x \in A_{j}, 1 \leq j \leq N$,
the rule is

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i: X_{i} \in A_{j}} Y_{i}>\sum_{i: X_{i} \in A_{j}}\left(1-Y_{i}\right), x \in A_{j} \\ 0 & \text { otherwise. }\end{cases}
$$

$A(x)$ denotes the set of the partition $\left\{A_{1}, \ldots, A_{N}\right\}$ into which $x$ falls, and $N(x)$ is the number of data points falling in this set. Recall that the general consistency result given in Theorem 6.1 is applicable in such cases. Consider a natural classification tree as defined above and assume the $X$-property. Theorem 6.1 states that then $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ if
(1) $\operatorname{diam}(A(X)) \rightarrow 0$ in probability,
(2) $N(X) \rightarrow \infty$ in probability.

A more general, but also more complicated consistency theorem is proved in Chapter 21. Let us start with the simplest possible example. We verify the conditions of Theorem 6.1 for the $k$-spacing rule in one dimension. This rule partitions the real line by using the $k$-th, $2 k$-th (and so on) order statistics (Mahalanobis, (1961); see also Parthasarathy and Bhattacharya (1961)).


Figure 20.9. A 3-spacing classifier.

Formally, let $k<n$ be a positive integer, and let $X_{(1)}, X_{(2)}, \ldots, X_{(n)}$ be the order statistics of the data points. Recall that $X_{(1)}, X_{(2)}, \ldots, X_{(n)}$ are obtained by permuting $X_{1}, \ldots, X_{n}$ in such a way that $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(n)}$. Note that this ordering is unique with probability one as $X$ has a density. We partition $\mathcal{R}$ into $N$ intervals $A_{1}, \ldots, A_{N}$, where $N=\lceil n / k\rceil$, such that for $j=1, \ldots, N-1, A_{j}$ satisfies

$$
X_{(k(j-1)+1)}, \ldots, X_{(k j)} \in A_{j}
$$

and the rightmost cell $A_{N}$ satisfies

$$
X_{(k(N-1)+1)}, \ldots, X_{(n)} \in A_{N}
$$

We have not specified the endpoints of each cell of the partition. For simplicity, let the borders between $A_{j}$ and $A_{j+1}$ be put halfway between the rightmost data point in $A_{j}$ and leftmost data point in $A_{j+1}, j=1, \ldots, N-1$.

The classification rule $g_{n}$ is defined in the usual way:

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{n} I_{\left\{X_{i} \in A(x), Y_{i}=1\right\}}>\sum_{i=1}^{n} I_{\left\{X_{i} \in A(x), Y_{i}=0\right\}} \\ 0 & \text { otherwise } .\end{cases}
$$

Theorem 20.1. Let $g_{n}$ be the $k$-spacing classifier given above. Assume that the distribution of $X$ has a density $f$ on $\mathcal{R}$. Then the classification rule $g_{n}$ is consistent, that is, $\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}=L^{*}$, if $k \rightarrow \infty$ and $k / n \rightarrow 0$ as $n$ tends to infinity.

REMARK. We discuss various generalizations of this rule in Chapter 21.

Proof. We check the conditions of Theorem 6.1, as the partition has the $X$-property. Condition (2) is obvious from $k \rightarrow \infty$.

To establish condition (1), fix $\epsilon>0$. Note that by the invariance of the rule under monotone transformations, we may assume without loss of generality that $f$ is the uniform density on $[0,1]$. Among the intervals $A_{1}, \ldots, A_{N}$, there are at most $1 / \epsilon$ disjoint intervals of length greater than $\epsilon$ in $[0,1]$. Thus,

$$
\begin{aligned}
& \mathbf{P}\{\operatorname{diam}(A(X))>\epsilon\} \\
& \leq \frac{1}{\epsilon} \mathbf{E}\left\{\max _{1 \leq j \leq N} \mu\left(A_{j}\right)\right\} \\
& \leq \frac{1}{\epsilon} \mathbf{E}\left\{\left(\max _{1 \leq j \leq N} \mu_{n}\left(A_{j}\right)+\max _{1 \leq j \leq N}\left|\mu\left(A_{j}\right)-\mu_{n}\left(A_{j}\right)\right|\right)\right\} \\
& \leq \frac{1}{\epsilon}\left(\frac{k}{n}+\mathbf{E}\left\{\sup _{A}\left|\mu(A)-\mu_{n}(A)\right|\right\}\right),
\end{aligned}
$$

where the supremum is taken over all intervals in $\mathcal{R}$. The first term within the parentheses converges to zero by the second condition of the theorem, while the second term goes to zero by an obvious extension of the classical Glivenko-Cantelli theorem (Theorem 12.4). This completes the proof.

We will encounter several trees in which the partition is determined by a small fraction of the data, such as binary k-d trees and quadtrees. In these cases, condition (2) of Theorem 6.1 may be verified with the help of the following lemma:

Lemma 20.1. Let $p_{1}, \ldots, p_{k}$ be a probability vector. Let $N_{1}, \ldots, N_{k}$ be multinomially distributed random variables with parameters $n$ and $p_{1}, \ldots, p_{k}$. Then if the random variable $X$ is independent of $N_{1}, \ldots, N_{k}$, and $\mathbf{P}\{X=i\}=p_{i}$, we have for any $M$,

$$
\mathbf{P}\left\{N_{X} \leq M\right\} \leq \frac{(2 M+4) k}{n}
$$

(Note: this probability goes to 0 if $k / n \rightarrow 0$ uniformly over all probability vectors with $k$ components!)

Proof. Let $Z_{i}$ be binomial with parameters $n$ and $p_{i}$. Then

$$
\begin{aligned}
\mathbf{P}\left\{N_{X} \leq M\right\} & \leq \sum_{i: n p_{i} \leq 2 M} p_{i}+\sum_{i: n p_{i}>2 M} p_{i} \mathbf{P}\left\{Z_{i} \leq M\right\} \\
& \leq \frac{2 M k}{n}+\sum_{i: n p_{i}>2 M} p_{i} \mathbf{P}\left\{Z_{i}-\mathbf{E} Z_{i} \leq M-n p_{i}\right\} \\
& \leq \frac{2 M k}{n}+\sum_{i: n p_{i}>2 M} p_{i} \mathbf{P}\left\{Z_{i}-\mathbf{E} Z_{i} \leq-\frac{\mathbf{E} Z_{i}}{2}\right\}
\end{aligned}
$$

$$
\leq \frac{2 M k}{n}+\sum_{i: n p_{i}>2 M} 4 p_{i} \frac{\operatorname{Var}\left\{Z_{i}\right\}}{\left(\mathbf{E}\left\{Z_{i}\right\}\right)^{2}}
$$

(by Chebyshev's inequality)

$$
\begin{aligned}
& \leq \frac{2 M k}{n}+\sum_{i: n p_{i}>2 M} 4 p_{i} \frac{1}{n p_{i}} \\
& \leq \frac{(2 M+4) k}{n} .
\end{aligned}
$$

The previous lemma implies that for any binary tree classifier constructed on the basis of $X_{1}, \ldots, X_{k}$ with $k+1$ regions, $N(X) \rightarrow \infty$ in probability whenever $k /(n-k) \rightarrow 0$ (i.e., $k / n \rightarrow 0$ ). It suffices to note that we may take $M$ arbitrarily large but fixed in Lemma 20.1. This remark saves us the trouble of having to verify just how large or small the probability mass of the region is. In fact, it also implies that we should not worry so much about regions with few data points. What matters more than anything else is the number of regions. Stopping rules based upon cardinalities of regions can effectively be dropped in many cases!

### 20.3 Balanced Search Trees

Balanced multidimensional search trees are computationally attractive. Binary trees with $n$ leaves have $O(\log n)$ height, for example, when at each node, the size of every subtree is at least $\alpha$ times the size of the other subtree rooted at the parent, for some constant $\alpha>0$. It is thus important to verify the consistency of balanced search trees used in classification. We again consider binary classification trees with the $X$-property and majority votes over the leaf regions. Take for example a tree in which we split every node perfectly, that is, if there are $n$ points, we find the median according to one coordinate, and create two subtrees of sizes $\lfloor(n-1) / 2\rfloor$ and $\lceil(n-1) / 2\rceil$. The median itself stays behind and is not sent down to the subtrees. Repeat this for $k$ levels of nodes, at each level cutting along the next coordinate axe in a rotational manner. This leads to $2^{k}$ leaf regions, each having at least $n / 2^{k}-k$ points and at most $n / 2^{k}$ points. Such a tree will be called a median tree.


Figure 20.10. Median tree with four leaf regions in $\mathcal{R}^{2}$.

Setting up such a tree is very easy, and hence such trees may appeal to certain programmers. In hypothesis testing, median trees were studied by Anderson (1966).

Theorem 20.2. Natural classifiers based upon median trees with $k$ levels $\left(2^{k}\right.$ leaf regions) are consistent $\left(\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}\right)$ whenever $X$ has a density, if

$$
\frac{n}{k 2^{k}} \rightarrow \infty \quad \text { and } \quad k \rightarrow \infty
$$

(Note: the conditions of $k$ are fulfilled if $k \leq \log _{2} n-2 \log _{2} \log _{2} n, k \rightarrow \infty$.)
We may prove the theorem by checking the conditions of Theorem 6.1. Condition (2) follows trivially by the fact that each leaf region contains at least $n / 2^{k}-k$ points and the condition $n /\left(k 2^{k}\right) \rightarrow \infty$. Thus, we need only verify the first condition of Theorem 6.1. To make the proof more transparent, we first analyze a closely related hypothetical tree, the theoretical median tree. Also, we restrict the analysis to $d=2$. The multidimensional extension is straightforward. The theoretical median tree rotates through the coordinates and cuts each hyperrectangle precisely so that the two new hyperrectangles have equal $\mu$-measure.

| 1/8 |  | 1/8 |
| :---: | :---: | :---: |
| 1/8 |  |  |
|  |  | 1/8 |
| 1/8 | 1/8 |  |
| 1/8 | 1/8 |  |

Figure 20.11. Theoretical median tree with three levels of cuts.

Observe that the rule is invariant under monotone transformations of the coordinate axes. Recall that in such cases there is no harm in assuming that the marginal distributions are all uniform on $[0,1]$. We let $\left\{H_{i}, V_{i}\right\}$ denote the horizontal and vertical sizes of the rectangles after $k$ levels of cuts. Of course, we begin with $H_{1}=V_{1}=1$ when $k=0$. We now show that, for the theoretical median tree, $\operatorname{diam}(A(X)) \rightarrow 0$ in probability, as $k \rightarrow \infty$. Note that $\operatorname{diam}(A(X)) \leq H(X)+$ $V(X)$, where $H(X)$ and $V(X)$ are the horizontal and vertical sizes of the rectangle $A(X)$. We show that if $k$ is even,

$$
\mathbf{E}\{H(X)+V(X)\}=\frac{2}{2^{k / 2}},
$$

from which the claim follows. After the $k$-th round of splits, since all $2^{k}$ rectangles have equal probability measure, we have

$$
\mathbf{E}\{H(X)+V(X)\}=\sum_{i=1}^{2^{k}} \frac{1}{2^{k}}\left(H_{i}+V_{i}\right) .
$$

Apply another round of splits, all vertical. Then each term $\frac{1}{2^{k}}\left(H_{i}+V_{i}\right)$ spawns, so to speak, two new rectangles with horizontal and vertical sizes ( $H_{i}^{\prime}, V_{i}$ ) and
( $H_{i}^{\prime \prime \prime}, V_{i}$ ) with $H_{i}=H_{i}^{\prime}+H_{i}^{\prime \prime \prime}$ that contribute

$$
\frac{1}{2^{k+1}}\left(H_{i}^{\prime}+V_{i}\right)+\frac{1}{2^{k+1}}\left(H_{i}^{\prime \prime \prime}+V_{i}\right)=\frac{1}{2^{k+1}} H_{i}+\frac{1}{2^{k}} V_{i}
$$

The next round yields horizontal splits, with total contribution now (see Figure 20.12)

$$
\begin{aligned}
\frac{1}{2^{k+2}}\left(H_{i}^{\prime}+V_{i}^{\prime}+H_{i}^{\prime}+V_{i}^{\prime \prime}+H_{i}^{\prime \prime \prime}+V_{i}^{\prime \prime \prime}+H_{i}^{\prime \prime \prime}+V_{i}^{\prime \prime \prime \prime}\right) & =\frac{1}{2^{k+2}}\left(2 H_{i}+2 V_{i}\right) \\
& =\frac{1}{2^{k+1}}\left(H_{i}+V_{i}\right)
\end{aligned}
$$

Thus, over two iterations of splits, we see that $\mathbf{E}\{H(X)+V(X)\}$ is halved, and the claim follows by simple induction.

We show now what happens in the real median tree when cuts are based upon a random sample. We deviate of course from the theoretical median tree, but consistency is preserved. The reason, seen intuitively, is that if the number of points in a cell is large, then the sample median will be close to the theoretical median, so that the shrinking-diameter property is preserved. The methodology followed here shows how one may approach the analysis in general by separating the theoretical model from the sample-based model.

Proof of Theorem 20.2. As we noted before, all we have to show is that $\operatorname{diam}(A(X)) \rightarrow 0$ in probability. Again, we assume without loss of generality that the marginals of $X$ are uniform $[0,1]$, and that $d=2$. Again, we show that $\mathbf{E}\{H(X)+V(X)\} \rightarrow 0$.


Figure 20.12. A rectangle after two rounds of splits.

If a rectangle of probability mass $p_{i}$ and sizes $H_{i}, V_{i}$ is split into four rectangles as in Figure 20.12, with probability masses $p_{i}^{\prime}, p_{i}^{\prime \prime}, p_{i}^{\prime \prime \prime}, p_{i}^{\prime \prime \prime \prime}$, then the contribution $p_{i}\left(H_{i}+V_{i}\right)$ to $\mathbf{E}\{H(X)+V(X)\}$ becomes

$$
p_{i}^{\prime}\left(H_{i}^{\prime}+V_{i}^{\prime}\right)+p_{i}^{\prime \prime}\left(H_{i}^{\prime \prime}+V_{i}^{\prime \prime}\right)+p_{i}^{\prime \prime \prime}\left(H_{i}^{\prime \prime \prime}+V_{i}^{\prime \prime \prime}\right)+p_{i}^{\prime \prime \prime \prime}\left(H_{i}^{\prime \prime \prime \prime}+V_{i}^{\prime \prime \prime \prime}\right)
$$

after two levels of cuts. This does not exceed

$$
\frac{p_{i}}{2}(1+\epsilon)\left(H_{i}+V_{i}\right)
$$

if

$$
\begin{gathered}
\max \left(p_{i}^{\prime}+p_{i}^{\prime \prime}, p_{i}^{\prime \prime \prime}+p_{i}^{\prime \prime \prime \prime}\right) \leq \frac{1}{2} \sqrt{1+\epsilon} p_{i}, \\
\max \left(p_{i}^{\prime}, p_{i}^{\prime \prime}\right) \leq \frac{1}{2} \sqrt{1+\epsilon}\left(p_{i}^{\prime}+p_{i}^{\prime \prime}\right),
\end{gathered}
$$

and

$$
\max \left(p_{i}^{\prime \prime \prime}, p_{i}^{\prime \prime \prime \prime}\right) \leq \frac{1}{2} \sqrt{1+\epsilon}\left(p_{i}^{\prime \prime \prime}+p_{i}^{\prime \prime \prime \prime}\right)
$$

that is, when all three cuts are within $(1 / 2) \sqrt{1+\epsilon}$ of the true median. We call such " $(1 / 2) \sqrt{1+\epsilon}$ " cuts good. If all cuts are good, we thus note that in two levels of cuts, $\mathbf{E}\{H(X)+V(X)\}$ is reduced by $(1+\epsilon) / 2$. Also, all $p_{i}$ 's decrease at a controlled rate. Let $G$ be the event that all $1+2+\cdots+2^{k-1}$ cuts in a median tree with $k$ levels are good. Then, at level $k$, all $p_{i}$ 's are at most $(\sqrt{1+\epsilon} / 2)^{k}$. Thus,

$$
\begin{aligned}
\sum_{i=1}^{2^{k}} p_{i}\left(H_{i}+V_{i}\right) & \leq\left(\frac{\sqrt{1+\epsilon}}{2}\right)^{k} \sum_{i=1}^{2^{k}}\left(H_{i}+V_{i}\right) \\
& \leq\left(3+2^{k / 2+1}\right)\left(\frac{\sqrt{1+\epsilon}}{2}\right)^{k}
\end{aligned}
$$

since $\sum_{i=1}^{2^{k}}\left(H_{i}+V_{i}\right) \leq 4+2+4+8+\cdots+2^{k / 2} \leq 2^{k / 2+1}+3$ if $k$ is even. Hence, after $k$ levels of cuts,

$$
\mathbf{E}\{H(X)+V(X)\} \leq\left(3+2^{k / 2+1}\right) \mathbf{P}\left\{G^{c}\right\}+\left(3+2^{k / 2+1}\right)\left(\frac{\sqrt{1+\epsilon}}{2}\right)^{k}
$$

The last term tends to zero if $\epsilon$ is small enough. We bound $\mathbf{P}\left\{G^{c}\right\}$ by $2^{k}$ times the probability that one cut is bad. Let us cut a cell with $N$ points and probability content $p$ in a given direction. A quick check of the median tree shows that given the position and size of the cell, the $N$ points inside the cell are distributed in an i.i.d. manner according to the restriction of $\mu$ to the cell. After the cut, we have $\lfloor(N-1) / 2\rfloor$ and $\lceil(N-1) / 2\rceil$ points in the new cells, and probability contents $p^{\prime}$ and $p^{\prime \prime}$. It is clear that we may assume without loss of generality that $p=1$. Thus, if all points are projected down in the direction of the cut, and $F$ and $F_{N}$ denote the distribution function and empirical distribution function of the obtained one-dimensional data, then

$$
\begin{aligned}
& \mathbf{P}\{\text { cut is not } \operatorname{good} \mid N\} \\
& \leq \mathbf{P}\left\{p^{\prime}>\frac{\sqrt{1+\epsilon}}{2} \text { or } \left.p^{\prime \prime}>\frac{\sqrt{1+\epsilon}}{2} \right\rvert\, N\right\} \\
& \leq \mathbf{P}\left\{\text { for some } x, F(x)>\frac{\sqrt{1+\epsilon}}{2}, \text { and } \left.F_{N}(x) \leq \frac{1}{2} \right\rvert\, N\right\}
\end{aligned}
$$

$$
\begin{aligned}
& \leq \mathbf{P}\left\{\left.\sup _{x}\left(F(x)-F_{N}(x)\right)>\frac{1}{2}(\sqrt{1+\epsilon}-1) \right\rvert\, N\right\} \\
& \leq 2 \exp \left(-\frac{1}{2} N(\sqrt{1+\epsilon}-1)^{2}\right)
\end{aligned}
$$

(by Theorem 12.9)

$$
\begin{aligned}
\leq & 2 \exp \left(-\left(\frac{n}{2^{k+1}}-\frac{k}{2}\right)(\sqrt{1+\epsilon}-1)^{2}\right) \\
& \left(\text { as } N \geq n /\left(2^{k}\right)-k\right)
\end{aligned}
$$

Hence, for $n$ large enough,

$$
\mathbf{P}\left\{G^{c}\right\} \leq 2^{k+1} e^{-\left(n / 2^{k+2}\right)(\sqrt{1+\epsilon}-1)^{2}}
$$

and $2^{k / 2} \mathbf{P}\left\{G^{c}\right\} \rightarrow 0$ if $n /\left(k 2^{k}\right) \rightarrow \infty$.

### 20.4 Binary Search Trees

The simplest trees to analyze are those whose structure depends in a straightforward way on the data. To make this point, we begin with the binary search tree and its multivariate extension, the k-d tree (see Cormen, Leiserson, and Rivest (1990) for the binary search tree; for multivariate binary trees, we refer to Samet (1990b)). A full-fledged k -d tree is defined as follows: we promote the first data point $X_{1}$ to the root and partition $\left\{X_{2}, \ldots, X_{n}\right\}$ into two sets: those whose first coordinate exceeds that of $X_{1}$, and the remaining points. Within each set, points are ordered by original index. The former set is used to build the right subtree of $X_{1}$ and the latter to construct the left subtree of $X_{1}$. For each subtree, the same construction is applied recursively with only one variant: at depth $l$ in the tree, the $(l \bmod d+1)$-st coordinate is used to split the data. In this manner, we rotate through the coordinate axes periodically.

Attach to each leaf two new nodes, and to each node with one child a second child. Call these new nodes external nodes. Each of the $n+1$ external nodes correspond to a region of $\mathcal{R}^{d}$, and collectively the external nodes define a partition of $\mathcal{R}^{d}$ (if we define exactly what happens on the boundaries between regions).


Figure 20.13. Ak-d tree of 15 random points on the plane and the induced partition.
Put differently, we may look at the external nodes as the leaves of a new tree with $2 n+1$ nodes and declare this new tree to be our new binary classification tree. As there are $n+1$ leaf regions and $n$ data points, the natural binary tree classifier it induces is degenerate-indeed, all external regions contain very few points. Clearly, we must have a mechanism for trimming the tree to insure better populated leaves. Let us look at just three naïve strategies. For convenience, we assume that the data points determining the tree are not counted when taking a majority vote over the cells. As the number of these points is typically much smaller than $n$, this restriction does not make a significant difference.
(1) Fix $k<n$, and construct a k-d tree with $k$ internal nodes and $k+1$ external nodes, based on the first $k$ data points $X_{1}, \ldots, X_{k}$. Classify by majority vote over all $k+1$ regions as in natural classification tees (taking the data pairs $\left(X_{k+1}, Y_{k+1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ into account). Call this the chronological $k-d$ tree.
(2) Fix $k$ and truncate the k-d tree to $k$ levels. All nodes at level $k$ are declared leaves and classification is again by majority vote over the leaf regions. Call this the deep $k$-d tree.
(3) Fix $k$ and trim the tree so that each node represents at least $k$ points in the original construction. Consider the maximal such tree. The number of regions here is random, with between 1 and $n / k$ regions. Call this the wellpopulated $k$-d tree.

Let the leaf regions be $\left\{A_{1}, \ldots, A_{N}\right\}$ (with $N$ possibly random), and denote the leaf nodes by $u_{1}, \ldots, u_{N}$. The strict descendants of $u_{i}$ in the full k-d tree have indices that we will collect in an index set $I_{i}$. Define $\left|I_{i}\right|=N_{i}$. As the leaf regions form a partition, we have

$$
\sum_{i=1}^{N} N_{i}=n-N
$$

because the leaf nodes themselves are not counted in $I_{i}$. Voting is by majority vote, so the rule is

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i \in I_{i}} Y_{i}>\sum_{i \in I_{i}}\left(1-Y_{i}\right), x \in A_{i} \\ 0 & \text { otherwise } .\end{cases}
$$

### 20.5 The Chronological k-d Tree

Here we have $N=k+1$. Also, $\left(\mu\left(A_{1}\right), \ldots, \mu\left(A_{k+1}\right)\right)$ are distributed as uniform spacings. That is, if $U_{1}, \ldots, U_{k}$ are i.i.d. uniform [0, 1] random variables defining $k+1$ spacings $U_{(1)}, U_{(2)}-U_{(1)}, \ldots, U_{(k)}-U_{(k-1)}, 1-U_{(k)}$ by their order statistics $U_{(1)} \leq U_{(2)} \leq \cdots \leq U_{(k)}$, then these spacings are jointly distributed as $\left(\mu\left(A_{1}\right), \ldots, \mu\left(A_{k+1}\right)\right)$. This can be shown by induction. When $U_{k+1}$ is added, $U_{k+1}$ first picks a spacing with probability equal to the size of the spacing. Then it cuts that spacing in a uniform manner. As the same is true when the chronological k -d tree grows by one leaf, the property follows by induction on $k$.

Theorem 20.3. We have $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ for all distributions of $X$ with a density for the chronological $k$-d tree classifier whenever

$$
k \rightarrow \infty, \quad \text { and } \quad k / n \rightarrow 0
$$

Proof. We verify the conditions of Theorem 6.1. As the number of regions is $k+1$, and the partition is determined by the first $k$ data points, condition (2) immediately follows from Lemma 20.1 and the remark following it.

Condition (1) of Theorem 6.1 requires significantly more work. We verify condition (1) for $d=2$, leaving the straightforward extension to $\mathcal{R}^{d}, d>2$, to the reader. Throughout, we assume without loss of generality that the marginal distributions are uniform $[0,1]$. We may do so by the invariance properties discussed earlier. Fix a point $x \in \mathcal{R}^{d}$. We insert points $X_{1}, \ldots, X_{k}$ into an initially empty k-d tree and let $R_{1}, \ldots, R_{k}$ be the rectangles containing $x$ just after these points were inserted. Note that $R_{1} \supseteq R_{2} \supseteq \cdots \supseteq R_{k}$. Assume for simplicity that the integer $k$ is a perfect cube and set $l=k^{1 / 3}, m=k^{2 / 3}$. Define the distances from $x$ to the sides of $R_{i}$ by $H_{i}, H_{i}^{\prime}, V_{i}, V_{i}^{\prime}$ (see Figure 20.14).


Figure 20.14. A rectangle $R_{i}$ containing $x$ with its distances to the sides.

We construct a sequence of events that forces the diameter of $R_{k}$ to be small with high probability. Let $\epsilon$ be a small positive number to be specified later.

Denote the four squares with opposite vertices $x, x+( \pm \epsilon, \pm \epsilon)$ by $C_{1}, C_{2}, C_{3}$, $C_{4}$. Then define the following events:

$$
\begin{aligned}
& E_{1}=\bigcap_{i=1}^{4}\left\{C_{i} \cap\left\{X_{1}, \ldots, X_{l}\right\} \neq \emptyset\right\}, \\
& E_{2}=\left\{\max \left(V_{l}, V_{l}^{\prime}, H_{l}, H_{l}^{\prime}\right) \leq \epsilon\right\} \text {, } \\
& E_{3}=\left\{\min \left(\max \left(V_{l}, V_{l}^{\prime}\right), \max \left(H_{l}, H_{l}^{\prime}\right)\right) \leq \epsilon<\max \left(V_{l}, V_{l}^{\prime}, H_{l}, H_{l}^{\prime}\right)\right\} \text {, } \\
& E_{4}=\left\{\text { at least three of } V_{m}, V_{m}^{\prime}, H_{m}, H_{m}^{\prime} \text { are } \leq \epsilon\right\} \text {, } \\
& E_{5}=\left\{\max \left(V_{m}, V_{m}^{\prime}, H_{m}, H_{m}^{\prime}\right) \leq \epsilon\right\}, \\
& E_{6}=\left\{\max \left(V_{k}, V_{k}^{\prime}, H_{k}, H_{k}^{\prime}\right) \leq \epsilon\right\} .
\end{aligned}
$$

If $E_{2}, E_{5}$, or $E_{6}$ hold, then $\operatorname{diam}\left(R_{k}\right) \leq \epsilon \sqrt{8}$. Assume that we find a set $B \subset \mathcal{R}^{d}$ such that $\mathbf{P}\{X \in B\}=1$, and for all $x \in B$, and sufficiently small $\epsilon>0$, $\mathbf{P}\left\{E_{2} \cup E_{5} \cup E_{6}\right\} \rightarrow 1$ as $k \rightarrow \infty$. Then, by the Lebesgue dominated convergence theorem, $\operatorname{diam}(A(X)) \rightarrow 0$ in probability, and condition (1) of Theorem 6.1 would follow. In the remaining part of the proof we define such a set $B$, and show that $\mathbf{P}\left\{E_{2} \cup E_{5} \cup E_{6}\right\} \rightarrow 1$. This will require some work.

We define the set $B$ in terms of the density $f$ of $X . x \in B$ if and only if
(1) $\min _{1 \leq i \leq 4} \int_{C_{i}} f(z) d z>0$ for all $\epsilon>0$ small enough;
(2) $\inf _{\substack{\text { rectangles } R \text { containing } x, \\ \text { of diameter } \leq \epsilon \sqrt{8}}} \frac{\int_{R} f(z) d z}{\lambda(R)} \geq \frac{f(x)}{2}$ for all $\epsilon>0$ small enough;
(3) $f(x) \geq 0$.

That $\mathbf{P}\{X \in B\}=1$ follows from a property of the support (Problem 20.6), a corollary of the Jessen-Marcinkiewicz-Zygmund theorem (Corollary A.2; this implies (2) for almost all $x$ ), and the fact that for $\mu$-almost all $x, f(x)>0$.

It is easy to verify the following:

$$
\begin{aligned}
\mathbf{P}\left\{E_{6}^{c}\right\}= & \mathbf{P}\left\{E_{1}^{c}\right\} \\
& +\mathbf{P}\left\{E_{1} \cap E_{2}^{c} \cap E_{3}^{c}\right\}
\end{aligned}
$$

$$
\begin{aligned}
& +\mathbf{P}\left\{E_{1} \cap E_{2}^{c} \cap E_{3} \cap E_{4}^{c}\right\} \\
& +\mathbf{P}\left\{E_{1} \cap E_{2}^{c} \cap E_{3} \cap E_{4} \cap E_{5}^{c} \cap E_{6}^{c}\right\}
\end{aligned}
$$

We show that each term tends to 0 at $x \in B$.
Term 1. By the union bound,

$$
\begin{aligned}
\mathbf{P}\left\{E_{1}^{c}\right\} & \leq \sum_{i=1}^{4} \mathbf{P}\left\{X_{1} \notin C_{i}, \ldots, X_{l} \notin C_{i}\right\} \\
& \leq \sum_{i=1}^{4}\left(1-\mu\left(C_{i}\right)\right)^{l} \\
& \leq \exp \left(-l \min _{1 \leq i \leq 4} \mu\left(C_{i}\right)\right) \\
& \rightarrow 0
\end{aligned}
$$

by part (1) of the definition of $B$.
TERM 2. $E_{1} \subseteq E_{3}$ by a simple geometric argument. Hence $\mathbf{P}\left\{E_{1} \cap E_{3}^{c}\right\}=0$.
TERM 3. To show that $\mathbf{P}\left\{E_{1} \cap E_{2}^{c} \cap E_{3} \cap E_{4}^{c}\right\} \rightarrow 0$, we assume without loss of generality that $\max \left(V_{l}, V_{l}^{\prime}\right) \leq \epsilon$ while $\max \left(H_{l}, H_{l}^{\prime}\right)=a>\epsilon$. Let $\left\{X_{i}^{\prime}\right\}$ be a subset of $\left\{X_{i}, l<i \leq m\right\}$ consisting of those $X_{i}$ 's that fall in $R_{l}$. We introduce three notions in this sequence: first, $Z_{i}$ is the absolute value of the difference of the $x^{(2)}$-coordinates of $x$ and $X_{i}^{\prime}$. Let $W_{i}$ be the absolute value of the difference of the $x^{(1)}$-coordinates of $x$ and $X_{i}^{\prime}$. We re-index the sequence $X_{i}^{\prime}$ (and $W_{i}$ and $Z_{i}$ ) so that $i$ runs from 1 to $N$, where

$$
N=\sum_{i=l+1}^{m} I_{\left\{X_{i} \in R_{l}\right\}} .
$$

To avoid trivialities, assume that $N \geq 1$ (this will be shown to happen with probability tending to one). Call $X_{i}^{\prime}$ a record if $Z_{i}=\min \left(Z_{1}, \ldots, Z_{i}\right)$. Call $X_{i}^{\prime}$ a good point if $W_{i} \leq \epsilon$. An $X_{i}^{\prime}$ causes $\min \left(H_{m}, H_{m}^{\prime}\right) \leq \epsilon$ if that $X_{i}^{\prime}$ is a good point and a record, and if it defines a vertical cut. The alternating nature of the cuts makes our analysis a bit heavier than needed. We show here what happens when all directions are picked independently of each other, leaving the rotating-cuts case to the reader (Problem 20.8). Thus, if we set

$$
S_{i}=I_{\left\{X_{i}^{\prime} \text { is a record, } X_{i}^{\prime} \text { defines a vertical cut }\right\}}
$$

we have,

$$
E_{1} \cap E_{2}^{c} \cap E_{3} \cap E_{4}^{c} \subseteq\{N=0\} \cup\left\{\left\{\sum_{i=1}^{N} S_{i} I_{\left\{W_{i} \leq \epsilon\right\}}=0\right\} \cap\{N=0\}\right\}
$$

Re-index again and let $X_{1}^{\prime}, X_{2}^{\prime}, \ldots$ all be records. Note that given $X_{i}^{\prime}, X_{i+1}^{\prime}$ is distributed according to $f$ restricted to the rectangle $R^{\prime}$ of

$$
\begin{aligned}
& \text { height } \min \left(V_{l}, Z_{i}\right) \text { above } x, \\
& \text { height } \min \left(V_{l}^{\prime}, Z_{i}\right) \text { below } x \text {, } \\
& \text { width } H_{l} \text { to the left of } x \text {, } \\
& \text { width } H_{l}^{\prime} \text { to the right of } x .
\end{aligned}
$$

Call these four quantities $v, v^{\prime}, h, h^{\prime}$, respectively. Then

$$
\mathbf{P}\left\{W_{i+1} \leq \epsilon \mid X_{i}^{\prime}\right\} \geq \frac{\left(v+v^{\prime}\right) \epsilon f(x) / 2}{v+v^{\prime}}=\frac{\epsilon f(x)}{2}
$$

because the marginal distribution of an independent $X_{1}$ is uniform and thus, $\mathbf{P}\left\{X_{1} \in\right.$ $\left.R^{\prime} \mid R^{\prime}\right\} \leq v+v^{\prime}$, while, by property (2) of $B, \mathbf{P}\left\{X_{1} \in R^{\prime}, W_{1} \leq \epsilon \mid R^{\prime}\right\} \geq(v+$ $\left.v^{\prime}\right) \in f(x) / 2$.

Recall the re-indexing. Let $M$ be the number of records (thus, $M$ is the length of our sequence $X_{i}^{\prime}$ ). Then

$$
E_{1} \cap E_{2}^{c} \cap E_{3} \cap E_{4}^{c} \cap\{N>0\} \subseteq\left\{\sum_{i=1}^{M} I_{\left\{W_{i} \leq \epsilon\right\}} I_{\left\{X_{i}^{\prime} \text { is a vertical cut }\right\}}=0\right\}
$$

But as cuts have independently picked directions, and since $\mathbf{P}\left\{W_{i+1} \leq \epsilon \mid X_{i}^{\prime}\right\} \geq$ $\epsilon f(x) / 2$, we see that

$$
\mathbf{P}\left\{E_{1} \cap E_{2}^{c} \cap E_{3} \cap E_{4}^{c}, N>0\right\} \leq \mathbf{E}\left\{\left(1-\frac{\epsilon f(x)}{4}\right)^{M} I_{\{N>0\}}\right\} .
$$

We rewrite $M=\sum_{i=1}^{N} I_{\left\{X_{i}^{\prime} \text { is a record\} }\right.}$ and recall that the indicator variables in this sum are independent and are of mean $1 / i$ (Problem 20.9). Hence, for $c>0$,

$$
\begin{aligned}
\mathbf{E}\left\{c^{M}\right\} & =\mathbf{E}\left\{\prod_{i=1}^{N}\left(\frac{c}{i}+1-\frac{1}{i}\right)\right\} \\
& \leq \mathbf{E}\left\{e^{-(1-c) \sum_{i=1}^{N} 1 / i}\right\} \quad\left(\text { use } 1-x \leq e^{-x}\right) \\
& \leq \mathbf{E}\left\{\frac{1}{(N+1)^{1-c}}\right\} \quad\left(\text { use } \sum_{i=1}^{N} 1 / i \geq \log (N+1)\right)
\end{aligned}
$$

The latter formula remains valid even if $N=0$. Thus, with $c=1-\epsilon f(x) / 4$,

$$
\begin{aligned}
\mathbf{P}\left\{E_{1} \cap E_{2}^{c} \cap E_{3} \cap E_{4}^{c}\right\} & \leq \mathbf{E}\left\{\left(1-\frac{\epsilon f(x)}{4}\right)^{M}\right\} \\
& \leq \mathbf{E}\left\{\frac{1}{(N+1)^{1-c}}\right\}
\end{aligned}
$$

$N$ is binomial with parameters ( $m-l, \mu\left(R_{l}\right)$ ). We know from the introduction of this section that $\mu\left(R_{l}\right)$ is distributed as the minimum of $l$ i.i.d. uniform $[0,1]$ random variables. Thus, for $\delta>0, \mathbf{E}\left\{\mu\left(R_{l}\right)\right\}=1 /(l+1)$, and $\mathbf{P}\left\{\mu\left(R_{l}\right)<\delta / l\right\} \leq \delta$. Hence,

$$
\begin{aligned}
& \mathbf{E}\left\{\frac{1}{(N+1)^{1-c}}\right\} \\
& \quad \leq \mathbf{P}\left\{\mu\left(R_{l}\right)<\delta / l\right\}+\mathbf{E}\left\{\frac{1}{\left.(\operatorname{Binomial}(m-l, \delta / l)+1)^{1-c}\right\}}\right. \\
& \quad \leq \delta+\left(\frac{l}{2(m-l) \delta}\right)^{1-c}+\mathbf{P}\left\{\operatorname{Binomial}(m-l, \delta / l)<\frac{(m-l) \delta}{2 l}\right\} .
\end{aligned}
$$

The first term is small by choice of $\delta$. The second one is $O\left(k^{-(1-c) / 3}\right)$. The third one is bounded from above, by Chebyshev's inequality, by

$$
\frac{(m-l)(\delta / l)(1-\delta / l)}{((m-l) \delta / 2 l)^{2}} \leq \frac{4 l}{(m-l) \delta} \rightarrow 0 .
$$

Term 4. This term is handled exactly as Term 3. Note, however, that $l$ and $m$ now become $m$ and $k$ respectively. The convergence to 0 requires now $m /(k-m) \rightarrow 0$, which is still the case.

This concludes the proof of Theorem 20.3.

### 20.6 The Deep k-d Tree

Theorem 20.4. The deep $k$-d tree classifier is consistent (i.e., $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ ) for all distributions such that $X$ has a density, whenever

$$
\lim _{n \rightarrow \infty} k=\infty \quad \text { and } \quad \limsup _{n \rightarrow \infty} \frac{k}{\log n}<2
$$

Proof. In Problem 20.10, you are asked to show that $k \rightarrow \infty$ implies diam $(A(X))$ $\rightarrow 0$ in probability. Theorem 20.3 may be invoked here. We now show that the assumption $\lim \sup _{n \rightarrow \infty} k / \log n<2$ implies that $N(X) \rightarrow \infty$ in probability. Let $D$ be the depth (distance from the root) of $X$ when $X$ is inserted into a k-d tree having $n$ elements. Clearly, $N(X) \geq D-k$, so it suffices to show that $D-k \rightarrow \infty$ in probability. We know that $D /(2 \log n) \rightarrow 1$ in probability (see, e.g., Devroye (1988a) and Problem 20.10). This concludes the proof of the theorem.

### 20.7 Quadtrees

Quadtrees or hyperquadtrees are unquestionably the most prominent trees in computer graphics. Easy to manipulate and compact to store, they have found their way into mainstream computer science. Discovered by Finkel and Bentley (1974) and surveyed by Samet (1984), (1990a), they take several forms. We are given $d$ dimensional data $X_{1}, \ldots, X_{n}$. The tree is constructed as the k -d tree. In particular, $X_{1}$ becomes the root of the tree. It partitions $X_{2}, \ldots, X_{n}$ into $2^{d}$ (possibly empty) sets according to membership in one of the $2^{d}$ (hyper-) quadrants centered at $X_{1}$ (see Figure 20.15).


Figure 20.15. Quadtree and the induced partition of $\mathcal{R}^{2}$. The points on the right are shown in the position in space. The root is specially marked.
The partitioning process is repeated at the $2^{d}$ child nodes until a certain stopping rule is satisfied. In analogy with the k -d tree, we may define the chronological quadtree (only $k$ splits are allowed, defined by the first $k$ points $X_{1}, \ldots, X_{k}$ ), and the deep quadtree ( $k$ levels of splits are allowed). Other, more balanced versions may also be introduced. Classification is by majority vote over all ( $X_{i}, Y_{i}$ )'s-with $k<i \leq n$ in the chronological quadtree-that fall in the same region as $x$. Ties are broken in favor of class 0 . We will refer to this as the (chronological, deep) quadtree classifier.

Theorem 20.5. Whenever $X$ has a density, the chronological quadtree classifier is consistent $\left(\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}\right)$ provided that $k \rightarrow \infty$ and $k / n \rightarrow 0$.

Proof. Assume without loss of generality that all marginal distributions are uniform $[0,1]$. As the $X$-property holds, Theorem 6.1 applies. By Lemma 20.1, since we have $k\left(2^{d}-1\right)+1$ external regions,

$$
\mathbf{P}\{N(X) \leq M\} \leq \frac{(2 M+4)\left(k\left(2^{d}-1\right)+1\right)}{n-k} \rightarrow 0
$$

for all $M>0$, provided that $k / n \rightarrow 0$. Hence, we need only verify the condition
$\operatorname{diam}(A(X)) \rightarrow 0$ in probability. This is a bit easier than in the proof of Theorem 20.3 for the k -d tree and is thus left to the reader (Problem 20.18).

REMARK. Full-fledged random quadtrees with $n$ nodes have expected height $O(\log n)$ whenever $X$ has a density (see, e.g., Devroye and Laforest (1990)). With $k$ nodes, every region is thus reached in only $O(\log k)$ steps on the average. Furthermore, quadtrees enjoy the same monotone transformation invariance that we observed for k - d trees and median trees.

### 20.8 Best Possible Perpendicular Splits

For computational reasons, classification trees are most often produced by determining the splits recursively. At a given stage of the tree-growing algorithm, some criterion is used to determine which node of the tree should be split next, and where the split should be made. As these criteria typically use all the data, the resulting trees no longer have the $X$-property. In this section we examine perhaps the most natural criterion. In the following sections we introduce some alternative splitting criteria.

A binary classification tree can be obtained by associating with each node a splitting function $\phi(x)$ obtained in a top-down fashion from the data. For example, at the root, we may select the function

$$
\phi(x)=s x^{(i)}-\alpha,
$$

where $i$, the component cut, $\alpha \in \mathcal{R}$, the threshold, and $s \in\{-1,+1\}$, a polarization, are all dependent upon the data. The root then splits the data $D_{n}=$ $\left\{\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right\}$ into two sets, $D_{n}^{\prime}, D_{n}^{\prime \prime}$, with $D_{n}^{\prime} \cup D_{n}^{\prime \prime}=D_{n},\left|D_{n}^{\prime}\right|+\left|D_{n}^{\prime \prime}\right|=n$, such that

$$
\begin{aligned}
D_{n}^{\prime} & =\left\{(x, y) \in D_{n}: \phi(x) \geq 0\right\} \\
D_{n}^{\prime \prime} & =\left\{(x, y) \in D_{n}: \phi(x)<0\right\}
\end{aligned}
$$

A decision is made whether to split a node or not, and the procedure is applied recursively to the subtrees. Natural majority vote decisions are taken at the leaf level. All such trees will be called perpendicular splitting trees.

In Chapter 4, we introduced univariate Stoller splits, that is, splits that minimize the empirical error. This could be at the basis of a perpendicular splitting tree. One realizes immediately that the number of possibilities for stopping is endless. To name two, we could stop after $k$ splitting nodes have been defined, or we could make a tree with $k$ full levels of splits (so that all leaves are at distance $k$ from the root). We first show that for $d>1$, any such strategy is virtually doomed to fail. To make this case, we will argue on the basis of distribution functions only. For convenience, we consider a two-dimensional problem. Given a rectangle $R$ now
assigned to one class $y \in\{0,1\}$, we see that the current probability of error in $R$, before splitting is $\mathbf{P}\{X \in R, Y \neq y\}$. Let $R^{\prime}$ range over all rectangles of the form $R \cap((-\infty, \alpha] \times \mathcal{R}), R \cap([\alpha, \infty) \times \mathcal{R}), R \cap(\mathcal{R} \times(-\infty, \alpha])$, or $R \cap(R \times[\alpha, \infty))$, and let $R^{\prime \prime}=R-R^{\prime}$. Then after a split based upon ( $R^{\prime}, R^{\prime \prime}$ ), the probability of error over the rectangle $R$ is

$$
\mathbf{P}\left\{X \in R^{\prime}, Y=1\right\}+\mathbf{P}\left\{X \in R^{\prime \prime}, Y=0\right\}
$$

as we assign class 0 to $R^{\prime}$ and class 1 to $R^{\prime \prime}$. Call $\Delta(R)$ the decrease in probability of error if we minimize over all $R^{\prime}$. Compute $\Delta(R)$ for all leaf rectangles and then proceed to split that rectangle (or leaf) for which $\Delta(R)$ is maximal. The data-based rule based upon this would proceed similarly, if $\mathbf{P}\{A\}$ is replaced everywhere by the empirical estimate $(1 / n) \sum_{i=1}^{n} I_{\left\{\left\{X_{i}, Y_{i}\right) \in A\right\}}$, where $A$ is of the form $R^{\prime} \times\{1\}$, $R^{\prime \prime} \times\{0\}$, or $R \times\{1-y\}$, as the case may be.
Let us denote by $L_{0}, L_{1}, L_{2}, \ldots$ the sequence of the overall probabilities of error for the theoretically optimal sequence of cuts described above. Here we start with $R=\mathcal{R}^{2}$ and $y=0$, for example. For fixed $\epsilon>0$, we now construct a simple example in which

$$
L^{*}=0
$$

and

$$
L_{k} \downarrow \frac{1-\epsilon}{2} \quad \text { as } k \rightarrow \infty .
$$

Thus, applying the best split incrementally, even if we use the true probability of error as our criterion for splitting, is not advisable.
The example is very simple: $X$ has uniform distribution on $[0,1]^{2}$ with probability $\epsilon$ and on $[1,2]^{2}$ with probability $1-\epsilon$. Also, $Y$ is a deterministic function of $X$, so that $L^{*}=0$.


Figure 20.16. Repeated Stoller splits are not consistent in this two-dimensional example. Cuts will always be made in the leftmost square.

The way $Y$ depends on $X$ is shown in Figure 20.16: $Y=1$ if

$$
X \in[1,3 / 2]^{2} \cup[3 / 2,2]^{2} \cup\left(A_{2} \cup A_{4} \cup A_{6} \ldots\right) \times[0,1],
$$

where

$$
\begin{aligned}
A_{1} & =[0,1 / 4) \\
A_{2} & =[1 / 4,1 / 4+3 / 8) \\
A_{3} & =[1 / 4+3 / 8,1 / 4+3 / 8+3 / 16) \\
& \vdots \\
A_{k} & =\left[1 / 4+3 / 8+\cdots+3 / 2^{k}, 1 / 4+3 / 8+\cdots+3 / 2^{k+1}\right)
\end{aligned}
$$

and so forth. We verify easily that $\mathbf{P}\{Y=1\}=1 / 2$. Also, the error probability before any cut is made is $L_{0}=1 / 2$. The best split has $R^{\prime}=(-\infty, 1 / 4) \times \mathcal{R}$ so that $A_{1}$ is cut off. Therefore, $L_{1}=\epsilon / 4+(1-\epsilon) / 2$. We continue and split off $A_{2}$ and so forth, leading to the tree of Figure 20.17.


Figure 20.17. The tree obtained by repeated Stoller splits.

Verify that $L_{2}=\epsilon / 8+(1-\epsilon) / 2$ and in general that

$$
L_{k}=\frac{\epsilon}{2^{k+1}}+\frac{1-\epsilon}{2} \downarrow \frac{1-\epsilon}{2}
$$

as claimed.

### 20.9 Splitting Criteria Based on Impurity Functions

In 1984, Breiman, Friedman, Olshen, and Stone presented their CART program for constructing classification trees with perpendicular splits. One of the key ideas in their approach is the notion that trees should be constructed from the bottom up, by combining small subtrees. The starting point is a tree with $n+1$ leaf regions defined by a partition of the space based on the $n$ data points. Such a tree is much too large and is pruned by some methods that will not be explored here. When
constructing a starting tree, a certain splitting criterion is applied recursively. The criterion determines which rectangle should be split, and where the cut should be made. To keep the classifier invariant under monotone transformation of the coordinate axes, the criterion should only depend on the coordinatewise ranks of the points, and their labels. Typically the criterion is a function of the numbers of points labeled by 0 and 1 in the rectangles after the cut is made. One such class of criteria is described here.

Let $\alpha \in \mathcal{R}$, and let $i$ be a given coordinate $(1 \leq i \leq d)$. Let $R$ be a hyperrectangle to be cut. Define the following quantities for a split at $\alpha$, perpendicular to the $i$-th coordinate:

$$
\begin{aligned}
& \mathcal{X}_{i}(R, \alpha)=\left\{\left(X_{j}, Y_{j}\right): X_{j} \in R, X_{j}^{(i)} \leq \alpha\right\}, \\
& \mathcal{X}_{i}^{c}(R, \alpha)=\left\{\left(X_{j}, Y_{j}\right): X_{j} \in R, X_{j}^{(i)}>\alpha\right\}
\end{aligned}
$$

are the sets of pairs falling to the left and to the right of the cut, respectively.

$$
N_{i}(R, \alpha)=\left|\mathcal{X}_{i}(R, \alpha)\right|, \quad N_{i}^{\prime}(R, \alpha)=\left|\mathcal{X}_{i}^{c}(R, \alpha)\right|
$$

are the numbers of such pairs. Finally, the numbers of points with label 0 and label 1 in these sets are denoted, respectively, by

$$
\begin{aligned}
& N_{i, 0}(R, \alpha)=\left|\mathcal{X}_{i}(R, \alpha) \cap\left\{\left(X_{j}, Y_{j}\right): Y_{j}=0\right\}\right| \\
& N_{i, 1}(R, \alpha)=\left|\mathcal{X}_{i}(R, \alpha) \cap\left\{\left(X_{j}, Y_{j}\right): Y_{j}=1\right\}\right| \\
& N_{i, 0}^{\prime}(R, \alpha)=\left|\mathcal{X}_{i}^{c}(R, \alpha) \cap\left\{\left(X_{j}, Y_{j}\right): Y_{j}=0\right\}\right| \\
& N_{i, 1}^{\prime}(R, \alpha)=\left|\mathcal{X}_{i}^{c}(R, \alpha) \cap\left\{\left(X_{j}, Y_{j}\right): Y_{j}=1\right\}\right|
\end{aligned}
$$

Following Breiman, Friedman, Olshen, and Stone (1984), we define an impurity function for a possible split $(i, \alpha)$ by

$$
I_{i}(R, \alpha)=\psi\left(\frac{N_{i 0}}{N_{i 0}+N_{i 1}}, \frac{N_{i 1}}{N_{i 0}+N_{i 1}}\right) N_{i}+\psi\left(\frac{N_{i 0}^{\prime}}{N_{i 0}^{\prime}+N_{i 1}^{\prime}}, \frac{N_{i 1}^{\prime}}{N_{i 0}^{\prime}+N_{i 1}^{\prime}}\right) N_{i}^{\prime}
$$

where we dropped the $\operatorname{argument}(R, \alpha)$ throughout. Here $\psi$ is a nonnegative function with the following properties:
(1) $\psi\left(\frac{1}{2}, \frac{1}{2}\right) \geq \psi(p, 1-p)$ for any $p \in[0,1]$;
(2) $\psi(0,1)=\psi(1,0)=0$;
(3) $\psi(p, 1-p)$ increases in $p$ on $[0,1 / 2]$ and decreases in $p$ on $[1 / 2,1]$.

A rectangle $R$ is split at $\alpha$ along the $i$-th coordinate if $I_{i}(R, \alpha)$ is minimal. $I_{i}$ penalizes splits in which the subregions have about equal proportions from both classes. Examples of such functions $\psi$ include
(1) The entropy function $\psi(p, 1-p)=-p \log p-(1-p) \log (1-p)$ (Breiman et al. (1984, pp.25,103)).
(2) The Gini function $\psi(p, 1-p)=2 p(1-p)$, leading to the Gini index of diversity $I_{i}$. (Breiman et al. (1984, pp.38,103)).
(3) The probability of misclassification $\psi(p, 1-p)=\min (p, 1-p)$. In this case the splitting criterion leads to the empirical Stoller splits studied in the previous section.

We have two kinds of splits:
(A) The forced split: force a split along the $i$-th coordinate, but minimize $I_{i}(\alpha, R)$ over all $\alpha$ and rectangles $R$.
(B) The free split: choose the most advantageous coordinate for splitting, that is, minimize $I_{i}(\alpha, R)$ over all $\alpha, i$, and $R$.

Unfortunately, regardless of which kind of policy we choose, there are distributions for which no splitting based on an impurity function leads to a consistent classifier. To see this, note that the two-dimensional example of the previous section applies to all impurity functions. Assume that we had infinite sample size. Then $I_{i}(\alpha, R)$ would approach $a \psi(p, 1-p)+b \psi(q, 1-q)$, where $p$ is the probability content of $R^{\prime}$ (one of the rectangles obtained after the cut is made), $b$ is that of $R^{\prime \prime}$ (the other rectangle), $p=\mathbf{P}\left\{Y=1 \mid X \in R^{\prime}\right\}$ and $q=\mathbf{P}\left\{Y=1 \mid X \in R^{\prime \prime}\right\}$. If $X$ is uniformly distributed on the checkerboard shown in Figure 20.18, regardless where we try to cut, $p=q=1 / 2$, and every cut seems to be undesirable.

| class | class |
| :---: | :---: |
| 0 | 1 |
| clase | class |
| 1 | 0 |

Figure 20.18. No cut decreases the value of the impurity function.

This simple example may be made more interesting by mixing it with a distribution with much less weight in which $x^{(1)}$ - and $x^{(2)}$-direction splits are alternatingly encouraged all the time. Therefore, impurity functions should be avoided in their raw form for splitting. This may explain partially why in CART, the original tree is undesirable and must be pruned from the bottom up. See Problems 20.22 to 20.24 for more information. In the next section and in the last section of this chapter we propose other ways of growing trees with much more desirable properties.

The derivation shown above does not indicate that the empirical version will not work properly, but simple versions of it will certainly not. See Problem 20.22.

Remark. Malicious splitting. The impurity functions suggested above all avoid leaving the proportions of zeros and ones intact through splitting. They push towards more homogeneous regions. Assume now that we do the opposite. Through such splits, we can in fact create hyperplane classification trees that are globally poor, that is, that are such that every trimmed version of the tree is also a poor classifier. Such splitting methods must of course use the $Y_{i}$ 's. Our example shows
that any general consistency theorem for hyperplane classification trees must come with certain restrictions on the splitting process-the $X$ property is good; sometimes it is necessary to force cells to shrink; sometimes the position of the split is restricted by empirical error minimization or some other criterion.

The ham-sandwich theorem (see Edelsbrunner (1987)) states that given $2 n$ class0 points and $2 m$ class- 1 points in $\mathcal{R}^{d}, d>1$, there exists a hyperplane cut that leaves $n$ class -0 points and $m$ class- 1 points in each halfspace. So, assume that $X$ has a density and that $p=\mathbf{P}\{Y=1\}=1 / 2$. In a sample of size $n$, let $y$ be the majority class (ties are broken in favor of class 0 ).


Figure 20.19. Ham-sandwich cut: Each halfspace contains exactly half the points from each class.

Regardless of the sample make-up, if $y=0$, we may construct a hyperplane-tree classifier in which, during the construction, every node represents a region in which the majority vote would be 0 . This property has nothing to do with the distribution of $(X, Y)$, and therefore, for any trimmed version of the tree classifier,

$$
L_{n} \geq p I_{\{y=0\}}
$$

and $\mathbf{P}\left\{L_{n} \geq p\right\} \geq \mathbf{P}\{y=0\} \geq 1 / 2$ if $p=1 / 2$. Obviously, as we may take $L^{*}=0$, these classifiers are hopeless.

BibLIOGRAPHIC REMARKS. Empirical Stoller splits without forced rotation were recommended by Payne and Meisel (1977) and Rounds (1980), but their failure to be universally good was noted by Gordon and Olshen (1978). The last two authors recommended a splitting scheme that combined several ideas, but roughly speaking, they perform empirical Stoller splits with forced rotation through the coordinate axes (Gordon and Olshen (1978), (1980)). Other splitting criteria include the aID criterion of Morgan and Sonquist (1963), which is a predecessor of the Gini index of diversity used in CART (Breiman, Friedman, Olshen, and Stone (1984), see also Gelfand and Delp (1991), Guo and Gelfand (1992) Gelfand, Ravishankar, and Delp (1989), (1991), and Ciampi (1991)). Michel-Briand and Milhaud (1994) also observed the failure of multivariate classification trees based on the AID criterion.

The Shannon entropy or modifications of it are recommended by Talmon (1986), Sethi and Sarvarayudu (1982), Wang and Suen (1984), Goodman and Smyth (1988), and Chou (1991). Permutation statistics are used in Li and Dubes (1986), still without forced rotations through the coordinate axes. Quinlan (1993) has a more involved splitting criterion. A general discussion on tree splitting may be
found in the paper by Sethi (1991). A class of impurity functions is studied in Burshtein, Della Pietra, Kanevsky, and Nádas (1992). Among the pioneers of tree splitting (with perpendicular cuts) are Sebestyen (1962), and Henrichon and Fu (1969). For related work, we refer to Stoffel (1974), Sethi and Chatterjee (1977), Argentiero, Chin and Beaudet (1982), You and Fu (1976), Anderson and Fu (1979), Qing-Yun and Fu (1983), Hartmann, Varshney, Mehrotra, and Gerberich (1982), and Casey and Nagy (1984). References on nonperpendicular splitting methods are given below in the section on BSP trees.

### 20.10 A Consistent Splitting Criterion

There is no reason for pessimism after the previous sections. Rest assured, there are several consistent splitting strategies that are fully automatic and depend only upon the populations of the regions. In this section we provide a solution for the simple case when $X$ is univariate and nonatomic. It is possible to generalize the method for $d>1$ if we force cuts to alternate directions. We omit here the detailed analysis for multidimensional cases for two reasons. First, it is significantly heavier than for $d=1$. Secondly, in the last section of this chapter we introduce a fully automatic way of building up consistent trees, that is, without forcing the directions of the splits.

To a partition $A_{1}, \ldots, A_{N}$ of $\mathcal{R}$, we assign the quantity

$$
Q=\sum_{i=1}^{N} N_{0}\left(A_{i}\right) N_{1}\left(A_{i}\right),
$$

where

$$
N_{0}(A)=\sum_{j=1}^{n} I_{\left\{X_{j} \in A, Y_{j}=0\right\}}, \quad N_{1}(A)=\sum_{j=1}^{n} I_{\left\{X_{j} \in A, Y_{j}=1\right\}}
$$

are the respective numbers of points labeled with 0 and 1 falling in the region $A$. The tree-growing algorithm starts with the trivial partition $\{\mathcal{R}\}$, and at each step it makes a cut that yields the minimal value of $Q$. It proceeds recursively until the improvement in the value of $Q$ falls below a threshold $\Delta_{n}$.
Remark. Notice that this criterion always splits a cell that has many points from both classes (see the proof of the theorem below). Thus, it avoids the anomalies of impurity-function criteria described in the previous section. On the other hand, it does not necessarily split large cells, if they are almost homogeneous. For a comparison, recall that the Gini criterion minimizes the quantity

$$
Q^{\prime}=\sum_{i=1}^{N} \frac{N_{0}\left(A_{i}\right) N_{1}\left(A_{i}\right)}{N_{0}\left(A_{i}\right)+N_{1}\left(A_{i}\right)},
$$

thus favoring cutting cells with very few points. We realize that the criterion $Q$ introduced here is just one of many with similarly good properties, and albeit probably imperfect, it is certainly one of the simplest.

Theorem 20.6. Let $X$ have a nonatomic distribution on the real line, and consider the tree classifier obtained by the algorithm described above. If the threshold satisfies

$$
\frac{\Delta_{n}}{n} \rightarrow \infty \text { and } \frac{\Delta_{n}}{n^{2}} \rightarrow 0
$$

then the classification rule is strongly consistent.
Proof. There are two key properties of the algorithm that we exploit:
Property 1. If $\min \left(N_{0}(A), N_{1}(A)\right)>\sqrt{2 \Delta_{n}}$ for a cell $A$, then $A$ gets cut by the algorithm. To see this, observe that (dropping the argument $A$ from the notation) if $N_{0} \geq N_{1}$, and we cut $A$ so that the number of 0 -labeled points in the two child regions differ by at most one, then the contribution of these two new regions to $Q$ is

$$
\left\lceil N_{0} / 2\right\rceil N_{1}^{\prime}+\left\lfloor N_{0} / 2\right\rfloor N_{1}^{\prime \prime} \leq\left\lceil N_{0} N_{1} / 2\right\rceil,
$$

where $N_{1}^{\prime}$ and $N_{1}^{\prime \prime}$ are the numbers of class-1 points in the two child regions. Thus, the decrease of $Q$ if $A$ is split is at least $\left\lfloor N_{0} N_{1} / 2\right\rfloor$. If $\min \left(N_{0}, N_{1}\right)>\sqrt{2 \Delta_{n}}$, then $\Delta_{n} \leq\left\lfloor N_{0} N_{1} / 2\right\rfloor$, and a cut is made.
Property 2. No leaf node has less than $\Delta_{n} / n$ points. Assume that after a region is cut, in one of the child regions, the total number of points is $N_{0}^{\prime}+N_{1}^{\prime} \leq k$. Then the improvement in $Q$ caused by the split is bounded by

$$
N_{0} N_{1}-\left(N_{0}^{\prime} N_{1}^{\prime}+N_{0}^{\prime \prime} N_{1}^{\prime \prime}\right) \leq N_{0} N_{1}-\left(N_{0}-k\right)\left(N_{1}-k\right) \leq\left(N_{0}+N_{1}\right) k \leq n k .
$$

Therefore, if $k<\Delta_{n} / n$, then the improvement is smaller than $\Delta_{n}$. Thus, no cut is made that leaves behind a child region with fewer than $\Delta_{n} / n$ points.
It follows from Property 1 that if a leaf region has more than $4 \sqrt{2 \Delta_{n}}$ points, then since the class in minority has less than $\sqrt{2 \Delta_{n}}$ points in it, it may be cut into intervals containing between $2 \sqrt{2 \Delta_{n}}$ and $4 \sqrt{2 \Delta_{n}}$ points without altering the decision, since the majority vote within each region remains the same.

Summarizing, we see that the tree classifier is equivalent to a classifier that partitions the real line into intervals, each containing at least $\Delta_{n} / n$, and at most $4 \sqrt{2 \Delta_{n}}$ data points. Thus, in this partition, each interval has a number of points growing to infinity as $o(n)$. We emphasize that the number of points in a region of the studied tree classifier may be large, but such regions are almost homogeneous, and therefore the classifier is equivalent to another classifier which has $o(n)$ points in each region. Consistency of such partitioning classifiers is proved in the next chapter-see Theorem 21.3.

### 20.11 BSP Trees

Binary space partition trees (or BSP trees) partition $\mathcal{R}^{d}$ by hyperplanes. Trees of this nature have evolved in the computer graphics literature via the work of Fuchs,

Kedem, and Naylor (1980) and Fuchs, Abram, and Grant (1983) (see also Samet (1990b), Kaplan (1985), and Sung and Shirley (1992)). In discrimination they are at the same time generalizations of linear discriminants, of histograms, and of binary tree classifiers. BSP trees were recommended for use in discrimination by Henrichon and Fu (1969), Mizoguchi, Kizawa, and Shimura (1977) and Friedman (1977). Further studies include Sklansky and Michelotti (1980), Argentiero, Chin, and Beaudet (1982), Qing-Yun and Fu (1983), Breiman, Friedman, Olshen, and Stone (1984), Loh and Vanichsetakul (1988), and Park and Sklansky (1990).

There are numerous ways of constructing BSP trees. Most methods of course use the $Y$-values to determine good splits. Nevertheless, we should mention first simple splits with the $X$-property, if only to better understand the BSP trees.


Figure 20.20. A raw BSP tree and its induced partition in the plane. Every region is split by a line determined by the two data points with smallest index in the region.

We call the raw BSP tree the tree obtained by letting $X_{1}, \ldots, X_{d}$ determine the first hyperplane. The $d$ data points remain associated with the root, and the others $\left(X_{d+1}, \ldots, X_{k}\right)$ are sent down to the subtrees, where the process is repeated as far as possible. The remaining points $X_{k+1}, \ldots, X_{n}$ are used in a majority vote in the external regions. Note that the number of regions is not more than $k / d$. Thus, by Lemma 20.1, if $k / n \rightarrow 0$, we have $N(X) \rightarrow \infty$ in probability. Combining this with Problem 20.25, we have our first result:

Theorem 20.7. The natural binary tree classifier based upon a raw BSP tree with $k \rightarrow \infty$ and $k / n \rightarrow 0$ is consistent whenever $X$ has a density.

Hyperplanes may also be selected by optimization of a criterion. Typically, this would involve separating the classes in some way. All that was said for perpendicular splitting remains valid here. It is worthwhile recalling therefore that there
are distributions for which the best empirical Stoller split does not improve the probability of error. Take for example the uniform distribution in the unit ball of $\mathcal{R}^{d}$ in which

$$
Y= \begin{cases}1 & \text { if }\|X\| \geq 1 / 2 \\ 0 & \text { if }\|X\|<1 / 2\end{cases}
$$

Figure 20.21. No single split improves on the error probability for this distribution.

Here, no linear split would be helpful as the 1's would always hold a strong majority. Minimizing other impurity functions such as the Gini criterion may be helpful, however (Problem 20.27).
Bibliographic remarks. Hyperplane splits may be generalized to include quadratic splits (Henrichon and Fu (1969)). For example, Mui and Fu (1980) suggest taking $d^{\prime}<d$ and forming quadratic classifiers as in normal discrimination (see Chapter 4) based upon vectors in $\mathcal{R}^{d^{\prime}}$. The cuts are thus perpendicular to $d-d^{\prime}$ axes but quadratic in the subspace $\mathcal{R}^{d^{\prime}}$. Lin and Fu (1983) employ the Bhattacharyya distance or 2-means clustering for determining splits. As a novelty, within each leaf region, the decision is not by majority vote, but rather by a slightly more sophisticated rule such as the $k$-nN rule or linear discrimination. Here, no optimization is required along the way. Loh and Vanichsetakul (1988) allow linear splits but use $F$ ratios to select desirable hyperplanes.

### 20.12 Primitive Selection

There are two reasons for optimizing a tree configuration. First of all, it just does not make sense to ignore the class labels when constructing a tree classifier, so the $Y_{i}$ 's must be used to help in the selection of a best tree. Secondly, some trees may not be consistent (or provably consistent), yet when optimized over a family of trees, consistency drops out. We take the following example: let $\mathcal{G}_{k}$ be a class of binary tree classifiers with the $X$-property, with the space partitioned into $k+1$ regions determined by $X_{1}, \ldots, X_{k}$ only. Examples include the chronological k-d tree and some kinds of BSP trees. We estimate the error for $g \in \mathcal{G}_{k}$ by

$$
\widehat{L}_{n}(g)=\frac{1}{n-k} \sum_{i=k+1}^{n} I_{\left\{g\left(X_{i}\right) \neq Y_{i}\right\}},
$$

realizing the danger of using the same data that were used to obtain majority votes to estimate the error. An optimistic bias will be introduced. (For more on such error estimates, see Chapter 23.) Let $g_{n}^{*}$ be the classifier (or one of the classifiers) in $\mathcal{G}_{k}$ for which $\widehat{L}_{n}$ is minimum. Assume that $\left|\mathcal{G}_{k}\right|<\infty$-for example, $\mathcal{G}_{k}$ could
be all $k$ ! chronological k-d trees obtained by permuting $X_{1}, \ldots, X_{k}$. We call $g_{n}^{*}$ then the permutation-optimized chronological k -d classifier. When $k=\lfloor\sqrt{\log n}\rfloor$, one can verify that $k!=o\left(n^{\epsilon}\right)$ for any $\epsilon>0$, so that the computational burden-at least on paper-is not out of sight. We assume that $\mathcal{G}_{k}$ has a consistent classifier sequence, that is, as $n \rightarrow \infty, k$ usually grows unbounded, and $\mathbf{E}\left\{L_{n}\left(g_{k}\right)\right\} \rightarrow L^{*}$ for a sequence $g_{k} \in \mathcal{G}_{k}$.

Example. Among the chronological k-d trees modulo permutations, the first one (i.e., the one corresponding to the identity permutation) was shown to be consistent in Theorem 20.3, if $X$ has a density, $k \rightarrow \infty$, and $k / n \rightarrow 0$.

EXAMPLE. Let $\mathcal{G}_{k}$ be the class of BSP trees in which we take as possible hyperplanes for splitting the root:
(1) the hyperplane through $X_{1}, \ldots, X_{d}$;
(2) the $d$ hyperplanes through $X_{1}, \ldots, X_{d-1}$ that are parallel to one axis;
(3) the $\binom{d}{2}$ hyperplanes through $X_{1}, \ldots, X_{d-2}$ that are parallel to two axes;
(d) the $\binom{d}{d-1}=d$ hyperplanes through $X_{1}$ that are parallel to $d-1$ axes.

Thus, conservatively estimated, $\left|\mathcal{G}_{k}\right| \leq k!2^{d}$ because there are at most $2^{d}$ possible choices at each node and there are $k!$ permutations of $X_{1}, \ldots, X_{k}$. Granted, the number of external regions is very variable, but it remains bounded by $k+1$ in any case. As $\mathcal{G}_{k}$ contains the chronological $\mathrm{k}-\mathrm{d}$ tree, it has a consistent sequence of classifiers when $k \rightarrow \infty$ and $n /(k \log k) \rightarrow \infty$. $\square$

Theorem 20.8. Let $g_{n}^{*}=\arg \min _{g \in \mathcal{G}_{k}} \widehat{L}_{n}(g)$. Then, if $\mathcal{G}_{k}$ has a consistent sequence of classifiers, if the number of regions in the partitions for all $g \in \mathcal{G}_{k}$ are at most $k+1$, and if $k \rightarrow \infty$ and $n / \log \left|\mathcal{G}_{k}\right| \rightarrow \infty$, then $\mathbf{E}\left\{L_{n}\left(g_{n}^{*}\right)\right\} \rightarrow L^{*}$, where $L_{n}\left(g_{n}^{*}\right)$ is the conditional probability of error of $g_{n}^{*}$.

In the examples cited above, we must take $k \rightarrow \infty, n / k \rightarrow \infty$. Furthermore, $\log \left|\mathcal{G}_{k}\right|=O(\log k!)=O(k \log k)$ in both cases. Thus, $g_{n}^{*}$ is consistent whenever $X$ has a density, $k \rightarrow \infty$ and $k=o(n / \log n)$. This is a simple way of constructing a basically universally consistent BSP tree.

Proof. Let $g^{+}=\arg \min _{g \in \mathcal{G}_{k}} L_{n}(g)$.

$$
\begin{aligned}
L_{n}\left(g_{n}^{*}\right)-L^{*}= & L_{n}\left(g_{n}^{*}\right)-\widehat{L}_{n}\left(g_{n}^{*}\right) \\
& +\widehat{L}_{n}\left(g_{n}^{*}\right)-\widehat{L}_{n}\left(g^{+}\right) \\
& +\widehat{L}_{n}\left(g^{+}\right)-L_{n}\left(g^{+}\right) \\
& +L_{n}\left(g^{+}\right)-L^{*}
\end{aligned}
$$

Clearly,

$$
\begin{aligned}
L_{n}\left(g_{n}^{*}\right)-L^{*} & \leq L_{n}\left(g^{+}\right)-L^{*}+2 \max _{g \in \mathcal{G}_{k}}\left|L_{n}(g)-\widehat{L}_{n}(g)\right| \\
& \stackrel{\text { def }}{=} I+I I .
\end{aligned}
$$

Obviously, $I \rightarrow 0$ in the mean by our assumption, and for $\epsilon>0$,

$$
\mathbf{P}\left\{2 \max _{g \in \mathcal{G}_{k}}\left|L_{n}(g)-\widehat{L}_{n}(g)\right|>\epsilon\right\} \leq\left|\mathcal{G}_{k}\right| \max _{g \in \mathcal{G}_{k}} \mathbf{P}\left\{\left|L_{n}(g)-\widehat{L}_{n}(g)\right|>\frac{\epsilon}{2}\right\}
$$

Next we bound the probabilities on the right-hand side. Let $p_{i 0}, p_{i 1}$ denote $\mathbf{P}\{X \in$ region $i, Y=0\}$ and $\mathbf{P}\{X \in$ region $i, Y=1\}$ respectively, with regions determined by $g, 1 \leq i \leq k+1$. Let

$$
N_{i 0}=\sum_{j=k+1}^{n} I_{\left\{X_{j} \in \text { region } i, Y_{j}=0\right\}} \quad \text { and } \quad N_{i 1}=\sum_{j=k+1}^{n} I_{\left\{X_{j} \in \text { region } i, Y_{j}=1\right\}} .
$$

Then

$$
L_{n}(g)=\sum_{i=1}^{k+1} p_{i 1} I_{\left\{N_{i 0} \geq N_{i 1}\right\}}+\sum_{i=1}^{k+1} p_{i 0} I_{\left\{N_{i 0}<N_{i 1}\right\}}
$$

and

$$
\widehat{L}_{n}(g)=\sum_{i=1}^{k+1} \frac{N_{i 1}}{n-k} I_{\left\{N_{i 0} \geq N_{i 1}\right\}}+\sum_{i=1}^{k+1} \frac{N_{i 0}}{n-k} I_{\left\{N_{i 0}<N_{i 1}\right\}}
$$

Thus,

$$
\left|L_{n}(g)-\widehat{L}_{n}(g)\right| \leq \sum_{i=1}^{k+1}\left|p_{i 1}-\frac{N_{i 1}}{n-k}\right|+\sum_{i=1}^{k+1}\left|p_{i 0}-\frac{N_{i 0}}{n-k}\right| .
$$

Introduce the notation $Z$ for the random variable on the right-hand side of the above inequality. By the Cauchy-Schwarz inequality,

$$
\begin{aligned}
\mathbf{E} & \left.Z \mid X_{1}, \ldots, X_{k}\right\} \\
\leq & \sum_{i=1}^{k+1} \sqrt{\mathbf{E}\left\{\left.\left(p_{i 1}-\frac{N_{i 1}}{n-k}\right)^{2} \right\rvert\, X_{1}, \ldots, X_{k}\right\}} \\
& +\sum_{i=1}^{k+1} \sqrt{\mathbf{E}\left\{\left.\left(p_{i 0}-\frac{N_{i 0}}{n-k}\right)^{2} \right\rvert\, X_{1}, \ldots, X_{k}\right\}}
\end{aligned}
$$

(as given $X_{1}, \ldots, X_{k}, N_{i 1}$ is binomial $\left(n-k, p_{i 1}\right)$ )

$$
\leq \sum_{i=1}^{k+1}\left(\sqrt{\frac{p_{i 1}}{n-k}}+\sqrt{\frac{p_{i 0}}{n-k}}\right)
$$

$$
\leq \frac{2 k+2}{\sqrt{n-k}} \sqrt{\frac{1}{2 k+2} \sum_{i=1}^{k+1}\left(p_{i 1}+p_{i 0}\right)}
$$

(by another use of the Cauchy-Schwarz inequality)

$$
=\sqrt{\frac{2 k+2}{n-k}} .
$$

Thus, $\mathbf{E}\left\{Z \mid X_{1}, \ldots, X_{k}\right\} \rightarrow 0$ as $k / n \rightarrow 0$. Note that

$$
\begin{aligned}
\mathbf{P}\{Z> & \left.\epsilon \mid X_{1}, \ldots, X_{k}\right\} \\
\leq & \mathbf{P}\left\{\left.Z-\mathbf{E}\left\{Z \mid X_{1}, \ldots, X_{k}\right\}>\frac{\epsilon}{2} \right\rvert\, X_{1}, \ldots, X_{k}\right\} \\
& \left(\text { if } \mathbf{E}\left\{Z \mid X_{1}, \ldots, X_{k}\right\}<\epsilon / 2, \text { which happens when } \sqrt{\frac{2 k+2}{n-k}} \leq \frac{\epsilon}{2}\right) \\
\leq & \exp \left(-\left(\frac{\epsilon}{2}\right)^{2} / 2(n-k) \frac{4}{(n-k)^{2}}\right)
\end{aligned}
$$

(by McDiarmid's inequality, as changing the value of an $X_{j}, j \geq k$, changes $Z$ by at most $2 /(n-k)$; see Theorem 9.2)

$$
=\exp \left(-\frac{(n-k) \epsilon^{2}}{32}\right)
$$

Thus, taking expected values, we see that

$$
\mathbf{P}\{I I>\epsilon\} \leq\left|\mathcal{G}_{k}\right| e^{-(n-k) \epsilon^{2} / 64}
$$

if $\sqrt{(2 k+2) /(n-k)} \leq \epsilon / 2$. This tends to 0 for all $\epsilon>0$ if $k / n \rightarrow 0$ and $n / \log \left|\mathcal{G}_{k}\right| \rightarrow \infty . \square$

### 20.13 Constructing Consistent Tree Classifiers

Thus far, we have taken you through a forest of beautiful trees and we have shown you a few tricks of the trade. When you read (or write) a research paper on tree classifiers, and try to directly apply a consistency theorem, you will get frustrated however-most real-life tree classifiers use the data in intricate ways to suit a certain application. It really helps to have a few truly general results that have universal impact. In this section we will point you to three different places in the book where you may find useful results in this respect. First of all, there is a consistency theorem-Theorem 21.2-that applies to rules that partition the space and decide by majority vote. The partition is arbitrary and may thus be generated by using some or all of $X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}$. If a rule satisfies the two (weak) conditions of Theorem 21.2, it must be universally consistent. To put it differently, even the worst rule within the boundaries of the theorem's conditions must perform well asymptotically.

Second, we will briefly discuss the design of tree classifiers obtained by minimizing the empirical error estimate ( $\widehat{L}_{n}$ ) over possibly infinite classes of classifiers. Such classifiers, however hard to find by an algorithm, have asymptotic properties that are related to the VC dimension of the class of rules. Consistency follows almost without work if one can calculate or bound the vc dimension appropriately.

While Chapters 13 and 21 deal in more detail with the vc dimension, it is necessary to give a few examples here.

Third, we point the reader to Chapter 22 on data splitting, where the previous approach is applied to the minimization of the holdout estimate, obtained by trees based upon part of the sample, and using another part to select the best tree in the bunch. Here too the VC dimension plays a crucial role.

Theorem 21.2 allows space partitions that depend quite arbitrarily on all the data and extends earlier universally applicable results of Gordon and Olshen (1978), (1980), (1984), and Breiman, Friedman, Olshen, and Stone (1984). A particularly useful format is given in Theorem 21.8. If the partition is by recursive hyperplane splits as in BSP trees and the number of splits is at most $m_{n}$, if $m_{n} \log n / n \rightarrow 0$, and if

$$
\operatorname{diam}\left(A(X) \cap S_{B}\right) \rightarrow 0
$$

with probability one for all $S_{B}$ (where $S_{B}$ is the ball of radius $B$ centered at the origin), then the classification rule is strongly consistent. The last condition forces a randomly picked region in the partition to be small. However, $m_{n} \log n / n \rightarrow 0$ guarantees that no devilish partition can be inconsistent. The latter condition is certainly satisfied if each region contains at least $k_{n}$ points, where $k_{n} / \log n \rightarrow \infty$.

Next we take a look at full-fledged minimization of $\widehat{L}_{n}$, the empirical error, over certain classes of tree classifiers. Here we are not concerned with the (often unacceptable) computational effort. For example, let $\mathcal{G}_{k}$ be the class of all binary tree classifiers based upon a tree consisting of $k$ internal nodes, each representing a hyperplane cut (as in the BSP tree), and all possible $2^{k+1}$ labelings of the $k+1$ leaf regions. Pick such a classifier for which

$$
\widehat{L}_{n}(g)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{g\left(X_{i}\right) \neq Y_{i}\right\}}
$$

is minimal. Observe that the chosen tree is always natural; that is, it takes majority votes over the leaf regions. Thus, the minimization is equivalent to the minimization of the resubstitution error estimate (defined in Chapter 23) over the corresponding class of natural tree classifiers. We say that a sequence $\left\{\mathcal{G}_{k}\right\}$ of classes is rich if we can find a sequence $g_{k} \in \mathcal{G}_{k}$ such that $L\left(g_{k}\right) \rightarrow L^{*}$. For hyperplanes, this is the case if $k \rightarrow \infty$ as $n \rightarrow \infty$-just make the hyperplane cuts form a regular histogram grid and recall Theorem 9.4. Let $\mathcal{S}\left(\mathcal{G}_{k}, n\right)$ be the shatter coefficient of $\mathcal{G}_{k}$ (for a definition, see Chapter 12, Definition 12.3). For example, for the hyperplane family,

$$
\mathcal{S}\left(\mathcal{G}_{k}, n\right) \leq n^{k(d+1)}
$$

Then by Corollary 12.1 we have $\mathbf{E}\left\{L_{n}\left(g_{n}^{*}\right)\right\} \rightarrow L^{*}$ for the selected classifier $g_{n}^{*}$ when $\mathcal{G}_{k}$ is rich $\left(k \rightarrow \infty\right.$ here) and $\log \mathcal{S}\left(\mathcal{G}_{k}, n\right)=o(n)$, that is,

$$
k=o\left(\frac{n}{\log n}\right) .
$$

Observe that no conditions are placed on the distribution of $(X, Y)$ here! Consistency follows from basic notions-one combinatorial to keep us from overfitting, and one approximation-theoretical (the richness).

The above result remains valid under the same conditions on $k$ in the following classes:
(1) All trees based upon $k$ internal nodes each representing a perpendicular split. (Note: $\mathcal{S}\left(\mathcal{G}_{k}, n\right) \leq((n+1) d)^{k}$.)
(2) All trees based upon $k$ internal nodes, each representing a quadtree split. $\left(\right.$ Note: $\mathcal{S}\left(\mathcal{G}_{k}, n\right) \leq\left(n^{d}+1\right)^{k}$.)

### 20.14 A Greedy Classifier

In this section, we define simply a binary tree classifier that is grown via optimization of a simple criterion. It has the remarkable property that it does not require a forced rotation through the coordinate axes or special safeguards against small or large regions or the like. It remains entirely parameter-free (nothing is picked by the user), is monotone transformation invariant, and fully automatic. We show that in $\mathcal{R}^{d}$ it is always consistent. It serves as a prototype for teaching about such rules and should not be considered as more than that. For fully practical methods, we believe, one will have to tinker with the approach.

The space is partitioned into rectangles as shown below:


Figure 20.22. A tree based on partitioning the plane into rectangles. The right subtree of each internal node belongs to the inside of a rectangle, and the left subtree belongs to the complement of the same rectangle ( $i^{c}$ denotes the complement of $i$ ). Rectangles are not allowed to overlap.

A hyperrectangle defines a split in a natural way. The theory presented here applies for many other types of cuts. These will be discussed after the main consistency theorem is stated.

A partition is denoted by $\mathcal{P}$, and a decision on a set $A \in \mathcal{P}$ is by majority vote. We write $g_{\mathcal{P}}$ for such a rule:

$$
g_{\mathcal{P}}(x)= \begin{cases}1 & \text { if } \sum_{i: X_{i} \in A} Y_{i}>\sum_{i: X_{i} \in A}\left(1-Y_{i}\right), x \in A \\ 0 & \text { otherwise } .\end{cases}
$$

Given a partition $\mathcal{P}$, a legal rectangle $A$ is one for which $A \cap B=\emptyset$ or $A \subseteq B$ for all sets $B \in \mathcal{P}$. If we refine $\mathcal{P}$ by adding a legal rectangle $T$ somewhere, then we obtain the partition $\mathcal{T}$. The decision $g_{\mathcal{T}}$ agrees with $g_{\mathcal{P}}$ except on the set $B \in \mathcal{P}$ that contains $T$.

We introduce the convenient notation

$$
\begin{aligned}
v_{j}(A) & =\mathbf{P}\{X \in A, Y=j\}, j \in\{0,1\}, \\
v_{j, n}(A) & =\frac{1}{n} \sum_{i=1}^{n} I_{\left\{X_{i} \in A, Y_{i}=j\right\}}, \quad j \in\{0,1\}
\end{aligned}
$$

An estimate of the quality of $g_{\mathcal{P}}$ is

$$
\widehat{L}_{n}(\mathcal{P}) \stackrel{\text { def }}{=} \sum_{R \in \mathcal{P}} \widehat{L}_{n}(R)
$$

where

$$
\begin{aligned}
\widehat{L}_{n}(R) & =\frac{1}{n} \sum_{i=1}^{n} I_{\left\{X_{i} \in R, g_{\mathcal{P}}\left(X_{i}\right) \neq Y_{i}\right\}} \\
& =\min \left(v_{0, n}(R), v_{1, n}(R)\right) .
\end{aligned}
$$

Here we use two different arguments for $\widehat{L}_{n}(R$ and $\mathcal{P})$, but the distinction should be clear. We may similarly define $\widehat{L}_{n}(\mathcal{T})$. Given a partition $\mathcal{P}$, the greedy classifier selects that legal rectangle $T$ for which $\widehat{L}_{n}(\mathcal{T})$ is minimal (with any appropriate policy for breaking ties). Let $R$ be the set of $\mathcal{P}$ containing $T$. Then the greedy classifier picks that $T$ for which

$$
\widehat{L}_{n}(T)+\widehat{L}_{n}(R-T)-\widehat{L}_{n}(R)
$$

is minimal. Starting with the trivial partition $\mathcal{P}_{0}=\left\{\mathcal{R}^{d}\right\}$, we repeat the previous step $k$ times, leading thus to $k+1$ regions. The sequence of partitions is denoted by $\mathcal{P}_{0}, \mathcal{P}_{1}, \ldots, \mathcal{P}_{k}$.

We put no safeguards in place-the rectangles are not forced to shrink. And in fact, it is easy to construct examples in which most rectangles do not shrink. The main result of the section, and indeed of this chapter, is that the obtained classifier is consistent:

Theorem 20.9. For the greedy classifier with $k \rightarrow \infty$ and $k=o(\sqrt[3]{n / \log n})$, assuming that $X$ has nonatomic marginals, we have $L_{n} \rightarrow L^{*}$ with probability one.

Remark. We note that with techniques presented in the next chapter, it is possible to improve the second condition on $k$ to $k=o(\sqrt{n / \log n})$ (see Problem 20.31).

Before proving the theorem, we mention that the same argument may be used to establish consistency of greedily grown trees with many other types of cuts. We have seen in Section 20.8 that repeated Stoller splits do not result in good classifiers. The reason is that optimization is over a collection of sets (halfspaces) that is not guaranteed to improve matters-witness the examples provided in previous sections. A good cutting method is one that includes somehow many (but not too many) small sets. For example, let us split at the root making $d+1$ hyperplane cuts at once, that is, by finding the $d+1$ cuts that together produce the largest decrease in the empirical error probability. Then repeat this step recursively in each region $k$ times. The procedure is consistent under the same conditions on $k$ as in Theorem 6.1, whenever $X$ has a density (see Problem 20.30). The $d+1$ hyperplane cuts may be considered as an elementary cut which is repeated in a greedy manner. In Figures 20.23 to 20.25 we show a few elementary cuts that may be repeated greedily for a consistent classifier. The straightforward proofs of consistency are left to the reader in Problem 20.30.

Proof of Theorem 20.9. We restrict ourselves to $\mathcal{R}^{2}$, but the proof remains similar in $\mathcal{R}^{d}$ (Problem 20.29). The notation $\widehat{L}_{n}(\cdot)$ was introduced above, where the argument is allowed to be a partition or a set in a partition. We similarly define

$$
\begin{aligned}
L(R) & =\min _{y \in\{0,1\}} \mathbf{P}\{X \in R, Y \neq y\} \\
& =\min \left(v_{0}(R), \nu_{1}(R)\right), \\
L(\mathcal{P}) & =\sum_{R \in \mathcal{P}} L(R)
\end{aligned}
$$



Figure 20.23. An elementary cut here is composed of $d+1$ hyperplane cuts. They are jointly optimized.

Figure 20.24. $2 d$ rectangular cuts determine an elementary cut. All $2 d$ cuts have arbitrary directions; there are no forced directions.


Figure 20.25. Simplex cuts. A cut is determined by a polyhedron with $d+1$ vertices. The simplices are not allowed to overlap, just as for rectangular cuts in the greedy classifier. Three consecutive simplex cuts in $\mathcal{R}^{2}$ are shown here.


Figure 20.26. At every iteration we refine the partition by selecting that rectangle $R$ in the partition and that $3 \times$ $3 \times \cdots \times 3$ rectangular grid cut of $R$ for which the empirical error is minimal. Two consecutive rectangular grid cuts are shown.

For example, let $\mathcal{G}_{l}$ denote a partition of $\mathcal{R}^{2}$ into a rectangular $l \times l$ grid.


Figure 20.27. $\mathcal{G}_{l}:$ An $l \times l$ grid (with $l=7$ here $)$.

It is clear that for all $\epsilon>0$, there exists an $l=l(\epsilon)$ and an $l \times l$ grid $\mathcal{G}_{l}$ such that

$$
L\left(\mathcal{G}_{l}\right) \leq L^{*}+\epsilon
$$

If $\mathcal{Q}$ is another finer partition into rectangles (i.e., each set of $\mathcal{Q}$ is a rectangle and intersects at most one rectangle of $\mathcal{G}_{l}$ ), then necessarily

$$
L(\mathcal{Q}) \leq L\left(\mathcal{G}_{l}\right) \leq L^{*}+\epsilon
$$

We will call $\mathcal{Q}$ a refinement of $\mathcal{G}_{l}$. The next lemma is a key property of partitioning classifiers. In our eyes, it is the main technical property of this entire chapter. We say that the partition $\mathcal{T}$ is an extension of $\mathcal{P}$ by a set $Q$-where $Q \subseteq R \in \mathcal{P}$-if $\mathcal{T}$ contains all cells of $\mathcal{P}$ other than $R$, plus $Q$ and $R-Q$.

Lemma 20.2. Let $\mathcal{G}_{l}$ be a finite partition with $L\left(\mathcal{G}_{l}\right) \leq L^{*}+\epsilon$. Let $\mathcal{P}$ be a finite partition of $\mathcal{R}^{d}$, and let $\mathcal{Q}$ be a refinement of both $\mathcal{P}$ and $\mathcal{G}_{l}$. Then there exists a set $Q \in \mathcal{Q}$ (note: $Q$ is contained in one set of $\mathcal{P}$ only) and an extension of $\mathcal{P}$ by $Q$ to $\mathcal{T}_{Q}$ such that, if $L(\mathcal{P}) \geq L^{*}+\epsilon$,

$$
L\left(\mathcal{T}_{Q}\right)-\left(L^{*}+\epsilon\right) \leq\left(1-\frac{1}{|\mathcal{Q}|}\right)\left(L(\mathcal{P})-\left(L^{*}+\epsilon\right)\right)
$$

Proof. First fix $R \in \mathcal{P}$ and let $Q_{1}, \ldots, Q_{N}$ be the sets of $\mathcal{Q}$ contained in $R$. Define

$$
\begin{gathered}
p_{i}=v_{0}\left(Q_{i}\right), \quad q_{i}=v_{1}\left(Q_{i}\right), \quad L\left(Q_{i}\right)=\min \left(p_{i}, q_{i}\right), \\
p=\sum_{i=1}^{N} p_{i}, \quad q=\sum_{i=1}^{N} q_{i}, \quad L(R)=\min (p, q), \\
L\left(R-Q_{i}\right)=\min \left(p-p_{i}, q-q_{i}\right)
\end{gathered}
$$

First we show that there exists an integer $i$ such that

$$
L(R)-L\left(Q_{i}\right)-L\left(R-Q_{i}\right) \geq \frac{L(R)-\sum_{i=1}^{N} L\left(Q_{i}\right)}{N}
$$

or equivalently,

$$
\max _{1 \leq i \leq N} \Delta_{i} \geq \frac{\min (p, q)-\sum_{i=1}^{N} \min \left(p_{i}, q_{i}\right)}{N}
$$

where $\Delta_{i}=\min (p, q)-\min \left(p_{i}, q_{i}\right)-\min \left(p-p_{i}, q-q_{i}\right)$. To see this, assume without loss of generality that $p \leq q$. If $p_{i} \leq q_{i}$ for all $i$, then

$$
\min (p, q)-\sum_{i=1}^{N} \min \left(p_{i}, q_{i}\right)=p-\sum_{i=1}^{N} p_{i}=0
$$

so we are done. Assume therefore that $p_{i}>q_{i}$ for $i \in A$, where $A$ is a set of indices with $|A| \geq 1$. For such $i$,

$$
\Delta_{i}=p-q_{i}-\left(p-p_{i}\right)=p_{i}-q_{i}
$$

and thus,

$$
\sum_{i \in A} \Delta_{i}=\sum_{i \in A}\left(p_{i}-q_{i}\right)=p-\sum_{i \notin A} p_{i}-\sum_{i \in A} q_{i}=p-\sum_{i=1}^{N} \min \left(p_{i}, q_{i}\right)
$$

But then,

$$
\max _{1 \leq i \leq N} \Delta_{i} \geq \frac{1}{|A|} \sum_{i \in A} \Delta_{i} \geq \frac{p-\sum_{i=1}^{N} \min \left(p_{i}, q_{i}\right)}{|A|} \geq \frac{p-\sum_{i=1}^{N} \min \left(p_{i}, q_{i}\right)}{N}
$$

and the claim follows. To prove the lemma, notice that since $L(\mathcal{Q}) \leq L^{*}+\epsilon$, it suffices to show that

$$
\max _{Q \in \mathcal{Q}}\left(L(\mathcal{P})-L\left(\mathcal{T}_{Q}\right)\right) \geq \frac{L(\mathcal{P})-L(\mathcal{Q})}{|\mathcal{Q}|}
$$

or equivalently, that

$$
\max _{Q \in \mathcal{Q}}\left(L\left(R_{Q}\right)-L(Q)-L\left(R_{Q}-Q\right)\right) \geq \frac{L(\mathcal{P})-L(\mathcal{Q})}{|\mathcal{Q}|}
$$

where $R_{Q}$ is the unique cell of $\mathcal{P}$ containing $Q$. However,

$$
\begin{aligned}
& \max _{Q \in \mathcal{Q}}\left(L\left(R_{Q}\right)-L(Q)-L\left(R_{Q}-Q\right)\right) \\
& \quad \geq \max _{R \in \mathcal{P}, Q \subseteq R}(L(R)-L(Q)-L(R-Q)) \\
& \quad \geq \max _{R \in \mathcal{P}} \frac{L(R)-\sum_{Q \subseteq R} L(Q)}{|R|}
\end{aligned}
$$

(by the inequality shown above, where $|R|$ denotes
the number of sets of $\mathcal{Q}$ contained in $R$ )

$$
\begin{aligned}
& \geq \frac{\sum_{R \in \mathcal{P}}\left(L(R)-\sum_{Q \subseteq R} L(Q)\right)}{\sum_{R \in \mathcal{P}}|R|} \\
& =\frac{\sum_{R \in \mathcal{P}} L(R)-\sum_{Q \in \mathcal{Q}} L(Q)}{|\mathcal{Q}|} \\
& =\frac{L(\mathcal{P})-L(\mathcal{Q})}{|\mathcal{Q}|}
\end{aligned}
$$

and the lemma is proved.

Let us return to the proof of Theorem 20.9. The previous lemma applied to our situation $\left(\mathcal{P}=\mathcal{P}_{i}\right)$ shows that we may extend $\mathcal{P}_{i}$ by a rectangle $Q \in \mathcal{Q}$ (as $\mathcal{Q}$ will be a collection of rectangles refining both $\mathcal{G}_{l}$ and $\left.\mathcal{P}_{i}\right)$ such that $L\left(\mathcal{T}_{Q}\right)-\left(L^{*}+\epsilon\right)$ is smaller by a guaranteed amount than $L\left(\mathcal{P}_{i}\right)-\left(L^{*}+\epsilon\right)$. $\left(\mathcal{T}_{Q}\right.$ is the partition obtained by extending $\mathcal{P}_{i}$.) To describe $\mathcal{Q}$, we do the following (note that $\mathcal{Q}$ must consist entirely of rectangles for otherwise the lemma is useless): take the rectangles of $\mathcal{P}_{i}$ and extend all four sides (in the order of birth of the rectangles) until they hit a side of another rectangle or an extended border of a previous rectangle if they hit anything at all. Figure 20.28 illustrates the partition into rectangles. Note that this partition consists of $4 i$ line segments, and at most $9 i$ rectangles (this can be seen by noting that we can write in each nonoriginal rectangle of $\mathcal{P}_{i}$ the number of an original neighboring rectangle, each number appearing at most 8 times).


Figure 20.28. Extensions of rectangles to get a rectangular grid.

The rectangular grid partition thus obtained is intersected with $\mathcal{G}_{l}$ to yield $\mathcal{Q}$. To apply Lemma 20.2 , we only need a bound on $|\mathcal{Q}|$. To the rectangular partition just created, we add each of the $2 l-2$ lines of the grid $\mathcal{G}_{l}$ one by one. Start with the horizontal lines. Each time one is added, it creates at most $9 i$ new rectangles. Then, each vertical line adds at most $9 i+l$ new rectangles for a total of not more than $9 i+9 i(l-1)+(l-1)(9 i+l) \leq l^{2}+18 l i$ (assuming $\left.l \geq 1\right)$. Hence,

$$
|\mathcal{Q}| \leq l^{2}+18 l i
$$

Apply the lemma, and observe that there is a rectangle $Q \in \mathcal{Q}$ such that the extension of $\mathcal{P}_{i}$ by $Q$ to $\mathcal{P}_{i}^{\prime}$ would yield

$$
L\left(\mathcal{P}_{i}^{\prime}\right)-\left(L^{*}+\epsilon\right) \leq\left(L\left(\mathcal{P}_{i}\right)-\left(L^{*}+\epsilon\right)\right)\left(1-\frac{1}{l^{2}+18 l i}\right)
$$

At this point in the proof, the reader can safely forget $\mathcal{Q}$. It has done what it was supposed to do.

Of course, we cannot hope to find $\mathcal{P}_{i}^{\prime}$ because we do not know the distribution. Let us denote the actual new partition by $\mathcal{P}_{i+1}$ and let $R_{i+1}$ be the selected rectangle by the empirical minimization. Observe that

$$
\begin{aligned}
& L\left(\mathcal{P}_{i+1}\right)-L\left(\mathcal{P}_{i}^{\prime}\right) \\
& \quad=\left(L\left(\mathcal{P}_{i+1}\right)-\widehat{L}_{n}\left(\mathcal{P}_{i+1}\right)\right)+\left(\widehat{L}_{n}\left(\mathcal{P}_{i+1}\right)-\widehat{L}_{n}\left(\mathcal{P}_{i}^{\prime}\right)\right)+\left(\widehat{L}_{n}\left(\mathcal{P}_{i}^{\prime}\right)-L\left(\mathcal{P}_{i}^{\prime}\right)\right) \\
& \quad \leq\left(L\left(\mathcal{P}_{i+1}\right)-\widehat{L}_{n}\left(\mathcal{P}_{i+1}\right)\right)+\left(\widehat{L}_{n}\left(\mathcal{P}_{i}^{\prime}\right)-L\left(\mathcal{P}_{i}^{\prime}\right)\right)
\end{aligned}
$$

As $\mathcal{P}_{i+1}$ and $\mathcal{R}_{i}^{\prime}$ have most sets in common, many terms cancel in the last double sum. We are left with just $L(\cdot)$ and $\widehat{L}_{n}(\cdot)$ terms for sets of the form $R-$ $R_{i+1}, R_{i+1}, R-Q, Q$ with $R \in \mathcal{P}_{i}$. Thus,

$$
L\left(\mathcal{P}_{i+1}\right)-L\left(\mathcal{P}_{i}^{\prime}\right) \leq 2 \sup _{R^{\prime}}\left|\widehat{L}_{n}\left(R^{\prime}\right)-L\left(R^{\prime}\right)\right|
$$

where the supremum is taken over all rectangles of the above described form. These are sets of the form obtainable in $\mathcal{P}_{i+1}$. Every such set can be written as the difference of a (possibly infinite) rectangle and at most $i+1$ nonoverlapping contained rectangles. As $i+1 \leq k$ in our analysis, we call $\mathcal{Z}_{k}$ the class of all sets
that may be described in this form: $Z_{0}-Z_{1}-\cdots-Z_{k}$, where $Z_{0}, Z_{1}, \ldots, Z_{k}$ are rectangles, each $Z_{i} \subseteq Z_{0}, Z_{1}, \ldots, Z_{k}$ are mutually disjoint. Rectangles may be infinite or half infinite. Hence, for $i<k$,

$$
L\left(\mathcal{P}_{i+1}\right)-L\left(\mathcal{P}_{i}^{\prime}\right) \leq 2 \sup _{Z \in \mathcal{Z}_{k}}\left|\widehat{L}_{n}^{\prime}(Z)-L(Z)\right|
$$

For a fixed $Z \in \mathcal{Z}_{k}$,

$$
\begin{aligned}
\left|\widehat{L}_{n}(Z)-L(Z)\right| & =\left|\min \left(v_{0, n}(Z), v_{1, n}(Z)\right)-\min \left(v_{0}(Z), v_{1}(Z)\right)\right| \\
& \leq\left|v_{0, n}(Z)-v_{0}(Z)\right|+\left|v_{1, n}(Z)-v_{1}(Z)\right|
\end{aligned}
$$

Thus,

$$
L\left(\mathcal{P}_{i+1}\right)-L\left(\mathcal{P}_{i}^{\prime}\right) \leq V_{n} \stackrel{\text { def }}{=} 2 \sup _{Z \in \mathcal{Z}_{k}}\left(\left|v_{0, n}(Z)-v_{0}(Z)\right|+\left|v_{1, n}(Z)-v_{1}(Z)\right|\right)
$$

Define the (good) event

$$
G=\left\{V_{n}<\delta\right\}
$$

where $\delta>0$ will be picked later. On $G$, we see that for all $i<k$,

$$
\begin{aligned}
& L\left(\mathcal{P}_{i+1}\right)-\left(L^{*}+\epsilon\right) \\
& \quad \leq L\left(\mathcal{P}_{i}^{\prime}\right)-\left(L^{*}+\epsilon\right)+\delta \\
& \quad \leq\left(L\left(\mathcal{P}_{i}\right)-\left(L^{*}+\epsilon\right)\right)\left(1-\frac{1}{l^{2}+18 l i}\right)+\delta \quad(\text { Lemma 20.2 })
\end{aligned}
$$

We now introduce a convenient lemma.
Lemma 20.3. Let $a_{n}, b_{n}$ be sequences of positive numbers with $b_{n} \downarrow 0, b_{0}<1$, and let $\delta>0$ be fixed. If $a_{n+1} \leq a_{n}\left(1-b_{n}\right)+\delta$, then

$$
a_{n+1} \leq a_{0} e^{-\sum_{j=0}^{n} b_{j}}+\delta(n+1)
$$

Proof. We have

$$
\begin{aligned}
a_{n+1} & \leq a_{0} \prod_{j=0}^{n}\left(1-b_{j}\right)+\delta\left(\sum_{m=1}^{n-1} \prod_{j=m}^{n}\left(1-b_{j}\right)\right)+\delta \\
& =I+I I+I I I .
\end{aligned}
$$

Clearly, $I \leq a_{0} \exp \left(-\sum_{j=0}^{n} b_{j}\right)$. A trivial bound for $I I+I I I$ is $\delta(n+1)$.
Conclude that on $G$,

$$
\begin{aligned}
L\left(\mathcal{P}_{k}\right)-\left(L^{*}+\epsilon\right) & \leq\left(L\left(\mathcal{P}_{0}\right)-\left(L^{*}+\epsilon\right)\right) e^{-\sum_{j=0}^{k-1} 1 /\left(l^{2}+18 l j\right)}+\delta k \\
& \leq e^{-\int_{0}^{k-1} 1 /\left(l^{2}+18 l u\right) d u}+\delta k \\
& \leq e^{-\frac{1}{18 l} \log \left(\left(l^{2}+18 l(k-1)\right) /\left(l^{2}\right)\right)}+\delta k \\
& =\frac{1}{\left(1+\frac{18(k-1)}{l}\right)^{1 /(18 l)}}+\delta k
\end{aligned}
$$

Thus,

$$
\mathbf{P}\left\{L\left(\mathcal{P}_{k}\right)>L^{*}+2 \epsilon\right\} \leq \mathbf{P}\left\{G^{c}\right\}=\mathbf{P}\left\{V_{n} \geq \delta\right\}
$$

when $\delta k<\epsilon / 2$ and

$$
\frac{1}{\left(1+\frac{18(k-1)}{l}\right)^{1 /(18 l)}}<\frac{\epsilon}{2}
$$

Finally, introduce the notation

$$
\begin{aligned}
L_{n}(R) & =\mathbf{P}\left\{X \in R, Y \neq g_{\mathcal{P}}(X)\right\} \\
L_{n}(\mathcal{P}) & =\sum_{R \in \mathcal{P}} L_{n}(R)
\end{aligned}
$$

As $L\left(\mathcal{P}_{k}\right)$ takes the partition into account, but not the majority vote, and $L_{n}\left(\mathcal{P}_{k}\right)$ does, we need a small additional argument:

$$
\begin{aligned}
L_{n}\left(\mathcal{P}_{k}\right) & =\left(L_{n}\left(\mathcal{P}_{k}\right)-\widehat{L}_{n}\left(\mathcal{P}_{k}\right)\right)+\left(\widehat{L}_{n}\left(\mathcal{P}_{k}\right)-L\left(\mathcal{P}_{k}\right)\right)+L\left(\mathcal{P}_{k}\right) \\
& \leq 2 \sum_{R \in \mathcal{P}_{k}}\left(\left|\nu_{0}(R)-v_{0, n}(R)\right|+\left|v_{1}(R)-v_{1, n}(R)\right|\right)+L\left(\mathcal{P}_{k}\right) \\
& \stackrel{\text { def }}{=} 2 W_{n}+L\left(\mathcal{P}_{k}\right) .
\end{aligned}
$$

Putting things together,

$$
\begin{aligned}
\mathbf{P} & \left\{L_{n}\left(\mathcal{P}_{k}\right)>L^{*}+4 \epsilon\right\} \\
& \leq \mathbf{P}\left\{W_{n}>\epsilon\right\}+\mathbf{P}\left\{L\left(\mathcal{P}_{k}\right)>L^{*}+2 \epsilon\right\} \\
& \leq \mathbf{P}\left\{W_{n}>\epsilon\right\}+\mathbf{P}\left\{V_{n} \geq \delta\right\}
\end{aligned}
$$

under the given conditions on $\delta$ and $k$. Observe that if for a set $Z$,

$$
U_{n}(Z) \stackrel{\text { def }}{=}\left|\nu_{0}(Z)-v_{0, n}(Z)\right|+\left|v_{1}(Z)-v_{1, n}(Z)\right|
$$

then

$$
\begin{aligned}
V_{n} & \leq \sup _{\text {rectangles } Z} U_{n}(Z)+\sup _{\begin{array}{c}
\text { disjoint sets of } k \\
\text { rectangles } Z_{1}, \ldots, Z_{k}
\end{array}} \sum_{i=1}^{k} U_{n}\left(Z_{i}\right) \\
& \leq(k+1) \sup _{\text {rectangles } Z} U_{n}(Z) \\
W_{n} & \leq(k+1) \sup _{\operatorname{rectangles} Z} U_{n}(Z) .
\end{aligned}
$$

We may now use the Vapnik-Chervonenkis inequality (Theorems 12.5 and 13.8) to conclude that for all $\epsilon>0$,

$$
\mathbf{P}\left\{\sup _{\text {rectangles } Z} U_{n}(Z)>\epsilon\right\} \leq 16 n^{2 d} e^{-n \epsilon^{2} / 32}
$$

Armed with this, we have

$$
\begin{aligned}
& \mathbf{P}\left\{W_{n}>\epsilon\right\} \leq 16 n^{2 d} e^{-n\left(\frac{\epsilon}{2(k+1)}\right)^{2} / 32} \\
& \mathbf{P}\left\{V_{n} \geq \delta\right\} \leq 16 n^{2 d} e^{-n\left(\frac{\delta}{2(k+1)}\right)^{2} / 32}
\end{aligned}
$$

Both terms tend to 0 with $n$ if $n /\left(k^{2} \log n\right) \rightarrow \infty$ and $n \delta^{2} /\left(k^{2} \log n\right) \rightarrow \infty$. Take $\delta=\epsilon /(3 k)$ to satisfy our earlier side condition, and note that we need $n /\left(k^{3} \log n\right) \rightarrow \infty$. We, in fact, have strong convergence to 0 for $L_{n}\left(\mathcal{P}_{k}\right)-L^{*}$ by the Borel-Cantelli lemma.

The crucial element in the proof of Theorem 20.9 is the fact that the number of sets in the partition $\mathcal{Q}$ grows at most linearly in $i$, the iteration number. Had it grown quadratically, say, it would not have been good enough - we would not have had guaranteed improvements of large enough sizes to push the error probability towards $L^{*}$. In the multidimensional version and in extension to other types of cuts (Problems 20.29, 20.30) this is virtually the only thing that must be verified. For hyperplane cuts, an additional inequality of the vC type is needed, extending that for classes of hyperplanes.

Our proof is entirely combinatorial and geometric and comes with explicit bounds. The only reference to the Bayes error we need is the quantity $l(\epsilon)$, which is the smallest value $l$ for which

$$
\inf _{\text {all } l \times l \text { grids } \mathcal{G}_{l}} L\left(\mathcal{G}_{l}\right) \leq L^{*}+\epsilon
$$

It depends heavily on the distribution of course. Call $l(\epsilon)$ the grid complexity of the distribution, for lack of a better term. For example, if $X$ is discrete and takes values on the hypercube $\{0,1\}^{d}$, then $l(\epsilon) \leq 2^{d}$. If $X$ takes values on $\left\{a_{1}, \ldots, a_{M}\right\}^{d}$, then $l(\epsilon) \leq(M+1)^{d}$. If $X$ has an arbitrary distribution on the real line and $\eta(x)$ is monotone, then $l(\epsilon) \leq 2$. However, if $\eta$ is unimodal, $l(\epsilon) \leq 3$. In one dimension, $l(\epsilon)$ is sensitive to the number of places where the Bayes decision changes. In two dimensions, however, $l(\epsilon)$ measures the complexity of the distribution, especially near regions with $\eta(x) \approx 1 / 2$. When $X$ is uniform on $[0,1]^{2}$ and $Y=1$ if and only if the components of $X$ sum to less than one, then $l(\epsilon)=\Theta(1 / \sqrt{\epsilon})$ as $\epsilon \downarrow 0$ (Problem 20.32), for example. The grid complexity is eminently suited to study rates of convergence, as it is explicitly featured in the inequalities of our proof. There is no room in this book for following this thread, however.

## Problems and Exercises

Problem 20.1. Let $X$ be a random vector with density $f$ on $\mathcal{R}^{d}$. Show that each component has a density as well.

Problem 20.2. Show that both condition (1) and condition (2) of Theorem 6.1 are necessary for consistency of trees with the $X$-property.

Problem 20.3. For a theoretical median tree in $\mathcal{R}^{d}$ with uniform marginals and $k$ levels of splitting, show that $\mathbf{E}\{H(X)+V(X)\}=1 /\left(2^{k / d}\right)$ when $k$ is a multiple of $d$.

Problem 20.4. $\alpha$-balanced trees. Consider the following generalization of the median tree: at a node in the tree that represents $n$ points, if a split occurs, both subtrees must be of size at least $\alpha(n-1)$ for some fixed $\alpha>0$. Repeat the splits for $k$ levels, resulting in $2^{k}$ leaf regions. The tree has the $X$-property and the classifier is natural. However, the points at which cuts are made are picked among the data points in an arbitrary fashion. The splits rotate through the coordinate axes. Generalize the consistency theorem for median trees.

Problem 20.5. Consider the following tree classifier. First we find the median according to one coordinate, and create two subtrees of sizes $\lfloor(n-1) / 2\rfloor$ and $\lceil(n-1) / 2\rceil$. Repeat this at each level, cutting the next coordinate axis in a rotational manner. Do not cut a node any further if either all data points in the corresponding region have identical $Y_{i}$ values or the region contains less than $k$ points. Prove that the obtained natural classifier is consistent whenever $X$ has a density if $k \rightarrow \infty$ and $\log n / k \rightarrow 0$.

Problem 20.6. Let $\mu$ be a probability measure with a density $f$ on $\mathcal{R}^{d}$, and define

$$
C=\left\{x=\left(x^{(1)}, \ldots, x^{(d)}\right): \mu\left(\left[x^{(1)}, x^{(1)}+\epsilon\right] \times \cdots \times\left[x^{(d)}, x^{(d)}+\epsilon\right]\right)>0 \text { all } \epsilon>0\right\} .
$$

Show that $\mu(C)=1$. Hint: Proceed as in the proof of Lemma A. 1 in the Appendix.
Problem 20.7. Let $U_{(1)}, \ldots, U_{(n)}$ be uniform order statistics defining spacings $S_{1}, \ldots, S_{n+1}$ with $S_{i}=U_{(i)}-U_{(i-1)}$, if $U_{(n+1)}=1$ and $U_{(0)}=0$. Show that
(1) $S_{1}, \ldots, S_{n+1}$ are identically distributed;
(2) $\mathbf{P}\left\{S_{1}>x\right\}=(1-x)^{n}, x \in[0,1]$;
(3) $\mathbf{E}\left\{S_{1}\right\}=1 /(n+1)$.

Problem 20.8. In the proof of Theorem 20.3, we assumed in the second part that horizontal and vertical cuts were meted out independently. Return to the proof and see how you can modify it to take care of the forced alternating cuts.

Problem 20.9. Let $X_{1}, X_{2}, \ldots, X_{n}$ be i.i.d. with common density $f$ on the real line. Call
 independent and that $\mathbf{P}\left\{R_{i}=1\right\}=1 / i$. Conclude that the expected number of records is $\sim \log n$.

Problem 20.10. The deep k-d tree. Assume that $X$ has a density $f$ and that all marginal densities are uniform.
(1) Show that $k \rightarrow \infty$ implies that $\operatorname{diam}(A(X)) \rightarrow 0$ in probability. Do this by arguing that $\operatorname{diam}(A(X)) \rightarrow 0$ in probability for the chronological k -d tree with the same parameter $k$.
(2) In a random k-d tree with $n$ elements, show that the depth $D$ of the last inserted node satisfies $D /(2 \log n) \rightarrow 1$ in probability. Argue first that you may restrict yourself to $d=1$. Then write $D$ as a sum of indicators of records. Conclude by computing $\mathbf{E}\{D\}$ and $\operatorname{Var}\{D\}$ or at least bounding these quantities in an appropriate manner.
(3) Improve the second condition of Theorem 20.4 to $(k-2 \log n) / \sqrt{\log n} \rightarrow-\infty$. (Note that it is possible that $\lim \sup _{n \rightarrow \infty} k / \log n=2$.) Hint: Show that $\mid D-$ $2 \log n \mid=O(\sqrt{\log n})$ in probability by referring to the previous part of this exercise.

Problem 20.11. Consider a full-fledged $\mathrm{k}-\mathrm{d}$ tree with $n+1$ external node regions. Let $L_{n}$ be the probability of error if classification is based upon this tree and if external node regions are assigned the labels (classes) of their immediate parents (see Figure 20.29).


Figure 20.29. The external nodes are assigned the classes of their parents.

Show that whenever $X$ has a density, $\mathbf{E}\left\{L_{n}\right\} \rightarrow L_{\mathrm{NN}} \leq 2 L^{*}$ just as for the 1-nearest neighbor rule.

Problem 20.12. Continuation. Show that $L_{n} \rightarrow L_{\mathrm{NN}}$ in probability.
Problem 20.13. Meisel and Michalopoulos (1973) propose a binary tree classifier with perpendicular cuts in which all leaf regions are homogeneous, that is, all $Y_{i}$ 's for the $X_{i}$ 's in the same region are identical.


Figure 20.30. Example of a tree partition into homogeneous regions.
(1) If $L^{*}=0$, give a stopping rule for constructing a consistent rule whenever $X$ has a density.
(2) If $L^{*}=0$, show that there exists a consistent homogeneous partition classifier with $o(n)$ expected number of cells.
(3) If $L^{*}>0$, show a (more complicated) way of constructing a homogeneous partition that yields a tree classifier with $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ whenever $X$ has a density. Hint: First make a consistent nonhomogeneous binary tree classifier, and refine it to make it homogeneous.
(4) If $L^{*}>0$, then show that the expected number of homogeneous regions is at least $c n$ for some $c>0$. Such rules are therefore not practical, unless $L^{*}=0$. Overfitting will occur.

Problem 20.14. Let $X$ be uniform on $[0,1]$ and let $Y$ be independent of $X$ with $\mathbf{P}\{Y=$ $1\}=p$. Find a tree classifier based upon simple interval splitting for which each region has one data point, and

$$
\liminf _{p \rightarrow 0} \frac{\liminf _{n \rightarrow \infty} \mathbf{E} L_{n}}{p} \geq 3
$$

We know that for the nearest neighbor rule,

$$
\lim _{p \rightarrow 0} \frac{\lim \sup _{n \rightarrow \infty} \mathbf{E} L_{n}}{p}=2
$$

so that the given interval splitting method is worse than the nearest neighbor method. (This provides a counterexample to a conjecture of Breiman, Friedman, Olshen, and Stone (1984, pp.89-90).) Hint: Sort the points, and adjust the sizes of the intervals given to the class 0 and class 1 points in favor of the 1 's. By doing so, the asymptotic probability of error can be made as close as desired to $\left(p^{2}+3 p(1-p)\right)(1-p)$.

Problem 20.15. Continuation. Let $X$ be uniform on $[0,1]^{2}$ and let $Y$ be independent of $X$ with $\mathbf{P}\{Y=1\}=p$. Find a data-based tree classifier based upon perpendicular cuts for which each region has one data point, $\operatorname{diam}(A(X)) \rightarrow 0$ in probability (recall Theorems 6.1 and 21.2), and

$$
\liminf _{p \rightarrow 0} \frac{\liminf _{n \rightarrow \infty} \mathbf{E} L_{n}}{p}=\infty
$$

Conclude that in $\mathcal{R}^{d}$, we can construct tree classifiers with one point per cell that are much worse than the nearest neighbor rule. Hint: The next two problems may help you with the construction and analysis.

Problem 20.16. Continuation. Draw an i.i.d. sample $X_{1}, \ldots, X_{n}$ from the uniform distribution on $[0,1]^{2}$, and let $Y_{1}, \ldots, Y_{n}$ be i.i.d. and independent of the $X_{i}$ 's with $\mathbf{P}\{Y=1\}=p$. Construct a binary tree partition with perpendicular cuts for $\left\{X_{i}: Y_{i}=1\right\}$ such that every leaf region has one and only one point and $\operatorname{diam}(A(X)) \rightarrow 0$ in probability.
(1) How would you proceed, avoiding putting $X_{i}$ 's on borders of regions?
(2) Prove $\operatorname{diam}(A(X)) \rightarrow 0$ in probability.
(3) Add the $\left(X_{i}, Y_{i}\right)$ pairs with $Y_{i}=0$ to the leaf regions, so that every region has one class-1 observation and zero or more class-0 observations. Give the class-1 observation the largest area containing no class-0 points, as shown in Figure 20.31. Show that this can always be done by adding perpendicular cuts and keeping at least one observation per region.


Figure 20.31. Cutting a rectangle by giving a large area to the single class- 1 point.
(4) Partition all rectangles with more than one point further to finally obtain a one-point-per-leaf-region partition. If there are $N$ points in a region of the tree before the class- 1 and class- 0 points are separated (thus, $N-1$ class- 0 points and one class-1 point), then show that the expected proportion of the region's area given to class 1 , given $N$, times $N$ tends to $\infty$. (An explicit lower bound will be helpful.) Hint: Use the next problem.
(5) Write the probability of error for the rule in terms of areas of rectangles, and use part (4) to get an asymptotic lower bound.
(6) Now let $p$ tend to zero and get an asymptotic expression for your lower bound in terms of $p$. Compare this with $2 p(1-p)$, the asymptotic error probability of the 1-nearest neighbor rule.

Problem 20.17. Continuation. Draw a sample $X_{1}, \ldots, X_{n}$ of $n$ i.i.d. observations uniformly distributed on $[0,1]^{d}$. The rectangles defined by the origin and $X_{1}, \ldots, X_{n}$ as opposite vertices are denoted by $R_{1}, \ldots, R_{n}$, respectively. The probability content of $R_{i}$ is clearly $\mu\left(R_{i}\right)=\prod_{i=1}^{d} X_{i}^{(j)}$. Study

$$
M_{n}=\max \left\{\mu\left(R_{i}\right): 1 \leq i \leq n, R_{i} \text { does not contain } R_{j} \text { for any } j \neq i\right\},
$$

the probability content of the largest empty rectangle with a vertex at the origin. For $d=1$, $M_{n}$ is just the minimum of the $X_{i}$ 's, and thus $n M_{n} \xrightarrow{\mathcal{L}} \mathcal{E}$, where $\mathcal{E}$ is the exponential distribution. Also $\mathbf{E}\left\{M_{n}\right\}=1 /(n+1)$. For $d>1, M_{n}$ is larger. Show that $n M_{n} \rightarrow \infty$ in probability and try obtaining the first term in the rate of increase.

Problem 20.18. Show that $\operatorname{diam}(A(X)) \rightarrow 0$ in probability for the chronological quadtree whenever $k \rightarrow 0$ and $X$ has a density. Hint: Mimic the proof for the chronological quadtree.

Problem 20.19. Show that the deep quadtree is consistent if $X$ has a density and $k$ levels of splits are applied, where $k \rightarrow \infty$ and $k / \log n \rightarrow 0$.

Problem 20.20. Consider a full-fledged quadtree with $n$ nodes (and thus $n\left(2^{d}-1\right)+1$ leaf regions). Assign to each region the $Y$-label (class) of its parent in the quadtree. With this simple classifier, show that whenever $X$ has a density, $\mathbf{E}\left\{L_{n}\right\} \rightarrow 2 \mathbf{E}\{\eta(X)(1-\eta(X))\}=L_{\mathrm{NN}}$. In particular, $\lim \sup _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\} \leq 2 L^{*}$ and the classifier is consistent when $L^{*}=0$.

Problem 20.21. In $\mathcal{R}^{2}$, partition the space as follows: $X_{1}, X_{2}$ define nine regions by vertical and horizontal lines through them. $X_{3}, \ldots, X_{k}$ are sent down to the appropriate subtrees in the 9 -ary tree, and within each subtree with at least two points, the process is repeated recursively. A decision at $x$ is by a majority vote (among $\left(X_{k+1}, Y_{k+1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ ) among those $X_{i}$ 's in the same rectangle of the partition as $x$ Show that if $k \rightarrow \infty, k / n \rightarrow 0$, the rule is consistent whenever $X$ has a density.

Problem 20.22. On the two-dimensional counterexample shown in the text for multivariate Stoller splits, prove that if splits are performed based upon a sample drawn from the distribution, and if we stop after $k$ splits with $k$ depending on $n$ in such a way that $k / \log n \rightarrow 0$, then $L_{n}$, the conditional probability of error, satisfies ${\lim \inf _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\} \geq(1-\epsilon) / 2 \text {. Hint: }}_{\text {I }}$. Bound the probability of ever splitting $[1,2]^{2}$ anywhere by noting that the maximal difference between the empirical distribution functions for the first coordinate of $X$ given $Y=0$ and $Y=1$ is $O(1 / \sqrt{n})$ when restricted to $[1,2]^{2}$.

Problem 20.23. Let $X$ have the uniform distribution on [0,5] and let $Y=I_{\{2<X<3\}}$, so that $L^{*}=0$. Construct a binary classification tree by selecting at each iteration the split that minimizes the impurity function $I$, where $\psi$ is the Gini criterion. Consider just the first three splits made in this manner. Let $L_{n}$ be the probability of error with the given rule (use a majority vote over the leaf regions). Show that $L_{n} \rightarrow 0$ in probability. Analyze the algorithm when the Gini criterion is replaced by the probability-of-error criterion.

Problem 20.24. Let $X$ be uniform on $[0,1]$ and let $Y$ be independent of $X$, with $\mathbf{P}\{Y=$ $1\}=2 / 3$. Draw a sample of size $n$ from this distribution. Investigate where the first cut
might take place, based upon minimizing the impurity function with $\psi(p, 1-p)$ (the Gini criterion). (Once this is established, you will have discovered the nature of the classification tree, roughly speaking.)

Problem 20.25. Complete the consistency proof of Theorem 20.7 for the raw BSP tree for $\mathcal{R}^{2}$ by showing that $\operatorname{diam}(A(X)) \rightarrow 0$ in probability.

Problem 20.26. Balanced bsp trees. We generalize median trees to allow splitting the space along hyperplanes.


Figure 20.32. A balanced BSP tree: Each hyperplane cut splits a region into two cells of the same cardinality.

Assume that $X$ has a density and that the tree possesses the $X$-property. Keep splitting until there are $2^{k}$ leaf regions, as with median trees. Call the trees balanced BSP trees.
(1) Show that there are ways of splitting in $\mathcal{R}^{2}$ that lead to nonconsistent rules, regardless how $k$ varies with $n$.
(2) If every splitting hyperplane is forced to contain $d$ data points (in $\mathcal{R}^{d}$ ) and these data points stay with the splitting node (they are not sent down to subtrees), then show that once again, there exists a splitting method that leads to nonconsistent rules, regardless of how $k$ varies with $n$.

Problem 20.27. Let $X$ have a uniform distribution on the unit ball of $\mathcal{R}^{d}$. Let $Y=I_{[\|X\| \geq 1 / 2]}$, so that $L^{*}=0$. Assume that we split the space by a hyperplane by minimizing an impurity function based upon the Gini criterion. If $n$ is very large, where approximately would the cut take place (modulo a rotation, of course)?

Problem 20.28. There exist singular continuous distributions that admit uniform [0, 1] marginals in $\mathcal{R}^{d}$. Show, for example, that if $X$ is uniformly distributed on the surface of the unit sphere of $\mathcal{R}^{3}$, then its three components are all uniformly distributed on $[-1,1]$.

Problem 20.29. Verify that Theorem 20.9 remains valid in $\mathcal{R}^{d}$.
Problem 20.30. Prove that Theorem 20.9 remains valid if rectangular cuts are replaced by any of the elementary cuts shown on Figures 20.23 to 20.25, and such cuts are performed recursively $k$ times, always by maximizing the decrease of the empirical error.

Problem 20.31. Show that Theorem 20.11 remains valid if $k \rightarrow \infty$ and $k=o(\sqrt{n / \log n})$. Hint: In the proof of the theorem, the bound on $V_{n}$ and $W_{n}$ is loose. You may get more efficient bounds by applying Theorem 21.1 from the next chapter.

Problem 20.32. Study the behavior of the grid complexity as $\epsilon \downarrow 0$ for the following cases:
(1) $X$ is uniform on the perimeter of the unit circle of $\mathcal{R}^{2}$ with probability $1 / 2$ and $X$ is uniform on $[0,1]^{2}$ otherwise. Let $Y=1$ if and only if $X$ is on the perimeter of that circle (so that $L^{*}=0$ ).
(2) $X=\left(X^{(1)}, X^{(2)}\right)$ is uniform on $[0,1]^{2}$ and $Y=1$ if and only if $X^{(1)}+X^{(2)} \leq 1$.

## 21

## Data-Dependent Partitioning

### 21.1 Introduction

In Chapter 9 we investigated properties of the regular histogram rule. Histogram classifiers partition the observation space $\mathcal{R}^{d}$ and classify the input vector $X$ according to a majority vote among the $Y_{i}$ 's whose corresponding $X_{i}$ 's fall in the same cell of the partition as $X$. Partitions discussed in Chapter 9 could depend on the sample size $n$, but were not allowed to depend on the data $D_{n}$ itself. We dealt mostly with grid partitions, but will now allow other partitions as well. Just consider clustered training observations $X_{1}, \ldots, X_{n}$. Near the cluster's center finer partitions are called for. Similarly, when the components have different physical dimensions, the scale of one coordinate axis is not related at all to the other scales, and some data-adaptive stretching is necessary. Sometimes the data are concentrated on or around a hyperplane. In all these cases, although consistent, the regular histogram method behaves rather poorly, especially if the dimension of the space is large. Therefore, it is useful to allow data-dependent partitions, while keeping a majority voting scheme within each cell.

The simplest data-dependent partitioning methods are based on statistically equivalent blocks in which each cell contains the same number of points. In onedimensional problems statistically equivalent blocks reduce to $k$-spacing estimates where the $k$-th, $2 k$-th, etc. order statistics determine the partition of the real line.

Sometimes, it makes sense to cluster the data points into groups such that points in a group are close to each other, and define the partition so that each group is in a different cell.

Many other data-dependent partitioning schemes have been introduced in the lit-
erature. In most of these algorithms the cells of the partition correspond to the leaves of a binary tree, which makes computation of the corresponding classification rule fast and convenient. Tree classifiers were dealt with in Chapter 20. Analysis of universal consistency for these algorithms and corresponding density estimates was begun by Abou-Jaoudé (1976b) and Gordon and Olshen (1984), (1978), (1980) in a general framework, and was extended, for example, by Breiman, Friedman, Olshen, and Stone (1984), Chen and Zhao (1987), Zhao, Krishnaiah, and Chen (1990), Lugosi and Nobel (1996), and Nobel (1994).

This chapter is more general than the chapter on tree classifiers, as every partition induced by a tree classifier is a valid partition of space, but not vice versa. The example below shows a rectangular partition of the plane that cannot be obtained by consecutive perpendicular cuts in a binary classification tree.


Figure 21.1. A rectangular partition of $[0,1]^{2}$ that cannot be achieved by a tree of consecutive cuts.

In this chapter we first establish general sufficient conditions for the consistency of data-dependent histogram classifiers. Because of the complicated dependence of the partition on the data, methods useful for handling regular histograms have to be significantly refined. The main tool is a large deviation inequality for families of partitions that is related to the Vapnik-Chervonenkis inequality for families of sets. The reader is asked for forgiveness-we want to present a very generally applicable theorem and have to sacrifice (temporarily) by increasing the density of the text. However, as you will discover, the rewards will be sweet.

### 21.2 A Vapnik-Chervonenkis Inequality for Partitions

This is a technical section. We will use its results in the next section to establish a general consistency theorem for histogram rules with data-dependent partitions. The main goal of this section is to extend the basic Vapnik-Chervonenkis inequality (Theorem 12.5) to families of partitions from families of sets. The line of thought followed here essentially appears in Zhao, Krishnaiah, and Chen (1990) for rectangular partitions, and more generally in Lugosi and Nobel (1996). A substantial simplification in the proof was pointed out to us by Andrew Barron.

By a partition of $\mathcal{R}^{d}$ we mean a countable collection $\mathcal{P}=\left\{A_{1}, A_{2}, \ldots\right\}$ of subsets of $\mathcal{R}^{d}$ such that $\cup_{j=1}^{\infty} A_{j}=\mathcal{R}^{d}$ and $A_{i} \cap A_{j}=\emptyset$ if $i \neq j$. Each set $A_{j}$ is called a cell of the partition $\mathcal{P}$.

Let $M$ be a positive number. For each partition $\mathcal{P}$, we define $\mathcal{P}^{(M)}$ as the restriction of $\mathcal{P}$ to the ball $S_{M}$ (recall that $S_{M} \subset \mathcal{R}^{d}$ denotes the closed ball of radius $M$ centered at the origin). In other words, $\mathcal{P}^{(M)}$ is a partition of $S_{M}$, whose cells are obtained by intersecting the cells of $\mathcal{P}$ with $S_{M}$. We assume throughout that $\mathcal{P}$ is such that $\left|\mathcal{P}^{(M)}\right|<\infty$ for each $M<\infty$. We denote by $\mathcal{B}\left(\mathcal{P}^{(M)}\right)$ the collection of all $2^{\left|{ }^{(M)}\right|}$ sets obtained by unions of cells of $\mathcal{P}^{(M)}$.
Just as we dealt with classes of sets in Chapter 12, here we introduce families of partitions. Let $\mathcal{F}$ be a (possibly infinite) collection of partitions of $\mathcal{R}^{d}$. $\mathcal{F}^{(M)}=\left\{\mathcal{P}^{(M)}: \mathcal{P} \in \mathcal{F}\right\}$ denotes the family of partitions of $S_{M}$ obtained by restricting members of $\mathcal{F}$ to $S_{M}$. For each $M$, we will measure the complexity of a family of partitions $\mathcal{F}^{(M)}$ by the shatter coefficients of the class of sets obtained as unions of cells of partitions taken from the family $\mathcal{F}^{(M)}$. Formally, we define the combinatorial quantity $\Delta_{n}\left(\mathcal{F}^{(M)}\right)$ as follows: introduce the class $\mathcal{A}^{(M)}$ of subsets of $\mathcal{R}^{d}$ by

$$
\mathcal{A}^{(M)}=\left\{A \in \mathcal{B}\left(\mathcal{P}^{(M)}\right) \text { for some } \mathcal{P}^{(M)} \in \mathcal{F}^{(M)}\right\},
$$

and define

$$
\Delta_{n}\left(\mathcal{F}^{(M)}\right)=s\left(\mathcal{A}^{(M)}, n\right),
$$

the shatter coefficient of $\mathcal{A}^{(M)} . \mathcal{A}^{(M)}$ is thus the class of all sets that can be obtained as unions of cells of some partition of $S_{M}$ in the collection $\mathcal{F}^{(M)}$. For example, if all members of $\mathcal{F}^{(M)}$ partition $S_{M}$ into two sets, then $\Delta_{n}\left(\mathcal{F}^{(M)}\right)$ is just the shatter coefficient of all sets in these partitions (with $\emptyset$ and $S_{M}$ included in the collection of sets).
Let $\mu$ be a probability measure on $\mathcal{R}^{d}$ and let $X_{1}, X_{2}, \ldots$ be i.i.d. random vectors in $\mathcal{R}^{d}$ with distribution $\mu$. For $n=1,2, \ldots$ let $\mu_{n}$ denote the empirical distribution of $X_{1}, \ldots, X_{n}$, which places mass $1 / n$ at each of $X_{1}, \ldots, X_{n}$. To establish the consistency of data-driven histogram methods, we require information about the large deviations of random variables of the form

$$
\sup _{\mathcal{P}(M) \in \mathcal{F}^{(M)}} \sum_{A \in \mathcal{P}^{(M)}}\left|\mu_{n}(A)-\mu(A)\right|,
$$

where $\mathcal{F}$ is an appropriate family of partitions.
Remark. Just as in Chapter 12, the supremum above is not guaranteed to be measurable. In order to insure measurability, it is necessary to impose regularity conditions on uncountable collections of partitions. It suffices to mention that in all our applications, the measurability can be verified by checking conditions given, e.g., in Dudley (1978), or Pollard (1984).

The following theorem is a consequence of the Vapnik-Chervonenkis inequality:
Theorem 21.1. (Lugosi and Nobel (1996)). Let $X_{1}, \ldots, X_{n}$ be i.i.d. random vectors in $\mathcal{R}^{d}$ with measure $\mu$ and empirical measure $\mu_{n}$. Let $\mathcal{F}$ be a collection of partitions of $\mathcal{R}^{d}$. Then for each $M<\infty$ and $\epsilon>0$,

$$
\mathbf{P}\left\{\sup _{\mathcal{P}^{(M)} \in \mathcal{F}(M)} \sum_{A \in \mathcal{P}(M)}\left|\mu_{n}(A)-\mu(A)\right|>\epsilon\right\} \leq 8 \Delta_{n}\left(\mathcal{F}^{(M)}\right) e^{-n \epsilon^{2} / 512}+e^{-n \epsilon^{2} / 2} .
$$

Proof. For a fixed partition $\mathcal{P}$, define $A^{\prime}$ as the set

$$
A^{\prime}=\bigcup_{A \in \mathcal{P}^{(M)}: \mu_{n}(A) \geq \mu(A)} A
$$

Then

$$
\begin{aligned}
& \sum_{A \in \mathcal{P}^{(M)}}\left|\mu_{n}(A)-\mu(A)\right| \\
& \quad=\sum_{A \in \mathcal{P}^{(M)}: \mu_{n}(A) \geq \mu(A)}\left(\mu_{n}(A)-\mu(A)\right) \\
& \quad+\sum_{A \in \mathcal{P}^{(M)}: \mu_{n}(A)<\mu(A)}\left(\mu(A)-\mu_{n}(A)\right) \\
& =2\left(\mu_{n}\left(A^{\prime}\right)-\mu\left(A^{\prime}\right)\right)+\mu\left(S_{M}\right)-\mu_{n}\left(S_{M}\right) \\
& \leq 2 \sup _{A \in \mathcal{B}\left(\mathcal{P}^{(M)}\right)}\left|\mu_{n}(A)-\mu(A)\right|+\mu\left(S_{M}\right)-\mu_{n}\left(S_{M}\right)
\end{aligned}
$$

Recall that the class of sets $\mathcal{B}\left(\mathcal{P}^{(M)}\right)$ contains all $2^{\left|\mathcal{P}^{(M)}\right|}$ sets obtained by unions of cells of $\mathcal{P}^{(M)}$. Therefore,

$$
\begin{aligned}
& \sup _{\mathcal{P}^{(M)} \in \mathcal{F}^{(M)}} \sum_{A \in \mathcal{P}^{(M)}}\left|\mu_{n}(A)-\mu(A)\right| \\
& \leq 2 \sup _{\mathcal{P}^{(M)} \in \mathcal{F}^{(M)}} \sup _{A \in \mathcal{B}\left(\mathcal{P}^{(M)}\right) \cup\left\{S_{M}\right\}}\left|\mu_{n}(A)-\mu(A)\right|+\mu\left(S_{M}\right)-\mu_{n}\left(S_{M}\right) .
\end{aligned}
$$

Observe that the first term on the right-hand side of the inequality is a uniform deviation of the empirical measure $\mu_{n}$ from $\mu$ over a specific class of sets. The class contains all sets that can be written as unions of cells of a partition $\mathcal{P}^{(M)}$ in the class of partitions $\mathcal{F}^{(M)}$. This class of sets is just $\mathcal{A}^{(M)}$, defined above. The theorem now follows from the Vapnik-Chervonenkis inequality (Theorem 12.5), the definition of $\Delta_{n}\left(\mathcal{F}^{(M)}\right)$, and Hoeffding's inequality (Theorem 8.1).

We will use a special application of Theorem 21.1, summarized in the following corollary:

Corollary 21.1. Let $\left(X_{1}, Y_{1}\right),\left(X_{2}, Y_{2}\right) \ldots$ be a sequence of i.i.d. random pairs in $\mathcal{R}^{d} \times\{0,1\}$ and let $\mathcal{F}_{1}, \mathcal{F}_{2}, \ldots$ be a sequence of families of partitions of $\mathcal{R}^{d}$. If for $M<\infty$

$$
\lim _{n \rightarrow \infty} \frac{\log \left(\Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right)\right)}{n}=0
$$

then

$$
\sup _{\mathcal{P}^{(M)} \in \mathcal{F}_{n}^{(M)}} \sum_{A \in \mathcal{P}^{(M)}}\left|\frac{1}{n} \sum_{i=1}^{n} Y_{i} I_{\left\{X_{i} \in A\right\}}-\mathbf{E}\left\{Y I_{\{X \in A\}}\right\}\right| \rightarrow 0
$$

with probability one as $n$ tends to infinity.

Proof. Let $v$ be the measure of $(X, Y)$ on $\mathcal{R}^{d} \times\{0,1\}$, and let $v_{n}$ be the empirical measure corresponding to the sequence $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$. Using $\mathcal{F}_{n}^{(M)}$ we define a family $\mathcal{G}_{n}^{(M)}$ of partitions of $\mathcal{R}^{d} \times\{0,1\}$ by

$$
\mathcal{G}_{n}^{(M)}=\left\{\mathcal{P}^{(M)} \times\{0\}: \mathcal{P}^{(M)} \in \mathcal{F}_{n}^{(M)}\right\} \cup\left\{\mathcal{P}^{(M)} \times\{1\}: \mathcal{P}^{(M)} \in \mathcal{F}_{n}^{(M)}\right\},
$$

where $\mathcal{P} \times\{0\}=\left\{A_{1}, A_{2}, \ldots\right\} \times\{0\} \stackrel{\text { def }}{=}\left\{A_{1} \times\{0\}, A_{2} \times\{0\}, \ldots\right\}$. We apply Theorem 21.1 for these families of partitions.

$$
\begin{aligned}
& \sup _{\mathcal{P}^{(M)} \in \mathcal{F}_{n}^{(M)}} \sum_{A \in \mathcal{P}^{(M)}}\left|\frac{1}{n} \sum_{i=1}^{n} Y_{i} I_{\left\{X_{i} \in A\right\}}-\mathbf{E}\left\{Y I_{\{X \in A\}}\right\}\right| \\
= & \sup _{\mathcal{P}^{(M)} \in \mathcal{F}_{n}^{(M)}} \sum_{A \in \mathcal{P}^{(M)}}\left|v_{n}(A \times\{1\})-v(A \times\{1\})\right| \\
& \leq \sup _{\overline{\mathcal{P}}^{(M)} \in \mathcal{G}_{n}^{(M)}} \sum_{\bar{A} \in \overline{\mathcal{P}}^{(M)}}\left|v_{n}(\bar{A})-v(\bar{A})\right| .
\end{aligned}
$$

It is easy to see that $\Delta_{n}\left(\mathcal{G}_{n}^{(M)}\right)=\Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right)$. Therefore the stated convergence follows from Theorem 21.1 and the Borel-Cantelli lemma. $\square$

Lemma 21.1. Assume that the family $\mathcal{F}^{(M)}$ is such that the number of cells of the partitions in the family are uniformly bounded, that is, there is a constant $N$ such that $\left|\mathcal{P}^{(M)}\right| \leq N$ for each $\mathcal{P}^{(M)} \in \mathcal{F}^{(M)}$. Then

$$
\Delta_{n}\left(\mathcal{F}^{(M)}\right) \leq 2^{N} \Delta_{n}^{*}\left(\mathcal{F}^{(M)}\right),
$$

where $\Delta_{n}^{*}\left(\mathcal{F}^{(M)}\right)$ is maximal number of different ways $n$ points can be partitioned by members of $\mathcal{F}^{(M)}$.

Example. Flexible grids. As a first simple example, let us take in $\mathcal{F}_{n}$ all partitions into $d$-dimensional grids (called flexible grids as they may be visualized as chickenwire fences with unequally spaced vertical and horizontal wires) in which cells are made up as Cartesian products of $d$ intervals, and each coordinate axis contributes one of $m_{n}$ intervals to these products. Clearly, if $\mathcal{P}$ is a member partition of $\mathcal{F}_{n}$, then $|\mathcal{P}|=m_{n}^{d}$. Nevertheless, there are uncountably many $\mathcal{P}$ 's, as there are uncountably many intervals of the real line. This is why the finite quantity $\Delta_{n}$ comes in so handy. We will verify later that for each $M$,

$$
\Delta_{n}\left(\mathcal{F}^{(M)}\right) \leq 2^{m_{n}^{d}}\binom{n+m_{n}}{n}^{d},
$$

so that the condition of Corollary 21.1 is fulfilled when

$$
\lim _{n \rightarrow \infty} \frac{m_{n}^{d}}{n}=0 .
$$



Figure 21.2. A flexible grid partition.

### 21.3 Consistency

In this section we establish a general consistency theorem for a large class of data-based partitioning rules. Using the training data $D_{n}$ we produce a partition $\mathcal{P}_{n}=\pi_{n}\left(X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right)$ according to a prescribed rule $\pi_{n}$. We then use the partition $\mathcal{P}_{n}$ in conjunction with $D_{n}$ to produce a classification rule based on a majority vote within the cells of the (random) partition. That is, the training set is used twice and it is this feature of data-dependent histogram methods that distinguishes them from regular histogram methods.

Formally, an $n$-sample partitioning rule for $\mathcal{R}^{d}$ is a function $\pi_{n}$ that associates every $n$-tuple of pairs $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right) \in \mathcal{R}^{d} \times\{0,1\}$ with a measurable partition of $\mathcal{R}^{d}$. Associated with every partitioning rule $\pi_{n}$ there is a fixed, nonrandom family of partitions

$$
\mathcal{F}_{n}=\left\{\pi_{n}\left(x_{1}, y_{1}, \ldots, x_{n}, y_{n}\right):\left(x_{1}, y_{1}\right) \ldots,\left(x_{n}, y_{n}\right) \in \mathcal{R}^{d} \times\{0,1\}\right\}
$$

$\mathcal{F}_{n}$ is the family of partitions produced by the partitioning rule $\pi_{n}$ for all possible realizations of the training sequence $D_{n}$. When a partitioning rule $\pi_{n}$ is applied to the sequence $D_{n}=\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, it produces a random partition $\mathcal{P}_{n}=$ $\pi_{n}\left(D_{n}\right) \in \mathcal{F}_{n}$. In what follows we suppress the dependence of $\mathcal{P}_{n}$ on $D_{n}$ for notational simplicity. For every $x \in \mathcal{R}^{d}$ let $A_{n}(x)$ be the unique cell of $\mathcal{P}_{n}$ that contains the point $x$.

Now let $\left\{\pi_{1}, \pi_{2}, \ldots\right\}$ be a fixed sequence of partitioning rules. The classification rule $g_{n}(\cdot)=g_{n}\left(\cdot, D_{n}\right)$ is defined by taking a majority vote among the classes appearing in a given cell of $\mathcal{P}_{n}$, that is,

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{X_{i} \in A_{n}(x), Y_{i}=1\right\}} \leq \sum_{i=1}^{n} I_{\left\{X_{i} \in A_{n}(x), Y_{i}=0\right\}} \\ 1 & \text { otherwise } .\end{cases}
$$

We emphasize here that the partition $\mathcal{P}_{n}$ can depend on the vectors $X_{i}$ and the labels $Y_{i}$. First we establish the strong consistency of the rules $\left\{g_{n}\right\}$ for a wide class of partitioning rules. As always, $\operatorname{diam}(A)$ denotes the diameter of a set $A$, that is, the maximum Euclidean distance between any two points of $A$ :

$$
\operatorname{diam}(A)=\sup _{x, y \in A}\|x-y\|
$$

Theorem 21.2. (Lugosi and Nobel (1996)). Let $\left\{\pi_{1}, \pi_{2}, \ldots\right\}$ be a fixed sequence of partitioning rules, and for each $n$ let $\mathcal{F}_{n}$ be the collection of partitions associated with the $n$-sample partitioning rule $\pi_{n}$. If
(i) for each $M<\infty$

$$
\lim _{n \rightarrow \infty} \frac{\log \left(\Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right)\right)}{n}=0,
$$

and
(ii) for all balls $S_{B}$ and all $\gamma>0$,

$$
\lim _{n \rightarrow \infty} \mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x) \cap S_{M}\right)>\gamma\right\}\right)=0
$$

with probability one,
then the classification rule $\left\{g_{n}\right\}$ corresponding to $\pi_{n}$ satisfies

$$
L\left(g_{n}\right) \rightarrow L^{*}
$$

with probability one. In other words, the rule $\left\{g_{n}\right\}$ is strongly consistent.
Remark. In some applications we need a weaker condition to replace (ii) in the theorem. The following condition will do: for every $\gamma>0$ and $\delta \in(0,1)$

$$
\lim _{n \rightarrow \infty} \inf _{T \subset \mathcal{R}^{d}: \mu(T) \geq 1-\delta} \mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x) \cap T\right)>\gamma\right\}\right)=0 \quad \text { with probability one. }
$$

The verification of this is left to the reader (Problem 21.2).
The proof, given below, requires quite some effort. The utility of the theorem is not immediately apparent. The length of the proof is indicative of the generality of the conditions in the theorem. Given a data-dependent partitioning rule, we must verify two things: condition (i) merely relates to the richness of the class of partitions of $\mathcal{R}^{d}$ that may possibly occur, such as flexible grids. Condition (ii) tells us that the rule should eventually make local decisions. From examples in earlier chapters, it should be obvious that (ii) is not necessary. Finite partitions of $\mathcal{R}^{d}$ necessarily have component sets with infinite diameter, hence we need a condition that states that such sets have small $\mu$-measure. Condition (ii) requires that a randomly chosen cell have a small diameter. Thus, it may be viewed as the "with-probability-one" version of condition (1) of Theorem 6.1. However, the weaker version of condition (ii) stated in the above remark is more subtle. By considering examples in which $\mu$ has bounded support, more than just balls $S_{M}$ are needed, as the sets of the partition near the boundary of the support may all have infinite diameter as well. Hence we introduced an infimum with respect to sets $T$ over all $T$ with $\mu(T) \geq 1-\delta$.
It suffices to mention that (ii) is satisfied for all distributions in some of the examples that follow. Theorem 21.2 then allows us to conclude that such rules are strongly universally consistent. The theorem has done most of the digestion of
such proofs, so we are left with virtually no work at all. Consistency results will follow like dominos falling.

Proof of Theorem 21.2. Observe that the partitioning classifier $g_{n}$ can be rewritten in the form

$$
g_{n}(x)= \begin{cases}0 & \text { if } \frac{\sum_{i=1}^{n} \frac{I_{\left\{Y_{i}=1\right]} I_{\left\{x_{i} \in A_{n}(x)\right]}}{\mu\left(A_{n}(x)\right)}}{1}\end{cases}
$$

Introduce the notation

$$
\eta_{n}(x)=\frac{\sum_{i=1}^{n} Y_{i} I_{\left\{X_{i} \in A_{n}(x)\right\}}}{n \mu\left(A_{n}(x)\right)} .
$$

For any $\epsilon>0$, there is an $M \in(0, \infty)$ such that $\mathbf{P}\left\{X \notin S_{M}\right\}<\epsilon$. Thus, by an application of Theorem 2.3 we see that

$$
\begin{aligned}
& L\left(g_{n}\right)-L^{*} \\
& \quad \leq \mathbf{P}\left\{g_{n}(X) \neq Y, X \in S_{M} \mid D_{n}\right\}-\mathbf{P}\left\{g^{*}(X) \neq Y, X \in S_{M}\right\}+\epsilon \\
& \quad \leq \int_{S_{M}}\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)+\int_{S_{M}}\left|(1-\eta(x))-\eta_{n}^{(0)}(x)\right| \mu(d x)+\epsilon,
\end{aligned}
$$

where

$$
\eta_{n}^{(0)}(x)=\frac{\sum_{i=1}^{n}\left(1-Y_{i}\right) I_{\left\{X_{i} \in A_{n}(x)\right\}}}{n \mu\left(A_{n}(x)\right)}
$$

By symmetry, since $\epsilon>0$ is arbitrary, it suffices to show that for each $M<\infty$,

$$
\int_{S_{M}}\left|\eta(x)-\eta_{n}(x)\right| \mu(d x) \rightarrow 0 \text { with probability one. }
$$

Fix $\epsilon>0$ and let $r: \mathcal{R}^{d} \rightarrow \mathcal{R}$ be a continuous function with bounded support such that $\int_{S_{M}}|\eta(x)-r(x)| \mu(d x)<\epsilon$. Such $r$ exists by Theorem A.8. Now define the auxiliary functions

$$
\tilde{\eta}_{n}(x)=\frac{\mathbf{E}\left\{Y I_{\left\{X \in A_{n}(x)\right\}} \mid D_{n}\right\}}{\mu\left(A_{n}(x)\right)}
$$

and

$$
\tilde{r}_{n}(x)=\frac{\mathbf{E}\left\{r(X) I_{\left\{X \in A_{n}(x)\right\}} \mid D_{n}\right\}}{\mu\left(A_{n}(x)\right)}
$$

Note that both $\tilde{\eta}_{n}$ and $\tilde{r}_{n}$ are piecewise constant on the cells of the random partition $\mathcal{P}_{n}$. We may decompose the error as follows:

$$
\begin{align*}
& \left|\eta(x)-\eta_{n}(x)\right|  \tag{21.1}\\
& \quad \leq|\eta(x)-r(x)|+\left|r(x)-\tilde{r}_{n}(x)\right|+\left|\tilde{r}_{n}(x)-\tilde{\eta}_{n}(x)\right|+\left|\tilde{\eta}_{n}(x)-\eta_{n}(x)\right|
\end{align*}
$$

The integral of the first term on the right-hand side is smaller than $\epsilon$ by the definition of $r$. For the third term we have

$$
\begin{aligned}
& \int_{S_{M}}\left|\tilde{r}_{n}(x)-\tilde{\eta}_{n}(x)\right| \mu(d x) \\
& \quad=\sum_{A \in \mathcal{P}_{h}^{(M)}}\left|\mathbf{E}\left\{Y I_{\{X \in A\}} \mid D_{n}\right\}-\mathbf{E}\left\{r(X) I_{\{X \in A\}} \mid D_{n}\right\}\right| \\
& \quad=\sum_{A \in \mathcal{P}_{n}^{(M)}}\left|\int_{A} \eta(x) \mu(d x)-\int_{A} r(x) \mu(d x)\right| \\
& \quad \leq \int_{S_{M}}|\eta(x)-r(x)| \mu(d x)<\epsilon
\end{aligned}
$$

Consider the fourth term on the right-hand side of (21.1). Clearly,

$$
\begin{aligned}
& \int_{S_{M}}\left|\tilde{\eta}_{n}(x)-\eta_{n}(x)\right| \mu(d x) \\
&=\sum_{A \in \mathcal{P}_{n}^{(M)}}\left|\frac{1}{n} \sum_{i=1}^{n} Y_{i} I_{\left\{X_{i} \in A\right\}}-\mathbf{E}\left\{Y I_{\{X \in A\}} \mid D_{n}\right\}\right| \\
& \leq \sup _{\mathcal{P}^{(M)} \in \mathcal{F}_{n}^{(M)}} \sum_{A \in \mathcal{P}^{(M)}}\left|\frac{1}{n} \sum_{i=1}^{n} Y_{i} I_{\left\{X_{i} \in A\right\}}-\mathbf{E}\left\{Y I_{\{X \in A\}}\right\}\right|
\end{aligned}
$$

It follows from the first condition of the theorem and Corollary 21.1 of Theorem 21.1 that

$$
\int_{S_{M}}\left|\tilde{\eta}_{n}(x)-\eta_{n}(x)\right| \mu(d x) \rightarrow 0 \text { with probability one. }
$$

Finally, we consider the second term on the right-hand side of (21.1). Using Fubini's theorem we have

$$
\begin{aligned}
& \int_{S_{M}}\left|r(x)-\tilde{r}_{n}(x)\right| \mu(d x) \\
&= \sum_{A \in \mathcal{P}_{n}^{(M)}: \mu(A) \neq 0} \int_{A}\left|r(x)-\frac{\mathbf{E}\left\{r(X) I_{\{X \in A\}} \mid D_{n}\right\}}{\mu(A)}\right| \mu(d x) \\
&= \sum_{A \in \mathcal{P}_{n}^{(M)}: \mu(A) \neq 0} \frac{1}{\mu(A)} \int_{A}\left|r(x) \mu(A)-\mathbf{E}\left\{r(X) I_{\{X \in A\}} \mid D_{n}\right\}\right| \mu(d x) \\
&=\sum_{A \in \mathcal{P}_{n}^{(M)}: \mu(A) \neq 0} \frac{1}{\mu(A)} \int_{A}\left|r(x) \int_{A} \mu(d y)-\int_{A} r(y) \mu(d y)\right| \mu(d x) \\
& \leq \sum_{A \in \mathcal{P}_{n}^{(M)}: \mu(A) \neq 0} \frac{1}{\mu(A)} \int_{A} \int_{A}|r(x)-r(y)| \mu(d x) \mu(d y) .
\end{aligned}
$$

Fix $\delta \in(0,1)$. As $r$ is uniformly continuous, there exists a number $\gamma>0$ such that if $\operatorname{diam}(A)<\gamma$, then $|r(x)-r(y)|<\delta$ for every $x, y \in A$. In addition, there is a constant $M<\infty$ such that $|r(x)| \leq M$ for every $x \in \mathcal{R}^{d}$. Fix $n$ now and consider the integrals

$$
\frac{1}{\mu(A)} \int_{A} \int_{A}|r(x)-r(y)| \mu(d x) \mu(d y)
$$

appearing in the sum above. We always have the upper bound

$$
\frac{1}{\mu(A)} \int_{A} \int_{A}|r(x)-r(y)| \mu(d x) \mu(d y) \leq 2 M \mu(A)
$$

Assume now that $A \in \mathcal{P}_{n}^{(M)}$ has $\operatorname{diam}(A)<\gamma$. Then we can write

$$
\frac{1}{\mu(A)} \int_{A} \int_{A}|r(x)-r(y)| \mu(d x) \mu(d y) \leq \delta \mu(A)
$$

Summing over the cells $A \in \mathcal{P}_{n}^{(M)}$ with $\mu(A)>0$, these bounds give

$$
\begin{aligned}
& \int_{S_{M}}\left|r(x)-\tilde{r}_{n}(x)\right| \mu(d x) \\
& \quad \leq \sum_{A \in \mathcal{P}_{n}^{(M)}: \operatorname{diam}(A) \geq \gamma} 2 M \mu(A)+\sum_{A \in \mathcal{P}_{n}^{(M)}: \operatorname{diam}(A)<\gamma} \delta \mu(A) \\
& \leq 2 M \mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x)\right\}\right) \geq \gamma\right)+\delta .
\end{aligned}
$$

Letting $n$ tend to infinity gives

$$
\limsup _{n \rightarrow \infty} \int_{S_{M}}\left|r(x)-\tilde{r}_{n}(x)\right| \mu(d x) \leq \delta \quad \text { with probability one }
$$

by the second condition of the theorem. Summarizing,

$$
\limsup _{n \rightarrow \infty} \int_{S_{M}}\left|\eta(x)-\eta_{n}(x)\right| \mu(d x) \leq 2 \epsilon+\delta \text { with probability one. }
$$

Since $\epsilon$ and $\delta$ are arbitrary, the theorem is proved.

### 21.4 Statistically Equivalent Blocks

In this section we apply Theorem 21.2 both to classifiers based on uniform spacings in one dimension, and to their extension to multidimensional problems. We refer to these as rules based on statistically equivalent blocks. The order statistics of the components of the training data are used to construct a partition into rectangles. All
such classifiers are invariant with respect to all strictly monotone transformations of the coordinate axes. The simplest such rule is the $k$-spacing rule studied in Chapter 20 (see Theorem 20.1). Generalizations are possible in several ways. Theorem 21.2 allows us to handle partitions depending on the whole data sequence-and not only on the $X_{i}$ 's. The next simple result is sometimes useful.

Theorem 21.3. Consider a data-dependent partitioning classifier on the real line that partitions $\mathcal{R}$ into intervals in such a way that each interval contains at least $a_{n}$ and at most $b_{n}$ points. Assume that $X$ has a nonatomic distribution. Then the classifier is strongly consistent whenever $a_{n} \rightarrow \infty$ and $b_{n} / n \rightarrow 0$ as $n \rightarrow \infty$.

Proof. We check the conditions of Theorem 21.2. Let $\mathcal{F}_{n}$ contain all partitions of $\mathcal{R}$ into $m=\left\lceil n / a_{n}\right\rceil$ intervals. Since for each $M$, all partitions in $\mathcal{F}_{n}^{(M)}$ have at most $m$ cells, we can bound $\Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right)$ according to the Lemma 21.1. By the lemma, $\Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right)$ does not exceed $2^{m}$ times the number of different ways $n$ points can be partitioned into $m$ intervals. A little thought confirms that this number is

$$
\binom{n+m}{n}
$$

and therefore,

$$
\Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right) \leq 2^{m}\binom{n+m}{n}
$$

Let $\mathcal{H}$ denote the binary entropy function, $\mathcal{H}(x)=-x \log (x)-(1-x) \log (1-x)$ for $x \in(0,1)$. Note that $\mathcal{H}$ is symmetric about $1 / 2$ and that $\mathcal{H}$ is increasing for $0<x \leq 1 / 2$. By the inequality of Theorem $13.4, \log \binom{s}{t} \leq s \mathcal{H}(t / s)$. Therefore, it is easy to see that

$$
\begin{aligned}
\log \Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right) & \leq m+(n+m) \mathcal{H}\left(\frac{m}{n+m}\right) \\
& \leq n / a_{n}+2 n \mathcal{H}\left(1 / a_{n}\right)+1
\end{aligned}
$$

As $\mathcal{H}(x) \rightarrow 0$ when $x \rightarrow 0$, the condition $a_{n} \rightarrow \infty$ implies that

$$
\frac{1}{n} \log \Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right) \rightarrow 0
$$

which establishes condition (i).
To establish condition (ii) of Theorem 21.2, we proceed similarly to the proof of Theorem 20.1. Fix $\gamma, \epsilon>0$. There exists an interval $[-M, M]$ such that $\mu\left([-M, M]^{c}\right)<\epsilon$, and consequently

$$
\mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x)\right)>\gamma\right\}\right) \leq \epsilon+\mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x)\right)>\gamma\right\} \cap[-M, M]\right)
$$

where $A_{n}(x)$ denotes the cell of the partition $\mathcal{P}_{n}$ containing $x$. Among the intervals of $\mathcal{P}_{n}$, there can be at most $2+2 M / \gamma$ disjoint intervals of length greater than $\gamma$ in
$[-M, M]$. Thus we may bound the second term on the right-hand side above by

$$
\begin{aligned}
& \mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x)\right)>\gamma\right\} \cap[-M, M]\right) \\
& \leq\left(2+\frac{2 M}{\gamma}\right) \max _{A \in \mathcal{P}_{n}} \mu(A) \\
& \leq\left(2+\frac{2 M}{\gamma}\right)\left(\max _{A \in \mathcal{P}_{n}} \mu_{n}(A)+\max _{A \in \mathcal{P}_{n}}\left|\mu(A)-\mu_{n}(A)\right|\right) \\
& \quad \leq\left(2+\frac{2 M}{\gamma}\right)\left(\frac{b_{n}}{n}+\sup _{A \in \mathcal{A}}\left|\mu(A)-\mu_{n}(A)\right|\right),
\end{aligned}
$$

where $\mathcal{A}$ is the set of all intervals in $\mathcal{R}$. The first term in the parenthesis converges to zero by the second condition of the theorem, while the second term goes to zero with probability one by an obvious extension of the classical Glivenko-Cantelli theorem (Theorem 12.4). Summarizing, we have shown that for any $\gamma, \epsilon>0$

$$
\limsup _{n \rightarrow \infty} \mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x)\right)>\gamma\right\}\right) \leq \epsilon \quad \text { with probability one. }
$$

This completes the proof.

The $d$-dimensional generalizations of $k$-spacing rules include rules based upon statistically equivalent blocks, that is, partitions with sets that contain $k$ points each. It is obvious that one can proceed in many ways, see, for example, Anderson (1966), Patrick (1966), Patrick and Fisher (1967), Quesenberry and Gessaman (1968) and Gessaman and Gessaman (1972).

As a first example, consider the following algorithm: the $k$-th smallest $x^{(1)}$. coordinate among the training data defines the first cut. The (infinite) rectangle with $n-k$ points is cut according to the $x^{(2)}$-axis, isolating another $k$ points. This can be repeated on a rotational basis for all coordinate axes. Unfortunately, the classifier obtained this way is not consistent. To see this, observe that if $k$ is much smaller than $n$-a clearly necessary requirement for consistency-then almost all cells produced by the cuts are long and thin. We sketched a distribution in Figure 21.3 for which the error probability of this classifier fails to converge to $L^{*}$. The details are left to the reader (Problem 21.3). This example highlights the importance of condition (ii) of Theorem 21.2, that is, that the diameters of the cells should shrink in some sense as $n \rightarrow \infty$.


Figure 21.3. A nonconsistent $k$-block algorithm (with $k=2$ in the picture). $\eta(x)=1$ in the shaded area and $\eta(x)=0$ in the white area.

Rules have been developed in which the rectangular partition depends not only upon the $X_{i}$ 's in the training sequence, but also upon the $Y_{i}$ 's (see, e.g., Henrichon and Fu (1969), Meisel and Michalopoulos (1973) and Friedman (1977)). For example, Friedman cuts the axes at the places where the absolute differences between the marginal empirical distribution functions are largest, to insure minimal empirical error after the cut. His procedure is based upon Stoller splits (see Chapters 4 and 20).

Rules depending on the coordinatewise ranks of data points are interesting because they are invariant under monotone transformations of the coordinate axes. This is particularly important in practice when the components are not physically comparable. "Distribution-free" is an adjective often used to point out a property that is universally valid. For such methods in statistics, see the survey of Das Gupta (1964). Statistically equivalent sets in partitions are called "distributionfree" because the measure $\mu(A)$ of a set in the partition does not depend upon the distribution of $X$. We already noted a similar distribution-free behavior for k-d trees and median trees (Chapter 20). There is no reason to stay with rectangular-shaped sets (Anderson and Benning (1970), Beakley and Tuteur (1972)) but doing so greatly simplifies the interpretation of a classifier. In this book, to avoid confusion, we reserve the term "distribution-free" for consistency results or other theoretical properties that hold for all distributions of $(X, Y)$.

It is possible to define consistent partitions that have statistically equivalent sets. To fix the ideas, we take Gessaman's rule (1970) as our prototype rule for further study (note: for hypothesis testing, this partition was already noted by Anderson (1966)). For each $n$, let $m=\left\lceil\left(n / k_{n}\right)^{1 / d}\right\rceil$. Project the vectors $X_{1}, \ldots, X_{n}$ onto the first coordinate axis, and then partition the data into $m$ sets using hyperplanes


Figure 21.4. Gessaman's partition with $m=3$.
perpendicular to that axis, in such a way that each set contains an equal number of points (except, possibly, the rightmost set, where fewer points may fall if $n$ is not a multiple of $m$ ). We obtain $m$ cylindrical sets. In the same fashion, cut each of these cylindrical sets, along the second axis, into $m$ boxes such that each box contains the same number of data points. Continuing in the same way along the remaining coordinate axes, we obtain $m^{d}$ rectangular cells, each of which (with the exception of those on the boundary) contains about $k_{n}$ points. The classification rule $g_{n}$ uses a majority vote among those $Y_{i}$ 's for which $X_{i}$ lies within a given cell. Consistency of this classification rule can be established by an argument similar to that used for the $k_{n}$-spacing rule above. One needs to check that the conditions of Theorem 21.2 are satisfied. The only minor difference appears in the computation of $\Delta_{n}$, which in this case is bounded from above by $2^{m^{d}}\binom{n+m}{n}^{d}$. The following theorem summarizes the result:

Theorem 21.4. Assume that the marginal distributions of $X$ in $\mathcal{R}^{d}$ are nonatomic. Then the partitioning classification rule based on Gessaman's rule is strongly consistent if $k_{n} \rightarrow \infty$ and $k_{n} / n \rightarrow 0$ as $n$ tends to infinity.

To consider distributions with possibly atomic marginals, the partitioning algorithm must be modified, since every atom has more than $k_{n}$ points falling on it for large $n$. With a proper modification, a strongly universally consistent rule can be obtained. We leave the details to the reader (Problem 21.4).

Remark. Consistency of Gessaman's classification scheme can also be derived from the results of Gordon and Olshen (1978) under the additional condition $k_{n} / \sqrt{n} \rightarrow \infty$. Results of Breiman, Friedman, Olshen, and Stone (1984) can be
used to improve this condition to $k_{n} / \log n \rightarrow \infty$. Theorem 21.4 guarantees consistency under the weakest possible condition $k_{n} \rightarrow \infty$.

### 21.5 Partitioning Rules Based on Clustering

Clustering is one of the most widely used methods in statistical data analysis. Typical clustering schemes divide the data into a finite number of disjoint groups by minimizing some empirical error measure, such as the average squared distance from cluster centers (see Hartigan (1975)). In this section we outline the application of our results to classification rules based on $k$-means clustering of unlabeled observations.

As a first step, we divide $X_{1}, \ldots, X_{n}$ into $k_{n}$ disjoint groups having cluster centers $a_{1}, \ldots, a_{k_{n}} \in \mathcal{R}^{d}$. The vectors $a_{1}, \ldots, a_{k_{n}}$ are chosen to minimize the empirical squared Euclidean distance error,

$$
e_{n}\left(b_{1}, \ldots, b_{k_{n}}\right)=\frac{1}{n} \sum_{i=1}^{n} \min _{1 \leq j \leq k_{n}}\left\|X_{i}-b_{j}\right\|^{2}
$$

over all the nearest-neighbor clustering rules having $k_{n}$ representatives $b_{1}, \ldots$, $b_{k_{n}} \in \mathcal{R}^{d}$, where $\|\cdot\|$ denotes the usual Euclidean norm. Note that the choice of cluster centers depends only on the vectors $X_{i}$, not on their labels. For the behavior of $e_{n}\left(a_{1}, \ldots, a_{k_{n}}\right)$, see Problem 29.4.

The vectors $a_{1}, \ldots, a_{k_{n}}$ give rise to a Voronoi partition $\mathcal{P}_{n}=\left\{A_{1}, \ldots, A_{k_{n}}\right\}$ in a natural way: for each $j \in\left\{1, \ldots, k_{n}\right\}$, let

$$
A_{j}=\left\{x:\left\|x-a_{j}\right\| \leq\left\|x-a_{i}\right\| \text { for all } i=1, \ldots, k_{n}\right\}
$$

Ties are broken by assigning points on the boundaries to the vector that has the smallest index.

The classification rule $g_{n}$ is defined in the usual way: $g_{n}(x)$ is a majority vote among those $Y_{j}$ 's such that $X_{j}$ falls in $A_{n}(x)$. If the measure $\mu$ of $X$ has a bounded support, Theorem 21.2 shows that the classification rule $\left\{g_{n}\right\}$ is strongly consistent if $k_{n}$ grows with $n$ at an appropriate rate. Note that this rule is just another of the prototype nearest neighbor rules that we discussed in Chapter 19.


Figure 21.5. Example of partitioning based on clustering with $k_{n}=3$. The criterion we minimize is the sum of the squares of the distances of the $X_{i}$ 's to the $a_{j}$ 's.

Theorem 21.5. (Lugosi and Nobel (1996)). Assume that there is a bounded set $A \subset \mathcal{R}^{d}$ such that $\mathbf{P}\{X \in A\}=1$. Let $\left\{k_{n}\right\}$ be a sequence of integers for which

$$
k_{n} \rightarrow \infty \quad \text { and } \frac{k_{n}^{2} \log n}{n} \rightarrow 0 \quad \text { as } n \rightarrow \infty
$$

Let $g_{n}\left(\cdot, D_{n}\right)$ be the histogram classification rule based on the Voronoi partition of $k_{n}$ cluster centers minimizing the empirical squared Euclidean distance error. Then

$$
L\left(g_{n}\right) \rightarrow L^{*}
$$

with probability one as $n$ tends to infinity. If $d=1$ or 2 , then the second condition on $k_{n}$ can be relaxed to

$$
\frac{k_{n} \log n}{n} \rightarrow 0
$$

Proof. Again, we check the conditions of Theorem 21.2. Let $\mathcal{F}_{n}$ consist of all Voronoi partitions of $k_{n}$ points in $\mathcal{R}^{d}$. As each partition consists of $k_{n}$ cells, we may use Lemma 21.1 to bound $\Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right)$. Clearly, boundaries between cells are subsets of hyperplanes. Since there are at most $k_{n}\left(k_{n}-1\right) / 2$ boundaries between the $k_{n}$ Voronoi cells, each cell of a partition in $\mathcal{F}_{n}$ is a polytope with at most $k_{n}\left(k_{n}-1\right) / 2<k_{n}^{2}$ faces. By Theorem 13.9, $n$ fixed points in $\mathcal{R}^{d}, d \geq 2$, can be split by hyperplanes in at most $n^{d+1}$ different ways. It follows that for each
$M<\infty, \Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right) \leq 2^{k_{n}} n^{(d+1) k_{n}^{2}}$, and consequently

$$
\frac{1}{n} \log \Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right) \leq \frac{k_{n}}{n}+\frac{(d+1) k_{n}^{2} \log n}{n} \rightarrow 0
$$

by the second condition on the sequence $\left\{k_{n}\right\}$. Thus condition (i) of Theorem 21.2 is satisfied.

It remains to establish condition (ii) of Theorem 21.2. This time we need the weaker condition mentioned in the remark after the theorem, that is, that for every $\gamma>0$ and $\delta \in(0,1)$

$$
\inf _{T: \mu(T) \geq 1-\delta} \mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x) \cap T\right)>\gamma\right\}\right) \rightarrow 0 \quad \text { with probability one. }
$$

Clearly, we are done if we can prove that there is a sequence of subsets $T_{n}$ of $\mathcal{R}^{d}$ (possibly depending on the data $D_{n}$ ) such that $\mu\left(T_{n}\right) \rightarrow 1$ with probability one, and

$$
\mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x) \cap T_{n}\right)>\gamma\right\}\right)=0 .
$$

To this end, let $a_{1}, \ldots, a_{k_{n}}$ denote the optimal cluster centers corresponding to $D_{n}$, and define

$$
T_{n}=\bigcup_{j=1}^{k_{n}} S_{a_{j}, \gamma / 2} \cap A_{j}
$$

where $S_{x, r}$ is the ball of radius $r$ around the vector $x$. Clearly, $x \in T_{n}$ implies that $\|x-a(x)\|<\gamma / 2$, where $a(x)$ denotes the closest cluster center to $x$ among $a_{1}, \ldots, a_{k_{n}}$. But since

$$
A_{n}(x) \cap T_{n}=A_{n}(x) \cap S_{a(x), \gamma / 2}
$$

it follows that

$$
\operatorname{diam}\left(A_{n}(x) \cap T_{n}\right) \leq \operatorname{diam}\left(S_{a(x), \gamma / 2}\right)=\gamma
$$

and

$$
\mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x) \cap T_{n}\right)>\gamma\right\}\right)=0 .
$$

It remains to show that $\mu\left(T_{n}\right) \rightarrow 1$ with probability one as $n \rightarrow \infty$. Using Markov's inequality, we may write

$$
\begin{aligned}
1-\mu\left(T_{n}\right) & =\mathbf{P}\left\{X \notin T_{n} \mid D_{n}\right\} \\
& \leq \mathbf{P}\left\{\left.\min _{1 \leq j \leq k_{n}}\left\|X-a_{j}\right\|^{2}>\left(\frac{\gamma}{2}\right)^{2} \right\rvert\, D_{n}\right\} \\
& \leq \frac{\mathbf{E}\left\{\min _{1 \leq j \leq k_{n}}\left\|X-a_{j}\right\|^{2} \mid D_{n}\right\}}{(\gamma / 2)^{2}}
\end{aligned}
$$

Using a large-deviation inequality for the empirical squared error of nearestneighbor clustering schemes, it can be shown (see Problem 29.4) that if $X$ has bounded support, then

$$
\mathbf{E}\left\{\min _{1 \leq j \leq k_{n}}\left\|X-a_{j}\right\|^{2} \mid D_{n}\right\}-\min _{b_{1}, \ldots, b_{k_{n}} \in \mathcal{R}^{d}} \mathbf{E}\left\{\min _{1 \leq j \leq k_{n}}\left\|X-b_{j}\right\|^{2}\right\} \rightarrow 0
$$

with probability one if $k_{n} \log n / n \rightarrow 0$ as $n \rightarrow \infty$. Moreover, it is easy to see that

$$
\min _{b_{1}, \ldots, b_{k_{n}} \in \mathcal{R}^{d}} \mathbf{E}\left\{\min _{1 \leq j \leq k_{n}}\left\|X-b_{j}\right\|^{2}\right\} \rightarrow 0
$$

as $k_{n} \rightarrow \infty$. It follows that $\mu\left(T_{n}\right) \rightarrow 1$, as desired.
If $d=1$, then the cells of the Voronoi partition are intervals on the real line, and therefore, $\Delta_{n}\left(\mathcal{F}_{n}^{(M)}\right) \leq 2^{k_{n}} n^{k_{n}}$. Similarly, if $d=2$, then the number of hyperplanes defining the Voronoi partition increases linearly with $k_{n}$. To see this, observe that if we connect centers of neighboring clusters by edges, then we obtain a planar graph. From Euler's theorem (see, e.g., Edelsbrunner (1987, p.242)), the number of edges in a planar graph is bounded by $3 N-6, N \geq 3$, where $N$ is the number of vertices. Thus, in order to satisfy condition (i), it suffices that $k_{n} \log n / n \rightarrow 0$ in both cases.

Remark. In the theorem above we assume that the cluster centers $a_{1}, \ldots, a_{k}$ are empirically optimal in the sense that they are chosen to minimize the empirical squared Euclidean distance error

$$
\frac{1}{n} \sum_{i=1}^{n} \min _{1 \leq j \leq k_{n}}\left\|X_{i}-a_{j}\right\|^{2}
$$

Practically speaking, it is hard to determine minimum. To get around this difficulty, several fast algorithms have been proposed that approximate the optimum (see Hartigan (1975) for a survey). Perhaps the most popular algorithm is the so-called $k$-means clustering method, also known in the theory of quantization as the LloydMax algorithm (Lloyd (1982), Max (1960), Linde, Buzo, and Gray (1980)). The iterative method works as follows:

STEP 1. Take $k$ initial centers $a_{1}^{(0)}, \ldots, a_{k}^{(0)} \in \mathcal{R}^{d}$, and set $i=0$.
STEP 2. Cluster the data points $X_{1}, \ldots, X_{n}$ around the centers $a_{1}^{(i)}, \ldots, a_{k}^{(i)}$ into $k$ sets such that the $m$-th set $C_{m}^{(i)}$ contains the $X_{j}$ 's that are closer to $a_{m}^{(i)}$ than to any other center. Ties are broken in favor of smaller indices.
STEP 3. Determine the new centers $a_{1}^{(i+1)}, \ldots, a_{k}^{(i+1)}$ as the averages of the data points within the clusters:

$$
a_{m}^{(i+1)}=\frac{\sum_{j: X_{j} \in C_{C}^{(i)}} X_{j}}{\left|C_{m}^{(i)}\right|} .
$$

StEP 4. Increase $i$ by one, and repeat Steps 1 and 2 until there are no changes in the cluster centers.

It is easy to see that each step of the algorithm decreases the empirical squared Euclidean distance error. On the other hand, the empirical squared error can take finitely many different values during the execution of the algorithm. Therefore
the algorithm halts in finite time. By inspecting the proof of Theorem 21.5, it is not hard to see that consistency can also be proved for partitions given by the suboptimal cluster centers obtained by the $k$-means method, provided that it is initialized appropriately (see Problem 21.6).

### 21.6 Data-Based Scaling

We now choose the grid size $h$ in a cubic histogram rule in a data-dependent manner and denote its value by $H_{n}$. Theorem 21.2 implies the following general result:

Theorem 21.6. Let $g_{n}$ be the cubic histogram classifier based on a partition into cubes of size $H_{n}$. If
(a) $\lim _{n \rightarrow \infty} H_{n}=0$ and
(b) $\lim _{n \rightarrow \infty} n H_{n}^{d}=\infty$ with probability one,
then the partitioning rule is strongly universally consistent.
To prove the theorem, we need the following auxiliary result:
Lemma 21.2. Let $Z_{1}, Z_{2}, \ldots$ be a sequence of nonnegative random variables. If $\lim _{n \rightarrow \infty} Z_{n}=0$ with probability one, then there exists a sequence $a_{n} \downarrow 0$ of positive numbers such that $\lim _{n \rightarrow \infty} I_{\left\{Z_{n} \geq a_{n}\right\}}=0$ with probability one.

Proof. Define $V_{n}=\sup _{m \geq n} Z_{m}$. Then clearly, $V_{n} \downarrow 0$ with probability one. We can find a subsequence $n_{1}, n_{2}, \ldots$ of positive integers such that for each $k$,

$$
\mathbf{P}\left\{V_{n_{k}} \geq 1 / k\right\} \leq 2^{-k}
$$

Then the Borel-Cantelli lemma implies that

$$
\begin{equation*}
\lim _{k \rightarrow \infty} I_{\left\{V_{n_{k}} \geq 1 / k\right\}}=0 \text { with probability one. } \tag{21.2}
\end{equation*}
$$

Define $a_{n}=1 / k$ for $n_{k} \leq n<n_{k+1}$. Then, for $n_{k} \leq n<n_{k+1}$,

$$
I_{\left\{V_{n} \geq a_{n}\right\}}=I_{\left\{V_{n} \geq 1 / k\right\}} \leq I_{\left\{V_{n_{k}} \geq 1 / k\right\}}
$$

The fact that $V_{n} \geq Z_{n}$ and (21.2) imply the statement.

Proof of Theorem 21.6. Let $\left\{a_{n}\right\}$ and $\left\{b_{n}\right\}$ be sequences of positive numbers with $a_{n}<b_{n}$. Then

$$
\left\{\lim _{n \rightarrow \infty} L\left(g_{n}\right)=L^{*}\right\} \subseteq\left\{\lim _{n \rightarrow \infty} L\left(g_{n}\right)=L^{*}, H_{n} \in\left[a_{n}, b_{n}\right]\right\} \bigcup\left\{H_{n} \notin\left[a_{n}, b_{n}\right]\right\}
$$

It follows from Lemma 21.2 that there exist sequences of positive numbers $\left\{a_{n}\right\}$ and $\left\{b_{n}\right\}$ satisfying $a_{n}<b_{n}, b_{n} \rightarrow 0$, and $n a_{n}^{d} \rightarrow \infty$ as $n \rightarrow \infty$ such that

$$
\lim _{n \rightarrow \infty} I_{\left\{H_{n} \notin\left[a_{n}, b_{n}\right]\right\}}=0 \quad \text { with probability one. }
$$

Therefore we may assume that for each $n, \mathbf{P}\left\{H_{n} \in\left[a_{n}, b_{n}\right]\right\}=1$. Since $b_{n} \rightarrow 0$, condition (ii) of Theorem 21.2 holds trivially, as all diameters of all cells are inferior to $b_{n} \sqrt{d}$. It remains to check condition (i). Clearly, for each $M<\infty$, each partition in $\mathcal{F}_{n}^{(M)}$ contains less than $c_{0}\left(1 / a_{n}\right)^{d}$ cells, where the constant $c_{0}$ depends on $M$ and $d$ only. On the other hand, it is easy to see that $n$ points can not be partitioned more than $c_{1} n\left(1 / a_{n}\right)^{d}$ different ways by cubic-grid partitions with cube size $h \geq a_{n}$ for some other constant $c_{1}$. Therefore, for each $M<\infty$,

$$
\Delta_{n}\left(\mathcal{F}^{(M)}\right) \leq 2^{c_{0}\left(1 / a_{n}\right)^{d}} \frac{c_{1} n}{a_{n}^{d}},
$$

and condition (i) is satisfied.

In many applications, different components of the feature vector $X$ correspond to different physical measurements. For example, in a medical application, the first coordinate could represent blood pressure, the second cholesterol level, and the third the weight of the patient. In such cases there is no reason to use cubic histograms, because the resolution of the partition $h_{n}$ along the coordinate axes depends on the apparent scaling of the measurements, which is rather arbitrary. Then one can use scale-independent partitions such as methods based on order statistics described earlier. Alternatively, one might use rectangular partitions instead of cubic ones, and let the data decide the scaling along the different coordinate axes. Again, Theorem 21.2 can be used to establish conditions of universal consistency of the classification rule corresponding to data-based rectangular partitions:

Theorem 21.7. Consider a data-dependent histogram rule when the cells of the partition are all rectangles of the form

$$
\left[k_{1} H_{n 1},\left(k_{1}+1\right) H_{n 1}\right) \times \ldots \times\left[k_{d} H_{n d},\left(k_{d}+1\right) H_{n d}\right)
$$

where $k_{1}, \ldots, k_{d}$ run through the set of integers, and the edge sizes of the rectangles $H_{n 1}, \ldots, H_{n d}$ are determined from the data $D_{n}$. If as $n \rightarrow \infty$
$H_{n i} \rightarrow 0$ for each $1 \leq i \leq d$, and $n H_{n 1} \cdots H_{n d} \rightarrow \infty$ with probability one, then the data-dependent rectangular partitioning rule is strongly universally consistent.

To prove this, just check the conditions of Theorem 21.2 (Problem 21.7). We may pick, for example, $H_{n 1}, \ldots, H_{n d}$ to minimize the resubstitution estimate

$$
\sum_{i=1}^{n} I_{\left\{g_{n}\left(X_{i}\right) \neq Y_{i}\right\}}
$$

subject of course to certain conditions, so that $\sum_{i=1}^{d} H_{n i} \rightarrow 0$ with probability one, yet $n \prod_{i=1}^{d} H_{n i} \rightarrow \infty$ with probability one. See Problem 21.8.

### 21.7 Classification Trees

Consider a partition of the space obtained by a binary classification tree in which each node dichotomizes its set by a hyperplane (see Chapter 20 for more on classification trees). The construction of the tree is stopped according to some unspecified rule, and classification is by majority vote over the convex polytopes of the partition.

The following corollary of Theorem 21.2 generalizes (somewhat) the consistency results in the book by Breiman, Friedman, Olshen, and Stone (1984):

Theorem 21.8. (Lugosi and Nobel (1996)). Let $g_{n}$ be a binary tree classifier based upon at most $m_{n}-1$ hyperplane splits, where $m_{n}=o(n / \log n)$. If, in addition, condition (ii) of Theorem 21.2 is satisfied, then $g_{n}$ is strongly consistent. In particular, the rule is strongly consistent if condition (ii) of Theorem 21.2 holds and every cell of the partition contains at least $k_{n}$ points, where $k_{n} / \log n \rightarrow \infty$.

Proof. To check condition (i) of Theorem 21.2, recall Theorem 13.9 which implies that $n \geq 2$ points in a $d$-dimensional Euclidean space can be dichotomized by hyperplanes in at most $n^{d+1}$ different ways. From this, we see that the number of different ways $n$ points of $\mathcal{R}^{d}$ can be partitioned by the rule $g_{n}$ can be bounded by

$$
\left(n^{d+1}\right)^{m_{n}}=2^{(d+1) m_{n} \log n},
$$

as there are not more than $m_{n}$ cells in the partition. Thus,

$$
\Delta_{n}\left(\mathcal{F}_{n}\right) \leq 2^{m_{n}} 2^{(d+1) m_{n} \log n}
$$

By the assumption that $m_{n} \log n / n \rightarrow 0$ we have

$$
\frac{1}{n} \log \Delta_{n}\left(\mathcal{F}_{n}\right) \leq \frac{m_{n}}{n}+\frac{m_{n}(d+1) \log n}{n} \rightarrow 0,
$$

so that condition (i) of Theorem 21.2 is satisfied.
For the second part of the statement, observe that there are no more than $n / k_{n}$ cells in any partition, and that the tree-structured nature of the partitions assures that $g_{n}$ is based on at most $n / k_{n}$ hyperplane splits. This completes the proof.

## Problems and Exercises

Problem 21.1. Let $\mathcal{P}$ be a partition of $\mathcal{R}^{d}$. Prove that

$$
\sum_{A \in \mathcal{P}}\left|\mu_{n}(A)-\mu(A)\right|=2 \sup _{B \in \mathcal{B}(\mathcal{P})}\left|\mu_{n}(B)-\mu(B)\right|,
$$

where the class of sets $\mathcal{B}(\mathcal{P})$ contains all sets obtained by unions of cells of $\mathcal{P}$. This is Scheffé's (1947) theorem for partitions. See also Problem 12.13.

Problem 21.2. Show that condition (ii) of Theorem 21.2 may be replaced by the following: for every $\gamma>0$ and $\delta \in(0,1)$

$$
\lim _{n \rightarrow \infty} \inf _{T \subset \mathcal{R}^{d}: \mu(T) \geq 1-\delta} \mu\left(\left\{x: \operatorname{diam}\left(A_{n}(x) \cap T\right)>\gamma\right\}\right)=0 \quad \text { with probability one. }
$$

Problem 21.3. Let $X$ be uniformly distributed on the unit square $[0,1]^{2}$. Let $\eta(x)=1$ if $x^{(1)} \leq 2 / 3$, and $\eta(x)=0$ otherwise (see Figure 21.3). Consider the algorithm when first the $x^{(1)}$-coordinate is cut at the $k$-th smallest $x^{(1)}$-value among the training data. Next the rectangle with $n-k$ points is cut according to the $x^{(2)}$-axis, isolating another $k$ points. This is repeated on a rotational basis for the two coordinate axes. Show that the error probability of the obtained partitioning classifier does not tend to $L^{*}=0$. Can you determine the asymptotic error probability?

Problem 21.4. Modify Gessaman's rule based on statistically equivalent blocks so that the rule is strongly universally consistent.

Problem 21.5. Cut each axis independently into intervals containing exactly $k$ of the (projected) data points. The $i$-th axis has intervals $A_{1, i}, A_{2, i}, \ldots$. Form a histogram rule that takes a majority vote over the product sets $A_{i_{1}, 1} \times \cdots \times A_{i_{d}, d}$.


Figure 21.6. A partition based upon the method obtained above with $k=6$.

This rule does not guarantee a minimal number of points in every cell. Nevertheless, if $k^{d}=o(n), k \rightarrow \infty$, show that this decision rule is consistent, i.e., that $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ in probability.

Problem 21.6. Show that each step of the $k$-means clustering algorithm decreases the empirical squared error. Conclude that Theorem 21.5 is also true if the clusters are given by the $k$-means algorithm. Hint: Observe that the only property of the clusters used in the proof of Theorem 21.5 is that

$$
\mathbf{E}\left\{\min _{1 \leq j \leq k_{n}}\left\|X-a_{j}\right\|^{2} \mid D_{n}\right\} \rightarrow 0
$$

with probability one. This can be proven for clusters given by the $k$-means algorithm if it is appropriately initialized. To this end, use the techniques of Problem 29.4.
Problem 21.7. Prove Theorem 21.7.
Problem 21.8. Consider the $H_{n i}$ 's in Theorem 21.7, the interval sizes for cubic histograms.
Let the $H_{n i}$ 's be found by minimizing

$$
\sum_{i=1}^{n} I_{\left\{g_{n}\left(X_{i}\right) \neq Y_{i}\right\}}
$$

subject to the condition that each marginal interval contain at least $k$ data points (that is, at least $k$ data points have that one coordinate in the interval). Under which condition on $k$ is the rule consistent?

Problem 21.9. Take a histogram rule with data-dependent sizes $H_{n 1}, \ldots, H_{n d}$ as in Theorem 21.7, defined as follows:
(1) $\quad H_{n i}=\left(\max _{1 \leq j \leq n} X_{j}^{(i)}-\min _{1 \leq j \leq n} X_{j}^{(i)}\right) / \sqrt{n}, 1 \leq i \leq d\left(\right.$ where $X_{j}=\left(X_{j}^{(1)}, \ldots, X_{j}^{(d)}\right)$;
(2) $H_{n i}=\left(W_{n i}-V_{n i}\right) / n^{1 /(2 d)}, 1 \leq i \leq d$, where $V_{n i}$ and $W_{n i}$ are 25 and 75 percentiles of $X_{1}^{(i)}, \ldots, X_{n}^{(i)}$.
Assume for convenience that $X$ has nonatomic marginals. Show that (1) leads sometimes to an inconsistent rule, even if $d=1$. Show that (2) always yields a scale-invariant consistent rule.

## 22

## Splitting the Data

### 22.1 The Holdout Estimate

Universal consistency gives us a partial satisfaction-without knowing the underlying distribution, taking more samples is guaranteed to push us close to the Bayes rule in the long run. Unfortunately, we will never know just how close we are to the Bayes rule unless we are given more information about the unknown distribution (see Chapter 7). A more modest goal is to do as well as possible within a given class of rules. To fix the ideas, consider all nearest neighbor rules based upon metrics of the form

$$
\|x\|^{2}=\sum_{i=1}^{d} a_{i} x^{(i)^{2}}
$$

where $a_{i} \geq 0$ for all $i$ and $x=\left(x^{(1)}, \ldots, x^{(d)}\right)$. Here the $a_{i}$ 's are variable scale factors. Let $\phi_{n}$ be a particular nearest neighbor rule for a given choice of $\left(a_{1}, \ldots, a_{d}\right)$, and let $g_{n}$ be a data-based rule chosen from this class. The best we can hope for now is something like

$$
\frac{L\left(g_{n}\right)}{\inf _{\phi_{n}} L\left(\phi_{n}\right)} \rightarrow 1 \quad \text { in probability }
$$

for all distributions, where $L\left(g_{n}\right)=\mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}$ is the conditional probability of error for $g_{n}$. This sort of optimality-within-a-class is definitely achievable. However, proving such optimality is generally not easy as $g_{n}$ depends on the data. In this chapter we present one possible methodology for selecting provably good rules from restricted classes. This is achieved by splitting the data into a training
sequence and a testing sequence. This idea was explored and analyzed in depth in Devroye (1988b) and is now formalized.

The data sequence $D_{n}$ is split into a training sequence $D_{m}=\left(X_{1}, Y_{1}\right), \ldots$, $\left(X_{m}, Y_{m}\right)$ and a testing sequence $T_{l}=\left(X_{m+1}, Y_{m+1}\right), \ldots,\left(X_{m+l}, Y_{m+l}\right)$, where $l+$ $m=n$. The sequence $D_{m}$ defines a class of classifiers $\mathcal{C}_{m}$ whose members are denoted by $\phi_{m}(\cdot)=\phi_{m}\left(\cdot, D_{m}\right)$. The testing sequence is used to select a classifier from $\mathcal{C}_{m}$ that minimizes the error count

$$
\widehat{L}_{m, l}\left(\phi_{m}\right)=\frac{1}{l} \sum_{i=1}^{l} I_{\left\{\phi_{m}\left(X_{m+i}\right) \neq Y_{m+i}\right\}} .
$$

This error estimate is called the holdout estimate, as the testing sequence is "held out" of the design of $\phi_{m}$. Thus, the selected classifier $g_{n} \in \mathcal{C}_{m}$ satisfies

$$
\widehat{L}_{m, l}\left(g_{n}\right) \leq \widehat{L}_{m, l}\left(\phi_{m}\right)
$$

for all $\phi_{m} \in \mathcal{C}_{m}$. The subscript $n$ in $g_{n}$ may be a little confusing, since $g_{n}$ is in $\mathcal{C}_{m}$, a class of classifiers depending on the first $m$ pairs $D_{m}$ only. However, $g_{n}$ depends on the entire data $D_{n}$, as the rest of the data is used for testing the classifiers in $\mathcal{C}_{m}$. We are interested in the difference between the error probability

$$
L\left(g_{n}\right)=\mathbf{P}\left\{g_{n}(X) \neq Y \mid D_{n}\right\}
$$

and that of the best classifier in $\mathcal{C}_{m}, \inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)$. Note that $L\left(\phi_{m}\right)=\mathbf{P}\left\{\phi_{m}(X) \neq\right.$ $\left.Y \mid D_{m}\right\}$ denotes the error probability conditioned on $D_{m}$. The conditional probability

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in C_{m}} L\left(\phi_{m}\right)>\epsilon \mid D_{m}\right\}
$$

is small when most testing sequences $T_{l}$ pick a rule $g_{n}$ whose error probability is within $\epsilon$ of the best classifier in $\mathcal{C}_{m}$. We have already addressed similar questions in Chapters 8 and 12. There we have seen (Lemma 8.2), that

$$
L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right) \leq 2 \sup _{\phi_{m} \in \mathcal{C}_{m}}\left|\widehat{L}_{m, l}\left(\phi_{m}\right)-L\left(\phi_{m}\right)\right| .
$$

If $\mathcal{C}_{m}$ contains finitely many rules, then the bound of Theorem 8.3 may be useful:

$$
\mathbf{P}\left\{\sup _{\phi_{m} \in \mathcal{C}_{m}}\left|\widehat{L}_{m, l}\left(\phi_{m}\right)-L\left(\phi_{m}\right)\right|>\epsilon \mid D_{m}\right\} \leq 2\left|\mathcal{C}_{m}\right| e^{-2 l \epsilon^{2}}
$$

If we take $m=l=n / 2$, then Theorem 8.3 shows (see Problem 12.1) that on the average we are within $\sqrt{\log \left(2 e\left|\mathcal{C}_{m}\right|\right) / n}$ of the best possible error rate, whatever it is.

If $\mathcal{C}_{m}$ is of infinite cardinality, then we can use the Vapnik-Chervonenkis theory to get similar inequalities. For example, from Theorem 12.8 we get

$$
\begin{equation*}
\mathbf{P}\left\{\sup _{\phi_{m} \in \mathcal{C}_{m}}\left|\widehat{L}_{m, l}\left(\phi_{m}\right)-L\left(\phi_{m}\right)\right|>\epsilon \mid D_{m}\right\} \leq 4 e^{8} \mathcal{S}\left(\mathcal{C}_{m}, l^{2}\right) e^{-2 l \epsilon^{2}} \tag{22.1}
\end{equation*}
$$

and consequently,

$$
\begin{equation*}
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)>\epsilon \mid D_{m}\right\} \leq 4 e^{8} \mathcal{S}\left(\mathcal{C}_{m}, l^{2}\right) e^{-l \epsilon^{2} / 2} \tag{22.2}
\end{equation*}
$$

where $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ is the $l$-th shatter coefficient corresponding to the class of classifiers $\mathcal{C}_{m}$ (see Theorem 12.6 for the definition). Since $\mathcal{C}_{m}$ depends on the training data $D_{m}$, the shatter coefficients $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ may depend on $D_{m}$, too. However, usually it is possible to find upper bounds on the random variable $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ that depend on $m$ and $l$ only, but not on the actual values of the random variables $X_{1}, Y_{1}, \ldots, X_{m}, Y_{m}$. Both upper bounds above are distribution-free, and the problem now is purely combinatorial: count $\left|\mathcal{C}_{m}\right|$ (this is usually trivial), or compute $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$.

REMARK. With much more effort it is possible to obtain performance bounds of the holdout estimate in the form of bounds for

$$
\mathbf{P}\left\{\widehat{L}_{m, l}\left(\phi_{m}\right)-L\left(\phi_{n}\right)>\epsilon\right\}
$$

for some special rules where $\phi_{m}$ and $\phi_{n}$ are carefully defined. For example, Devroye and Wagner (1979a) give upper bounds when both $\phi_{m}$ and $\phi_{n}$ are $k$-nearest neighbor classifiers with the same $k$ (but working with different sample size).

Remark. Minimizing the holdout estimate is not the only possibility. Other error estimates that do not split the data may be used in classifier selection as well. Such estimates are discussed in Chapters 23, 24, and 31. However, these estimates are usually tailored to work well for specific discrimination rules. The most general and robust method is certainly the data splitting described here.

### 22.2 Consistency and Asymptotic Optimality

Typically, $\mathcal{C}_{m}$ becomes richer as $m$ grows, and it is natural to ask whether the empirically best classifier in $\mathcal{C}_{m}$ is consistent.

Theorem 22.1. Assume that from each $\mathcal{C}_{m}$ we can pick one $\phi_{m}$ such that the sequence of $\phi_{m}$ 's is consistent for a certain class of distributions. Then the automatic rule $g_{n}$ defined above is consistent for the same class of distributions (i.e., $\mathbf{E} L\left(g_{n}\right) \rightarrow L^{*}$ as $\left.n \rightarrow \infty\right)$ if

$$
\lim _{n \rightarrow \infty} \frac{\log \left(\mathbf{E}\left\{\mathcal{S}\left(\mathcal{C}_{m}, l\right)\right\}\right)}{l}=0
$$

Proof. Decompose the difference between the actual error probability and the Bayes error as

$$
L\left(g_{n}\right)-L^{*}=\left(L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)\right)+\left(\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)-L^{*}\right) .
$$

The convergence of the first term on the right-hand side is a direct corollary of (22.2). The second term converges to zero by assumption.

Theorem 22.1 shows that a consistent rule is picked if the sequence of $\mathcal{C}_{m}$ 's contains a consistent rule, even if we do not know which functions from $\mathcal{C}_{m}$ lead to consistency. If we are just worried about consistency, Theorem 22.1 reassures us that nothing is lost as long as we take $l$ much larger than $\log \left(\mathbf{E}\left\{\mathcal{S}\left(\mathcal{C}_{m}, l\right)\right\}\right)$. Often, this reduces to a very weak condition on the size $l$ of the testing set.

Let us now introduce the notion of asymptotic optimality. A sequence of rules $g_{n}$ is said to be asymptotically optimal for a given distribution of $(X, Y)$ when

$$
\lim _{n \rightarrow \infty} \frac{\mathbf{E}\left\{L\left(g_{n}\right)\right\}-L^{*}}{\mathbf{E}\left\{\inf _{\phi_{m} \in \mathcal{C}_{n}} L\left(\phi_{m}\right)\right\}-L^{*}}=1 .
$$

Our definition is not entirely fair, because $g_{n}$ uses $n$ observations, whereas the family of rules in the denominator is restricted to using $m$ observations. If $g_{n}$ is not taken from the same $\mathcal{C}_{m}$, then it is possible to have a ratio which is smaller than one. But if $g_{n} \in \mathcal{C}_{m}$, then the ratio always is at least one. That is why the definition makes sense in our setup.

When our selected rule is asymptotically optimal, we have achieved something very strong: we have in effect picked a rule (or better, a sequence of rules) which has a probability of error converging at the optimal rate attainable within the sequence of $\mathcal{C}_{m}$ 's. And we do not even have to know what the optimal rate of convergence is. This is especially important in nonparametric rules, where some researchers choose smoothing factors based on theoretical results about the optimal attainable rate of convergence for certain classes of distributions.

We are constantly faced with the problem of choosing between parametric and nonparametric discriminators. Parametric discriminators are based upon an underlying model in which a finite number of unknown parameters is estimated from the data. A case in point is the multivariate normal distribution, which leads to linear or quadratic discriminators. If the model is wrong, parametric methods can perform very poorly; when the model is right, their performance is difficult to beat. The method based on splitting the data chooses among the best discriminator depending upon which happens to be best for the given data. We can throw in $\mathcal{C}_{m}$ a variety of rules, including nearest neighbor rules, a few linear discriminators, a couple of tree classifiers and perhaps a kernel rule. The probability bounds above can be used when the complexity of $\mathcal{C}_{m}$ (measured by its shatter coefficients) does not get out of hand.

The notion of asymptotic optimality can be too strong in many cases. The reason for this is that in some rare lucky situations $\mathbf{E}\left\{\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)\right\}-L^{*}$ may be very small. In these cases it is impossible to achieve asymptotic optimality. We can fix this problem by introducing the notion of $\epsilon_{m}$-optimality, where $\epsilon_{m}$ is a positive sequence decreasing to 0 with $m$. A rule is said to be $\epsilon_{m}$-optimal when

$$
\lim _{n \rightarrow \infty} \frac{\mathbf{E}\left\{L\left(g_{n}\right)\right\}-L^{*}}{\mathbf{E}\left\{\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)\right\}-L^{*}}=1
$$

for all distributions of $(X, Y)$ for which

$$
\lim _{m \rightarrow \infty} \frac{\mathbf{E}\left\{\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)\right\}-L^{*}}{\epsilon_{m}}=\infty .
$$

In what follows we apply the idea of data splitting to scaled nearest neighbor rules and to rules closely related to the data-dependent partitioning classifiers studied in Chapter 21. Many more examples are presented in Chapters 25 and 26.

### 22.3 Nearest Neighbor Rules with Automatic Scaling

Let us work out the simple example introduced above. The $\left(a_{1}, \ldots, a_{d}\right)$-nN rule is the nearest neighbor rule for the metric

$$
\|x\|=\sqrt{\sum_{i=1}^{d} a_{i}^{2} x^{(i)^{2}}},
$$

where $x=\left(x^{(1)}, \ldots, x^{(d)}\right)$. The class $\mathcal{C}_{m}$ is the class of all $\left(a_{1}, \ldots, a_{d}\right)$-NN rules for $D_{m}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{m}, Y_{m}\right)\right)$. The testing sequence $T_{l}$ is used to choose $\left(a_{1}, \ldots, a_{d}\right)$ so as to minimize the holdout error estimate. In order to have

$$
L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right) \rightarrow 0 \quad \text { in probability }
$$

it suffices that

$$
\lim _{n \rightarrow \infty} \frac{\log \mathbf{E}\left\{\mathcal{S}\left(\mathcal{C}_{m}, l\right)\right\}}{l}=0 .
$$

This puts a lower bound on $l$. To get this lower bound, one must compute $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$. Clearly, $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ is bounded by the number of ways of classifying $X_{m+1}, \ldots, X_{m+l}$ using rules picked from $\mathcal{C}_{m}$, that is, the total number of different values for

$$
\left\{\left(\phi_{m}\left(X_{m+1}\right), \ldots, \phi_{m}\left(X_{m+l}\right)\right), \phi_{m} \in \mathcal{C}_{m}\right\}
$$

We now show that regardless of $D_{m}$ and $X_{m+1}, \ldots, X_{m+l}$, for $n \geq 4$ we have

$$
\mathcal{S}\left(\mathcal{C}_{m}, l\right) \leq\left(l\binom{m}{2}\right)^{d}
$$

This sort of result is typical-the bound does not depend upon $D_{m}$ or $T_{l}$. When plugged back into the condition of convergence, it yields the simple condition

$$
\frac{\log m}{l} \rightarrow 0
$$

In fact, we may thus take $l$ slightly larger than $\log m$. It would plainly be silly to take $l=n / 2$, as we would thereby in fact throw away most of the data.

Set

$$
\mathcal{A}_{m}=\left\{\left(\phi_{m}\left(X_{m+1}\right), \ldots, \phi_{m}\left(X_{m+l}\right)\right), \phi_{m} \in \mathcal{C}_{m}\right\}
$$

To count the number of values in $\mathcal{A}_{m}$, note that all squared distances can be written as

$$
\left\|X_{i}-X_{j}\right\|^{2}=\sum_{k=1}^{d} a_{k}^{2} \rho_{k, i, j}
$$

where $\rho_{k, i, j}$ is a nonnegative number only depending upon $X_{i}$ and $X_{j}$. Note that each squared distance is linear in $\left(a_{1}^{2}, \ldots, a_{d}^{2}\right)$. Now consider the space of $\left(a_{1}^{2}, \ldots, a_{d}^{2}\right)$. Observe that, in this space, $\phi_{m}\left(X_{n+1}\right)$ is constant within each cell of the partition determined by the $\binom{m}{2}$ hyperplanes

$$
\sum_{k=1}^{d} a_{k}^{2} \rho_{k, i, m+1}=\sum_{k=1}^{d} a_{k}^{2} \rho_{k, i^{\prime}, m+1}, \quad 1 \leq i<i^{\prime} \leq m
$$

To see this, note that within each set in the partition, $\phi_{m}\left(X_{m+1}\right)$ keeps the same nearest neighbor among $X_{1}, \ldots, X_{m}$. It is known that $k>2$ hyperplanes in $\mathcal{R}^{d}$ create partitions of cardinality not exceeding $k^{d}$ (see Problem 22.1). Now, overlay the $l$ partitions obtained for $\phi_{m}\left(X_{m+1}\right), \ldots, \phi_{m}\left(X_{m+l}\right)$ respectively. This yields at most

$$
\left(l\binom{m}{2}\right)^{d}
$$

sets, as the overlays are determined by $l\binom{m}{2}$ hyperplanes. But clearly, on each of these sets, $\left(\phi_{m}\left(X_{m+1}\right), \ldots, \phi_{m}\left(X_{m+l}\right)\right)$ is constant. Therefore,

$$
\mathcal{S}\left(\mathcal{C}_{m}, l\right) \leq\left|\mathcal{A}_{m}\right| \leq\left(l\binom{m}{2}\right)^{d}
$$

### 22.4 Classification Based on Clustering

Recall the classification rule based on clustering that was introduced in Chapter 21. The data points $X_{1}, \ldots, X_{n}$ are grouped into $k$ clusters, where $k$ is a predetermined integer, and a majority vote decides within the $k$ clusters. If $k$ is chosen such that $k \rightarrow \infty$ and $k^{2} \log n / n \rightarrow 0$, then the rule is consistent. For a given finite $n$, however, these conditions give little guidance. Also, the choice of $k$ could dramatically affect the performance of the rule, as there may be a mismatch between $k$ and some unknown natural number of clusters. For example, one may construct distributions in which the optimal number of clusters does not increase with $n$.

Let us split the data and let the testing sequence decide the value of $k$. In the framework of this chapter, $\mathcal{C}_{m}$ contains the classifiers based on the first $m$ pairs $D_{m}$ of the data with all possible values of $k$. Clearly, $\mathcal{C}_{m}$ is a finite family with $\left|\mathcal{C}_{m}\right|=m$. In this case, by Problem 12.1, we have

$$
\mathbf{E}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right) \mid D_{m}\right\} \leq \sqrt{\frac{2 \log (2 m)}{l}}
$$

The consistency result Theorem 21.5 implies that

$$
\mathbf{E}\left\{\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)\right\} \rightarrow L^{*}
$$

for all distributions, if the $X_{i}$ 's have bounded support. Thus, we see that our strategy leads to a universally consistent rule whenever $(\log m) / l \rightarrow 0$. This is a very mild condition, since we can take $l$ equal to a small fraction of $n$, without sacrificing consistency. If $l$ is small compared to $n$, then $m$ is close to $n$, so $\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)$ is likely to be close to $\inf _{\phi_{n} \in \mathcal{C}_{n}} L\left(\phi_{n}\right)$. Thus, we do not lose much by sacrificing a part of the data for testing purposes, but the gain can be tremendous, as we are guaranteed to be within $\sqrt{(\log m) / l}$ of the optimum in the class $\mathcal{C}_{m}$.

That we cannot take $l=1$ and hope to obtain consistency should be obvious. It should also be noted that for $l=m$, we are roughly within $\sqrt{\log (m) / m}$ of the best possible probability of error within the family. Also the empirical selection rule is $\sqrt{\log (m) / l}$-optimal.

### 22.5 Statistically Equivalent Blocks

Recall from Chapter 21 that classifiers based on statistically equivalent blocks typically partition the feature space $\mathcal{R}^{d}$ into rectangles such that each rectangle contains $k$ points, where $k$ is a certain positive integer, the parameter of the rule. This may be done in several different ways. One of many such rules-the rule introduced by Gessaman (1970)- is consistent if $k \rightarrow \infty$ and $k / n \rightarrow 0$ (Theorem 21.4). Again, we can let the data pick the value of $k$ by minimizing the holdout estimate. Just as in the previous section, $\left|\mathcal{C}_{m}\right|=m$, and every remark mentioned there about consistency and asymptotic optimality remains true for this case as well.

We can enlarge the family $\mathcal{C}_{m}$ by allowing partitions without restrictions on cardinalities of cells. This leads very quickly to oversized families of rules, and we have to impose reasonable restrictions. Consider cuts into at most $k$ rectangles, where $k$ is a number picked beforehand. Recall that for a fixed partition, the class assigned to every rectangle is decided upon by a majority vote among the training points. On the real line, choosing a partition into at most $k$ sets is equivalent to choosing $k-1$ cut positions from $m+l+1=n+1$ spacings between all test and training points. Hence,

$$
\mathcal{S}\left(\mathcal{C}_{m}, l\right) \leq\binom{ n+1}{k-1} \leq(n+1)^{k-1}
$$

For consistency, $k$ has to grow as $n$ grows. It is easily seen that

$$
\mathbf{E}\left\{\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)\right\} \rightarrow L^{*}
$$

if $k \rightarrow \infty$ and $m \rightarrow \infty$. To achieve consistency of the selected rule, however, we also need

$$
\frac{\log \mathcal{S}\left(\mathcal{C}_{m}, l\right)}{l} \leq \frac{(k-1) \log (n+1)}{l} \rightarrow 0
$$

Now consider $d$-dimensional partitions defined by at most $k-1$ consecutive orthogonal cuts, that is, the first cut divides $\mathcal{R}^{d}$ into two halfspaces along a hyperplane perpendicular to one of the coordinate axes. The second cut splits one of the halfspaces into two parts along another orthogonal hyperplane, and so forth. This procedure yields a partition of the space into $k$ rectangles. We see that for the first cut, there are at most $1+d n$ possible combinations to choose from. This yields the loose upper bound

$$
\mathcal{S}\left(\mathcal{C}_{m}, l\right) \leq(1+d n)^{k-1}
$$

This bound is also valid for all grids defined by at most $k-1$ cuts. The main difference here is that every cut defines two halfspaces of $\mathcal{R}^{d}$, and not two hyperrectangles of a cells, so that we usually end up with $2^{k}$ rectangles in the partition.

Assume that $\mathcal{C}_{m}$ contains all histograms with partitions into at most $k$ (possibly infinite) rectangles. Then, considering that a rectangle in $\mathcal{R}^{d}$ requires choosing $2 d$ spacings between all test and training points, two per coordinate axis,

$$
\mathcal{S}\left(\mathcal{C}_{m}, l\right) \leq(n+1)^{2 d(k-1)}
$$

See Feinholz (1979) for more work on such partitions.

### 22.6 Binary Tree Classifiers

We can analyze binary tree classifiers from the same viewpoint. Recall that such classifiers are represented by binary trees, where each internal node corresponds to a split of a cell by a hyperplane, and the terminal nodes represent the cells of the partition.

Assume that there are $k$ cells (and therefore $k-1$ splits leading to the partition). If every split is perpendicular to one of the axes, then, the situation is the same as in the previous section,

$$
\mathcal{S}\left(\mathcal{C}_{m}, l\right) \leq(1+d n)^{k-1}
$$

For smaller families of rules whose cuts depend upon the training sequence only, the bound is pessimistic. Others have proposed generalizing orthogonal cuts by using general hyperplane cuts. Recall that there are at most

$$
\sum_{j=0}^{d+1}\binom{n}{j} \leq n^{d+1}+1
$$

ways of dichotomizing $n$ points in $\mathcal{R}^{d}$ by hyperplanes (see Theorem 13.9). Thus, if we allow up to $k-1$ internal nodes (or hyperplane cuts),

$$
\mathcal{S}\left(\mathcal{C}_{m}, l\right) \leq\left(1+n^{d+1}\right)^{k-1} .
$$

The number of internal nodes has to be restricted in order to obtain consistency from this bound: Refer to Chapter 20 for more details.

## Problems and Exercises

Problem 22.1. Prove that $n$ hyperplanes partition $\mathcal{R}^{d}$ into at most $\sum_{i=0}^{d}\binom{n}{i}$ contiguous regions when $d \leq n$ (Schläffli (1950), Cover (1965)). Hint: Proceed by induction.

Problem 22.2. Assume that $g_{n}$ is selected so as to minimize the holdout error estimate $\widehat{L}_{l, m}\left(\phi_{m}\right)$ over $\mathcal{C}_{m}$, the class of rules based upon the first $m$ data points. Assume furthermore that we vary $l$ over $\left[\log ^{2} n, n / 2\right]$, and that we pick the best $l$ (and $m=n-l$ ) by minimizing the holdout error estimate again. Show that if $\mathcal{S}\left(\mathcal{C}_{m}, l\right)=O\left(n^{\gamma}\right)$ for some $\gamma>0$, then the obtained rule is strongly universally consistent.

Problem 22.3. Let $\mathcal{C}_{m}$ be the class of $\left(a_{1}, \ldots, a_{d}\right)$-nN rules based upon $D_{m}$. Show that if $m / n \rightarrow 1$, then

$$
\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)-\inf _{\phi_{n} \in \mathcal{C}_{n}} L\left(\phi_{n}\right) \rightarrow 0 \quad \text { in probability }
$$

for all distributions of $(X, Y)$ for which $X$ has a density. Conclude that if $l / \log n \rightarrow \infty$, $m=n-l, l=o(n)$, then

$$
L\left(g_{n}\right)-\inf _{\phi_{n} \in \mathcal{C}_{n}} L\left(\phi_{n}\right) \rightarrow 0 \quad \text { in probability. }
$$

Problem 22.4. Finding the best split. This exercise is concerned with the automatic selection of $m$ and $l=n-m$. If $g_{n}$ is the selected rule minimizing the holdout estimate, then

$$
\begin{equation*}
\mathrm{E}\left\{L\left(g_{n}\right) \mid D_{m}\right\}-L^{*} \leq \sqrt{\frac{2 \log \left(4 \mathcal{S}\left(\mathcal{C}_{m}, l^{2}\right)\right)+16}{l}}+\left(\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)-L^{*}\right) \tag{22.3}
\end{equation*}
$$

Since in most interesting cases $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ is bounded from above as a polynomial of $m$ and $l$, the estimation error typically decreases as $l$ increases. On the other hand, the approximation error $\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)-L^{*}$ typically decreases as $m$ increases, as the class $\mathcal{C}_{m}$ gets richer. Some kind of balance between the two terms is required to get optimum performance. We may use the empirical estimates $\widehat{L}_{m, l}\left(\phi_{m}\right)$ again to decide which value of $m$ we wish to choose (Problem 22.2). However, as $m$ gets large-and therefore $l$ small—the class $\mathcal{C}_{m}$ will tend to overfit the data $T_{l}$, providing strongly optimistically biased estimates for $\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)$. To prevent overfitting, we may apply the method of complexity regularization (Chapter 18). By (22.1), we may define the penalty term by

$$
r(m, l)=\sqrt{\frac{2 \log \left(4 e^{8} \mathbf{E}\left\{\mathcal{S}\left(\mathcal{C}_{m}, l^{2}\right)\right\}\right)+\log n}{l}}
$$

and minimize the penalized error estimate $\tilde{L}_{m, l}(\phi)=\widehat{L}_{m, l}(\phi)+r(m, l)$ over all $\phi \in \cup_{m=1}^{n} \mathcal{C}_{m}$. Denote the selected rule by $\phi_{n}^{*}$. Prove that for every $n$ and all distributions of $(X, Y)$,

$$
\begin{aligned}
\mathbf{E}\left\{L\left(\phi_{n}^{*}\right)-L^{*}\right\} \leq & \min _{m, l}\left(3 \sqrt{\frac{2 \log \left(4 \mathbf{E}\left\{\mathcal{S}\left(\mathcal{C}_{m}, l^{2}\right)\right\}\right)+\log n+16}{l}}\right. \\
& \left.+\left(\mathbf{E}\left\{\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)-L^{*}\right\}\right)\right)+\frac{2}{\sqrt{n}}
\end{aligned}
$$

Hint: Proceed similarly to the proof of Theorem 18.2.

## 23

## The Resubstitution Estimate

Estimating the error probability is of primordial importance for classifier selection. The method explored in the previous chapter attempts to solve this problem by using a testing sequence to obtain a reliable holdout estimate. The independence of testing and training sequences leads to a rather straightforward analysis. For a good performance, the testing sequence has to be sufficiently large (although we often get away with testing sequences as small as about $\log n$ ). When data are expensive, this constitutes a waste. Assume that we do not split the data and use the same sequence for testing and training. Often dangerous, this strategy nevertheless works if the class of rules from which we select is sufficiently restricted. The error estimate in this case is appropriately called the resubstitution estimate and it will be denoted by $L_{n}^{(R)}$. This chapter explores its virtues and pitfalls. A third error estimate, the deleted estimate, is discussed in the next chapter. Estimates based upon other paradigms are treated briefly in Chapter 31.

### 23.1 The Resubstitution Estimate

The resubstitution estimate $L_{n}^{(R)}$ counts the number of errors committed on the training sequence by the classification rule. Expressed formally,

$$
L_{n}^{(R)}=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{g_{n}\left(X_{i}\right) \neq Y_{i}\right\}} .
$$

Sometimes $L_{n}^{(R)}$ is called the apparent error rate. It is usually strongly optimisti-
cally biased. Since the classifier $g_{n}$ is tuned by $D_{n}$, it is intuitively clear that $g_{n}$ may behave better on $D_{n}$ than on independent data.

The best way to demonstrate this biasedness is to consider the 1-nearest neighbor rule. If $X$ has a density, then the nearest neighbor of $X_{i}$ among $X_{1}, \ldots, X_{n}$ is $X_{i}$ itself with probability one. Therefore, $L_{n}^{(R)} \equiv 0$, regardless of the value of $L_{n}=L\left(g_{n}\right)$. In this case the resubstitution estimate is useless. For $k$-nearest neighbor rules with large $k, L_{n}^{(R)}$ is close to $L_{n}$. This was demonstrated by Devroye and Wagner (1979a), who obtained upper bounds on the performance of the resubstitution estimate for the $k$-nearest neighbor rule without posing any assumption on the distribution.

Also, if the classifier whose error probability is to be estimated is a member of a class of classifiers with finite Vapnik-Chervonenkis dimension (see the definitions in Chapter 12), then we can get good performance bounds for the resubstitution estimate. To see this, consider any generalized linear classification rule, that is, any rule that can be put into the following form:

$$
g_{n}(x)= \begin{cases}0 & \text { if } a_{0, n}+\sum_{i=1}^{d^{*}} a_{i, n} \psi_{i}(x) \leq 0 \\ 1 & \text { otherwise }\end{cases}
$$

where the $\psi_{i}$ 's are fixed functions, and the coefficients $a_{i, n}$ depend on the data $D_{n}$ in an arbitrary but measurable way. We have the following estimate for the performance of the resubstitution estimate $L_{n}^{(R)}$.

Theorem 23.1. (Devroye and Wagner (1976A)). For all $n$ and $\epsilon>0$, the resubstitution estimate $L_{n}^{(R)}$ of the error probability $L_{n}$ of a generalized linear rule satisfies

$$
\mathbf{P}\left\{\left|L_{n}^{(R)}-L_{n}\right| \geq \epsilon\right\} \leq 8 n^{d^{*}} e^{-n \epsilon^{2} / 32}
$$

Proof. Define the set $A_{n} \subset \mathcal{R}^{d} \times\{0,1\}$ as the set of all pairs $(x, y) \in \mathcal{R}^{d} \times\{0,1\}$, on which $g_{n}$ errs:

$$
A_{n}=\left\{(x, y): g_{n}(x) \neq y\right\} .
$$

Observe that

$$
L_{n}=\mathbf{P}\left\{(X, Y) \in A_{n} \mid D_{n}\right\}
$$

and

$$
L_{n}^{(R)}=\frac{1}{n} \sum_{j=1}^{n} I_{\left\{\left(X_{j}, Y_{j}\right) \in A_{n}\right\}},
$$

or denoting the measure of $(X, Y)$ by $v$, and the corresponding empirical measure by $v_{n}$,

$$
L_{n}=v\left(A_{n}\right),
$$

and

$$
L_{n}^{(R)}=v_{n}\left(A_{n}\right)
$$

The set $A_{n}$ depends on the data $D_{n}$, so that, for example, $\mathbf{E}\left\{v_{n}\left(A_{n}\right)\right\} \neq \mathbf{E}\left\{v\left(A_{n}\right)\right\}$. Fortunately, the powerful Vapnik-Chervonenkis theory comes to the rescue via the inequality

$$
\left|L_{n}-L_{n}^{(R)}\right| \leq \sup _{C \in \mathcal{C}}\left|v_{n}(C)-v(C)\right|
$$

where $\mathcal{C}$ is the family of all sets of the form $\{(x, y): \phi(x) \neq y\}$, where $\phi: \mathcal{R}^{d} \rightarrow$ $\{0,1\}$ is a generalized linear classifier based on the functions $\psi_{1}, \ldots, \psi_{d^{*}}$. By Theorems $12.6,13.1$, and 13.7 we have

$$
\mathbf{P}\left\{\sup _{C \in \mathcal{C}}\left|v_{n}(C)-v(C)\right| \geq \epsilon\right\} \leq 8 n^{d^{*}} e^{-n \epsilon^{2} / 32}
$$

Similar inequalities can be obtained for other classifiers. For example, for partitioning rules with fixed partitions we have the following:

Theorem 23.2. Let $g_{n}$ be a classifier whose value is constant over cells of a fixed partition of $\mathcal{R}^{d}$ into $k$ cells. Then

$$
\mathbf{P}\left\{\left|L_{n}^{(R)}-L_{n}\right| \geq \epsilon\right\} \leq 8 \cdot 2^{k} e^{-n \epsilon^{2} / 32}
$$

The proof is left as an exercise (Problem 23.1). From Theorems 23.1 and 23.2 we get bounds for the expected difference between the resubstitution estimate and the actual error probability $L_{n}$. For example, Theorem 23.1 implies (via Problem 12.1) that

$$
\mathbf{E}\left\{\left|L_{n}^{(R)}-L_{n}\right|\right\}=O\left(\sqrt{\frac{d^{*} \log n}{n}}\right)
$$

In some special cases the expected behavior of the resubstitution estimate can be analyzed in more detail. For example, McLachlan (1976) proved that if the conditional distributions of $X$ given $Y=0$ and $Y=1$ are both normal with the same covariance matrices, and the rule is linear and based on the estimated parameters, then the bias of the estimate is of the order $1 / n$ :

$$
\left|\mathbf{E}\left\{L_{n}^{(R)}\right\}-\mathbf{E}\left\{L_{n}\right\}\right|=O\left(\frac{1}{n}\right)
$$

McLachlan also showed for this case that for large $n$ the expected value of the resubstitution estimate is smaller than that of $L_{n}$, that is, the estimate is optimistically biased, as expected.

### 23.2 Histogram Rules

In this section, we explore the properties of the resubstitution estimate for histogram rules. Let $\mathcal{P}=\left\{A_{1}, A_{2}, \ldots\right\}$ be a fixed partition of $\mathcal{R}^{d}$, and let $g_{n}$ be the
corresponding histogram classifier (see Chapters 6 and 9). Introduce the notation

$$
v_{0, n}(A)=\frac{1}{n} \sum_{j=1}^{n} I_{\left\{Y_{j}=0\right\}} I_{\left\{X_{j} \in A\right\}}, \quad \text { and } \quad v_{1, n}(A)=\frac{1}{n} \sum_{j=1}^{n} I_{\left\{Y_{j}=1\right\}} I_{\left\{X_{j} \in A\right\}}
$$

The analysis is simplified if we rewrite the estimate $L_{n}^{(R)}$ in the following form (see Problem 23.2):

$$
\begin{equation*}
L_{n}^{(R)}=\sum_{i} \min \left\{v_{0, n}\left(A_{i}\right), v_{1, n}\left(A_{i}\right)\right\} \tag{23.1}
\end{equation*}
$$

It is also interesting to observe that $L_{n}^{(R)}$ can be put in the following form:

$$
L_{n}^{(R)}=\int\left(I_{\left\{\eta_{1, n}(x) \leq \eta_{0, n}(x)\right\}} \eta_{1, n}(x)+I_{\left\{\eta_{1, n}(x)>\eta_{0, n}(x)\right\}} \eta_{0, n}(x)\right) \mu(d x)
$$

where

$$
\eta_{0, n}(x)=\frac{\frac{1}{n} \sum_{j=1}^{n} I_{\left\{Y_{j}=0, X_{j} \in A(x)\right\}}}{\mu(A(x))} \quad \text { and } \quad \eta_{1, n}(x)=\frac{\frac{1}{n} \sum_{j=1}^{n} I_{\left\{Y_{j}=1, X_{j} \in A(x)\right\}}}{\mu(A(x))}
$$

We can compare this form with the following expression of $L_{n}$ :

$$
L_{n}=\int\left(I_{\left\{\eta_{1, n}(x) \leq \eta_{0, n}(x)\right\}} \eta(x)+I_{\left\{\eta_{1, n}(x)>\eta_{0, n}(x)\right\}}(1-\eta(x))\right) \mu(d x)
$$

(see Problem 23.3). We begin the analysis of performance of the estimate by showing that its mean squared error is not larger than a constant times the number of cells in the partition over $n$.

Theorem 23.3. For any distribution of $(X, Y)$ and for all $n$, the estimate $L_{n}^{(R)}$ of the error probability of any histogram rule satisfies

$$
\operatorname{Var}\left\{L_{n}^{(R)}\right\} \leq \frac{1}{n}
$$

Also, the estimate is optimistically biased, that is,

$$
\mathbf{E}\left\{L_{n}^{(R)}\right\} \leq \mathbb{E}\left\{L_{n}\right\}
$$

If, in addition, the histogram rule is based on a partition $\mathcal{P}$ of $\mathcal{R}^{d}$ into at most $k$ cells, then

$$
\mathbb{E}\left\{\left(L_{n}^{(R)}-L_{n}\right)^{2}\right\} \leq \frac{6 k}{n}
$$

Proof. The first inequality is an immediate consequence of Theorem 9.3 and (23.1). Introduce the auxiliary quantity

$$
R^{*}=\sum_{i} \min \left\{v_{0}\left(A_{i}\right), v_{1}\left(A_{i}\right)\right\}
$$

where $\nu_{0}(A)=\mathbf{P}\{Y=0, X \in A\}$, and $\nu_{1}(A)=\mathbf{P}\{Y=1, X \in A\}$. We use the decomposition

$$
\begin{equation*}
\mathbf{E}\left\{\left(L_{n}^{(R)}-L_{n}\right)^{2}\right\} \leq 2\left(\mathbb{E}\left\{\left(L_{n}^{(R)}-R^{*}\right)^{2}\right\}+\mathbb{E}\left\{\left(R^{*}-L_{n}\right)^{2}\right\}\right) \tag{23.2}
\end{equation*}
$$

First we bound the second term on the right-hand side of (23.2). Observe that $R^{*}$ is just the Bayes error corresponding to the pair of random variables $(T(X), Y)$, where the function $T$ transforms $X$ according to $T(x)=i$ if $x \in A_{i}$. Since the histogram classification rule $g_{n}(x)$ can be written as a function of $T(x)$, its error probability $L_{n}$ cannot be smaller than $R^{*}$. Furthermore, by the first identity of Theorem 2.2, we have

$$
\begin{aligned}
0 & \leq L_{n}-R^{*} \\
& =\sum_{i} I_{\left\{\operatorname{sign}\left(\nu_{1, n}\left(A_{i}\right)-v_{0, n}\left(A_{i}\right)\right) \neq \operatorname{sign}\left(v_{1}\left(A_{i}\right)-v_{0}\left(A_{i}\right)\right)\right\}}\left|v_{1}\left(A_{i}\right)-v_{0}\left(A_{i}\right)\right| \\
& \leq \sum_{i}\left|v_{1}\left(A_{i}\right)-v_{0}\left(A_{i}\right)-\left(v_{1, n}\left(A_{i}\right)-v_{0, n}\left(A_{i}\right)\right)\right| .
\end{aligned}
$$

If the partition has at most $k$ cells, then by the Cauchy-Schwarz inequality,

$$
\begin{aligned}
& \mathbb{E}\left\{\left(L_{n}-R^{*}\right)^{2}\right\} \\
& \quad \leq \mathbb{E}\left\{\left(\sum_{i}\left|v_{1}\left(A_{i}\right)-v_{0}\left(A_{i}\right)-\left(v_{1, n}\left(A_{i}\right)-v_{0, n}\left(A_{i}\right)\right)\right|\right)^{2}\right\} \\
& \leq k \sum_{i} \operatorname{Var}\left\{v_{1, n}\left(A_{i}\right)-v_{0, n}\left(A_{i}\right)\right\} \\
& \quad \leq k \sum_{i} \frac{\mu\left(A_{i}\right)}{n}=\frac{k}{n} .
\end{aligned}
$$

We bound the first term on the right-hand side of (23.2):

$$
\mathbb{E}\left\{\left(L_{n}^{(R)}-R^{*}\right)^{2}\right\} \leq \operatorname{Var}\left\{L_{n}^{(R)}\right\}+\left(R^{*}-\mathbb{E}\left\{L_{n}^{(R)}\right\}\right)^{2}
$$

As we have seen earlier, $\operatorname{Var}\left\{L_{n}^{(R)}\right\} \leq 1 / n$, so it suffices to bound $\left|R^{*}-\mathbb{E}\left\{L_{n}^{(R)}\right\}\right|$. By Jensen's inequality,

$$
\begin{aligned}
\mathbb{E}\left\{L_{n}^{(R)}\right\} & =\sum_{i} \mathbf{E}\left\{\min \left\{v_{0, n}\left(A_{i}\right), \nu_{1, n}\left(A_{i}\right)\right\}\right\} \\
& \leq \sum_{i} \min \left(\mathbf{E}\left\{v_{0, n}\left(A_{i}\right)\right\}, \mathbf{E}\left\{v_{1, n}\left(A_{i}\right)\right\}\right) \\
& =\sum_{i} \min \left(v_{0}\left(A_{i}\right), v_{1}\left(A_{i}\right)\right) \\
& =R^{*} .
\end{aligned}
$$

So we bound $R^{*}-\mathbf{E}\left\{L_{n}^{(R)}\right\}$ from above. By the inequality

$$
|\min (a, b)-\min (c, d)| \leq|a-c|+|b-d|
$$

we have

$$
\begin{aligned}
R^{*}-\mathbf{E}\left\{L_{n}^{(R)}\right\} \leq & \sum_{i} \mathbf{E}\left\{\left|v_{0}\left(A_{i}\right)-v_{0, n}\left(A_{i}\right)\right|+\left|v_{1}\left(A_{i}\right)-v_{1, n}\left(A_{i}\right)\right|\right\} \\
\leq & \sum_{i}\left(\sqrt{\operatorname{Var}\left\{v_{0, n}\left(A_{i}\right)\right\}}+\sqrt{\operatorname{Var}\left\{v_{1, n}\left(A_{i}\right)\right\}}\right) \\
& (\text { by the Cauchy-Schwarz inequality }) \\
= & \sum_{i}\left(\sqrt{\frac{v_{0}\left(A_{i}\right)\left(1-v_{0}\left(A_{i}\right)\right)}{n}}+\sqrt{\frac{v_{1}\left(A_{i}\right)\left(1-v_{1}\left(A_{i}\right)\right)}{n}}\right) \\
\leq & \sum_{i}\left(\sqrt{\frac{2\left(v_{0}\left(A_{i}\right)\left(1-v_{0}\left(A_{i}\right)\right)+v_{1}\left(A_{i}\right)\left(1-v_{1}\left(A_{i}\right)\right)\right)}{n}}\right) \\
\leq & \sum_{i} \sqrt{\frac{2 \mu\left(A_{i}\right)}{n}},
\end{aligned}
$$

where we used the elementary inequality $\sqrt{a}+\sqrt{b} \leq \sqrt{2(a+b)}$. Therefore,

$$
\left(R^{*}-\mathbf{E}\left\{L_{n}^{(R)}\right\}\right)^{2} \leq\left(\sum_{i} \sqrt{\frac{2 \mu\left(A_{i}\right)}{n}}\right)^{2}
$$

To complete the proof of the third inequality, observe that if there are at most $k$ cells, then

$$
\begin{aligned}
R^{*}-\mathbf{E}\left\{L_{n}^{(R)}\right\} & \leq \sum_{i} \sqrt{\frac{2 \mu\left(A_{i}\right)}{n}} \\
& \leq k \sqrt{\frac{1}{k} \sum_{i} \frac{2 \mu\left(A_{i}\right)}{n}} \text { (by Jensen's inequality) } \\
& =\sqrt{\frac{2 k}{n}}
\end{aligned}
$$

Therefore, $\left(R^{*}-\mathbf{E}\left\{L_{n}^{(R)}\right\}\right)^{2} \leq 2 k / n$. Finally,

$$
\mathbf{E}\left\{L_{n}\right\}-\mathbf{E}\left\{L_{n}^{(R)}\right\}=\left(\mathbf{E}\left\{L_{n}\right\}-R^{*}\right)+\left(R^{*}-\mathbf{E}\left\{L_{n}^{(R)}\right\}\right) \geq 0
$$

We see that if the partition contains a small number of cells, then the resubstitution estimate performs very nicely. However, if the partition has a large number of cells, then the resubstitution estimate of $L_{n}$ can be very misleading, as the next result indicates:

Theorem 23.4. (Devroye and Wagner (1976B)). For every $n$ there exists a partitioning rule and a distribution such that

$$
\mathbf{E}\left\{\left(L_{n}^{(R)}-L_{n}\right)^{2}\right\} \geq 1 / 4
$$

Proof. Let $A_{1}, \ldots, A_{2 n}$ be $2 n$ cells of the partition such that $\mu\left(A_{i}\right)=1 /(2 n)$, $i=1, \ldots, 2 n$. Assume further that $\eta(x)=1$ for every $x \in \mathcal{R}^{d}$, that is, $Y=1$ with probability one. Then clearly $L_{n}^{(R)}=0$. On the other hand, if a cell does not contain any of the data points $X_{1}, \ldots, X_{n}$, then $g_{n}(x)=0$ in that cell. But since the number of points is only half of the number of cells, at least half of the cells are empty. Therefore $L_{n} \geq 1 / 2$, and $\left|L_{n}^{(R)}-L_{n}\right| \geq 1 / 2$. This concludes the proof.

### 23.3 Data-Based Histograms and Rule Selection

Theorem 23.3 demonstrates the usefulness of $L_{n}^{(R)}$ for histogram rules with a fixed partition, provided that the number of cells in the partition is not too large. If we want to use $L_{n}^{(R)}$ to select a good classifier, the estimate should work uniformly well over the class from which we select a rule. In this section we explore such data-based histogram rules.

Let $\mathcal{F}$ be a class of partitions of $\mathcal{R}^{d}$. We will assume that each member of $\mathcal{F}$ partitions $\mathcal{R}^{d}$ into at most $k$ cells. For each partition $\mathcal{P} \in \mathcal{F}$, define the corresponding histogram rule by

$$
g_{n}^{(\mathcal{P})}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=1, X_{i} \in A(x)\right\}} \leq \sum_{i=1}^{n} I_{\left\{Y_{i}=0, X_{i} \in A(x)\right\}} \\ 1 & \text { otherwise },\end{cases}
$$

where $A(x)$ is the cell of $\mathcal{P}$ that contains $x$. Denote the error probability of $g_{n}^{(\mathcal{P})}(x)$ by

$$
L_{n}(\mathcal{P})=\mathbf{P}\left\{g_{n}^{(\mathcal{P})}(X) \neq Y \mid D_{n}\right\}
$$

The corresponding error estimate is denoted by

$$
L_{n}^{(R)}(\mathcal{P})=\sum_{A \in \mathcal{P}} \min \left\{v_{0, n}(A), v_{1, n}(A)\right\}
$$

By analogy with Theorems 23.1 and 23.2, we can derive the following result, which gives a useful bound for the largest difference between the estimate and the error probability within the class of histogram classifiers defined by $\mathcal{F}$. The combinatorial coefficient $\Delta_{n}^{*}(\mathcal{F})$ defined in Chapter 21 appears as a coefficient in the upper bound. The computation of $\Delta_{n}^{*}(\mathcal{F})$ for several different classes of partitions is illustrated in Chapter 21.

Theorem 23.5. (Feinholz (1979)). Assume that each member of $\mathcal{F}$ partitions $\mathcal{R}^{d}$ into at most $k$ cells. For every $n$ and $\epsilon>0$,

$$
\mathbb{P}\left\{\sup _{\mathcal{P} \in \mathcal{F}}\left|L_{n}^{(R)}(\mathcal{P})-L_{n}(\mathcal{P})\right|>\epsilon\right\} \leq 8 \cdot 2^{k} \Delta_{n}^{*}(\mathcal{F}) e^{-n \epsilon^{2} / 32}
$$

Proof. We can proceed as in the proof of Theorem 23.1. The shatter coefficient $\mathcal{S}(\mathcal{C}, n)$ corresponding to the class of histogram classifiers defined by partitions in $\mathcal{F}$ can clearly be bounded from above by the number of different ways in which $n$ points can be partitioned by members of $\mathcal{F}$, times $2^{k}$, as there are at most $2^{k}$ different ways to assign labels to cells of a partition of at most $k$ cells.

The theorem has two interesting implications. The error estimate $L_{n}^{(R)}$ can also be used to estimate the performance of histogram rules based on data-dependent partitions (see Chapter 21). The argument of the proof of Theorem 23.3 is not valid for these rules. However, Theorem 23.5 provides performance guarantees for these rules in the following corollaries:

COROLLARY 23.1. Let $g_{n}(x)$ be a histogram classifier based on a random partition $\mathcal{P}_{n}$ into at most $k$ cells, which is determined by the data $D_{n}$. Assume that for any possible realization ot the training data $D_{n}$, the partition $\mathcal{P}_{n}$ is a member of a class of partitions $\mathcal{F}$. If $L_{n}$ is the error probability of $g_{n}$, then

$$
\mathbf{P}\left\{\left|L_{n}^{(R)}-L_{n}\right|>\epsilon\right\} \leq 8 \cdot 2^{k} \Delta_{n}^{*}(\mathcal{F}) e^{-n \epsilon^{2} / 32}
$$

and

$$
\mathbb{E}\left\{\left(L_{n}^{(R)}-L_{n}\right)^{2}\right\} \leq \frac{32 k+32 \log \left(\Delta_{n}^{*}(\mathcal{F})\right)+128}{n}
$$

Proof. The first inequality follows from Theorem 23.5 by the obvious inequality

$$
\mathbb{P}\left\{\left|L_{n}^{(R)}-L_{n}\right|>\epsilon\right\} \leq \mathbb{P}\left\{\sup _{\mathcal{P} \in \mathcal{F}}\left|L_{n}^{(R)}(\mathcal{P})-L_{n}(\mathcal{P})\right|>\epsilon\right\} .
$$

The second inequality follows from the first one via Problem 12.1. $\square$

Perhaps the most important application of Theorem 23.5 is in classifier selection. Let $\mathcal{C}_{n}$ be a class of (possibly data-dependent) histogram rules. We may use the error estimate $L_{n}^{(R)}$ to select a classifier that minimizes the estimated error probability. Denote the selected histogram rule by $\phi_{n}^{*}$, that is,

$$
L_{n}^{(R)}\left(\phi_{n}^{*}\right) \leq L_{n}^{(R)}\left(\phi_{n}\right) \text { for all } \phi_{n} \in \mathcal{C}_{n}
$$

Here $L_{n}^{(R)}\left(\phi_{n}\right)$ denotes the estimated error probability of the classification rule $\phi_{n}$. The question is how well the selection method works, in other words, how close the error probability of the selected classifier $L\left(\phi_{n}^{*}\right)$ is to the error probability of the best rule in the class, $\inf _{\phi_{n} \in \mathcal{C}_{n}} L\left(\phi_{n}\right)$. It turns out that if the possible partitions are not too complex, then the method works very well:

COROLLARY 23.2. Assume that for any realization of the data $D_{n}$, the possible partitions that define the histogram classifiers in $\mathcal{C}_{n}$ belong to a class of partitions $\mathcal{F}$, whose members partition $\mathcal{R}^{d}$ into at most $k$ cells. Then

$$
\mathbf{P}\left\{L\left(\phi_{n}^{*}\right)-\inf _{\phi_{n} \in \mathcal{C}_{n}} L\left(\phi_{n}\right)>\epsilon\right\} \leq 8 \cdot 2^{k} \Delta_{n}^{*}(\mathcal{F}) e^{-n \epsilon^{2} / 128} .
$$

Proof. By Lemma 8.2,

$$
L\left(\phi_{n}^{*}\right)-\inf _{\phi_{n} \in \mathcal{C}_{n}} L\left(\phi_{n}\right) \leq 2 \sup _{\phi_{n} \in \mathcal{C}_{n}}\left|L_{n}^{(R)}\left(\phi_{n}\right)-L\left(\phi_{n}\right)\right| \leq 2 \sup _{\mathcal{P} \in \mathcal{F}}\left|L_{n}^{(R)}(\mathcal{P})-L(\mathcal{P})\right| .
$$

Therefore, the statement follows from Theorem 23.5.

## Problems and Exercises

problem 23.1. Prove Theorem 23.2. Hint: Proceed as in the proof of Theorem 23.1.
Problem 23.2. Show that for histogram rules, the resubstitution estimate may be written as

$$
L_{n}^{(R)}=\sum_{i} \min \left\{v_{0, n}\left(A_{i}\right), v_{1, n}\left(A_{i}\right)\right\} .
$$

Problem 23.3. Consider the resubstitution estimate $L_{n}^{(R)}$ of the error probability $L_{n}$ of a histogram rule based on a fixed sequence of partitions $\mathcal{P}_{n}$. Show that if the regression function estimate

$$
\eta_{n}(x)=\frac{\frac{1}{n} \sum_{j=1}^{n} I_{\left\{Y_{j}=1, X_{j} \in A(x)\right\}}}{\mu(A(x))}
$$

is consistent, that is, it satisfies $\mathbf{E}\left\{\int\left|\eta(x)-\eta_{n}(x)\right| \mu(d x)\right\} \rightarrow 0$ as $n \rightarrow \infty$, then

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{\left|L_{n}^{(R)}-L_{n}\right|\right\}=0
$$

and also $\mathbf{E}\left\{L_{n}^{(R)}\right\} \rightarrow L^{*}$.
Problem 23.4. Let $\left(X_{1}^{\prime}, Y_{1}^{\prime}\right), \ldots,\left(X_{m}^{\prime}, Y_{m}^{\prime}\right)$ be a sequence that depends in an arbitrary fashion on the data $D_{n}$, and let $g_{n}$ be the nearest neighbor rule with $\left(X_{1}^{\prime}, Y_{1}^{\prime}\right), \ldots,\left(X_{m}^{\prime}, Y_{m}^{\prime}\right)$, where $m$ is fixed. Let $L_{n}^{(R)}$ denote the resubstitution estimate of $L_{n}=L\left(g_{n}\right)$. Show that for all $\epsilon>0$ and all distributions,

$$
\mathbf{P}\left\{\left|L_{n}-L_{n}^{(R)}\right| \geq \epsilon\right\} \leq 8 n^{d m(m-1)} e^{-n \epsilon^{2} / 32}
$$

Problem 23.5. Find a rule for $(X, Y) \in \mathcal{R} \times\{0,1\}$ such that for all nonatomic distributions with $L^{*}=0$ we have $\mathbf{E}\left\{L_{n}\right\} \rightarrow 0$, yet $\mathbf{E}\left\{L_{n}^{(R)}\right\} \geq \mathbf{E}\left\{L_{n}\right\}$. (Thus, $L_{n}^{(R)}$ may be pessimistically biased even for a consistent rule.)

Problem 23.6. For histogram rules on fixed partitions (partitions that do not change with $n$ and are independent of the data), show that $\mathbf{E}\left\{L_{n}^{(R)}\right\}$ is monotonically nonincreasing.

Problem 23.7. Assume that $X$ has a density, and investigate the resubstitution estimate of the 3-NN rule. What is the limit of $\mathbf{E}\left\{L_{n}^{(R)}\right\}$ as $n \rightarrow \infty$ ?

## 24

## Deleted Estimates of the Error Probability

The deleted estimate (also called cross-validation, leave-one-out, or $U$-method) attempts to avoid the bias present in the resubstitution estimate. Proposed and developed by Lunts and Brailovsky (1967), Lachenbruch (1967), Cover (1969), and Stone (1974), the method deletes the first pair ( $X_{1}, Y_{1}$ ) from the training data and makes a decision $g_{n-1}$ using the remaining $n-1$ pairs. It tests for an an error on ( $X_{1}, Y_{1}$ ), and repeats this procedure for all $n$ pairs of the training data $D_{n}$. The estimate $L_{n}^{(D)}$ is the average the number of errors.

We formally denote the training set with $\left(X_{i}, Y_{i}\right)$ deleted by

$$
D_{n, i}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{i-1}, Y_{i-1}\right),\left(X_{i+1}, Y_{i+1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)
$$

Then we define

$$
L_{n}^{(D)}=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{g_{n-1}\left(X_{i}, D_{n, i}\right) \neq Y_{i}\right\}} .
$$

Clearly, the deleted estimate is almost unbiased in the sense that

$$
\mathbf{E} L_{n}^{(D)}=\mathbf{E} L_{n-1}
$$

Thus, $L_{n}^{(D)}$ should be viewed as an estimator of $L_{n-1}$, rather than of $L_{n}$. In most of the interesting cases $L_{n}$ converges with probability one so that the difference between $L_{n-1}$ and $L_{n}$ becomes negligible for large $n$.

The designer has the luxury of being able to pick the most convenient $g_{n-1}$. In some cases the choice is very natural, in other cases it is not. For example, if $g_{n}$ is the 1-nearest neighbor rule, then letting $g_{n-1}$ be the 1-nearest neighbor rule based on $n-1$ data pairs seems to be an obvious choice. We will see later
that this indeed yields an extremely good estimator. But what should $g_{n-1}$ be if $g_{n}$ is, for example, a $k$-NN rule? Should it be the $k$-nearest neighbor classifier, the $k$ - 1-nearest neighbor classifier, or maybe something else? Well, the choice is typically nontrivial and needs careful attention if the designer wants a distributionfree performance guarantee for the resulting estimate. Because of the variety of choices for $g_{n-1}$, we should not speak of the deleted estimate, but rather of $a$ deleted estimate.

In this chapter we analyze the performance of deleted estimates for a few prototype classifiers, such as the kernel, nearest neighbor, and histogram rules. In most cases studied here, deleted estimates have good distribution-free properties.

### 24.1 A General Lower Bound

We begin by exploring general limitations of error estimates for some important nonparametric rules. An error estimate $\widehat{L}_{n}$ is merely a function $\left(\mathcal{R}^{d} \times\{0,1\}\right)^{n} \rightarrow$ $[0,1]$, which is applied to the data $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$.

Theorem 24.1. Let $g_{n}$ be one of the following classifiers:
(a) The kernel rule

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} K\left(\frac{x-X_{i}}{h}\right) \geq \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} K\left(\frac{x-X_{i}}{h}\right) \\ 1 & \text { otherwise }\end{cases}
$$

with a nonnegative kernel $K$ of compact support and smoothing factor $h>0$.
(b) The histogram rule

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} I_{\left\{X_{i} \in A(x)\right\}} \leq \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} I_{\left\{X_{i} \in A(x)\right\}} \\ 1 & \text { otherwise }\end{cases}
$$

based on a fixed partition $\mathcal{P}=\left\{A_{1}, A_{2}, \ldots\right\}$ containing at least $n$ cells.
(c) The lazy histogram rule

$$
g_{n}(x)=Y_{j}, \quad x \in A_{i},
$$

where $X_{j}$ is the minimum-index point among $X_{1}, \ldots, X_{n}$ for which $X_{j} \in$ $A_{i}$, where $\mathcal{P}=\left\{A_{1}, A_{2}, \ldots\right\}$ is a fixed partition containing at least $n$ cells.
Denote the probability of error for $g_{n}$ by $L_{n}=\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y \mid D_{n}\right\}$. Then for every $n \geq 10$, and for every error estimator $\widehat{L}_{n}$, there exists a distribution of $(X, Y)$ with $L^{*}=0$ such that

$$
\mathbb{E}\left\{\left|\widehat{L}_{n}\left(D_{n}\right)-L_{n}\right|\right\} \geq \frac{1}{\sqrt{32 n}}
$$

The theorem says that for any estimate $\widehat{L}_{n}$, there exists a distribution with the property that $\mathbb{E}\left\{\left|\widehat{L}_{n}\left(D_{n}\right)-L_{n}\right|\right\} \geq 1 / \sqrt{32 n}$. For the rules $g_{n}$ given in the theorem,
no error estimate can possibly give better distribution-free performance guarantees. Error estimates are necessarily going to perform at least this poorly for some distributions.
Proof of Theorem 24.1. Let $\widehat{L}_{n}$ be an arbitrary fixed error estimate. The proof relies on randomization ideas similar to those of the proofs of Theorems 7.1 and 7.2. We construct a family of distributions for $(X, Y)$. For $b \in[0,1)$ let $b=0 . b_{1} b_{2} b_{3} \ldots$ be its binary expansion. In all cases the distribution of $X$ is uniform on $n$ points $\left\{x_{1}, \ldots, x_{n}\right\}$. For the histogram and lazy histogram rules choose $x_{1}, \ldots, x_{n}$ such that they fall into different cells. For the kernel rule choose $x_{1}, \ldots, x_{n}$ so that they are isolated from each other, that is, $K\left(\frac{x_{i}-x_{j}}{h}\right)=0$ for all $i, j \in\{1, \ldots, n\}, i \neq j$. To simplify the notation we will refer to these points by their indices, that is, we will write $X=i$ instead of $X=x_{i}$. For a fixed $b$ define

$$
Y=b_{X} .
$$

We may create infinitely many samples (one for each $b \in[0,1)$ ) drawn from the distribution of $(X, Y)$ as follows. Let $X_{1}, \ldots, X_{n}$ be i.i.d. and uniformly distributed on $\{1, \ldots, n\}$. All the samples share the same $X_{1}, \ldots, X_{n}$, but differ in their $Y_{i}$ 's. For given $X_{i}$, define $Y_{i}=b_{X_{i}}$. Write $Z_{n}=\left(X_{1}, \ldots, X_{n}\right)$ and $N_{i}=\sum_{j=1}^{n} I_{\left\{X_{j}=i\right\}}$. Observe that $D_{n}$ is a function of $Z_{n}$ and $b$. It is clear that for all classifiers covered by our assumptions, for a fixed $b$,

$$
L_{n}=\frac{1}{n} \sum_{i \in S} I_{\left\{b_{i}=1\right\}},
$$

where $S=\left\{i: 1 \leq i \leq n, N_{i}=0\right\}$ is the set of empty bins. We randomize the distribution as follows. Let $B=0 . B_{1} B_{2} \ldots$ be a uniform [ 0,1$]$ random variable independent of $(X, Y)$. Then clearly $B_{1}, B_{2}, \ldots$ are independent uniform binary random variables. Note that

$$
\begin{aligned}
\sup _{b} \mathbf{E}\left\{\left|\widehat{L}_{n}\left(D_{n}\right)-L_{n}\right|\right\} \geq & \mathbf{E}\left\{\left|\widehat{L}_{n}\left(D_{n}\right)-L_{n}\right|\right\} \\
& (\text { where } b \text { is replaced by } B) \\
= & \mathbf{E}\left\{\mathbf{E}\left\{\left|\widehat{L}_{n}\left(D_{n}\right)-L_{n}\right| \mid Z_{n}\right\}\right\} .
\end{aligned}
$$

In what follows we bound the conditional expectation within the brackets. To make the dependence upon $B$ explicit we write $\widehat{L}_{n}\left(Z_{n}, B\right)=\widehat{L}_{n}\left(D_{n}\right)$ and $L_{n}(B)=L_{n}$. Thus,

$$
\begin{aligned}
& \mathbf{E}\left\{\left|\widehat{L}_{n}\left(Z_{n}, B\right)-L_{n}(B)\right| \mid Z_{n}\right\} \\
&= \frac{1}{2} \mathbf{E}\left\{\left|\widehat{L}_{n}\left(Z_{n}, B\right)-L_{n}(B)\right|+\left|\widehat{L}_{n}\left(Z_{n}, B^{*}\right)-L_{n}\left(B^{*}\right)\right| \mid Z_{n}\right\} \\
& \text { (where } B_{i}^{*}=B_{i} \text { for } i \text { with } N_{i}>0 \text { and } B_{i}^{*}=1-B_{i} \text { for } i \text { with } N_{i}=0 \text { ) } \\
& \geq \frac{1}{2} \mathbf{E}\left\{\left|L_{n}(B)-L_{n}\left(B^{*}\right)\right| \mid Z_{n}\right\} \\
& \text { (since } \widehat{L}_{n}\left(Z_{n}, B^{*}\right)=\widehat{L}_{n}\left(Z_{n}, B\right) \text { ) }
\end{aligned}
$$

$$
=\frac{1}{2} \mathbf{E}\left\{\left.\frac{1}{n}\left|\sum_{i \in S}\left(I_{\left\{B_{i}=1\right\}}-I_{\left\{B_{i}=0\right\}}\right)\right| \right\rvert\, Z_{n}\right\}
$$

(recall the expression for $L_{n}$ given above)
$=\frac{1}{2 n} \mathbf{E}\{|2 B(|S|, 1 / 2)-1|| | S \mid\}$
(where $B(|S|, 1 / 2)$ is a binomial $(|S|, 1 / 2)$ random variable)
$\geq \frac{1}{2 n} \sqrt{\frac{|S|}{2}}$ (by Khintchine's inequality, see Lemma A.5).
In summary, we have

$$
\sup _{b} \mathbf{E}\left\{\left|\widehat{L}_{n}\left(D_{n}\right)-L_{n}\right|\right\} \geq \frac{1}{2 n} \mathbf{E}\left\{\sqrt{\frac{|S|}{2}}\right\}
$$

We have only to bound the right-hand side. We apply Lemma A. 4 to the random variable $\sqrt{|S|}=\sqrt{\sum_{i=1}^{n} I_{\left\{N_{i}=0\right\}}}$. Clearly, $\mathbf{E}|S|=n(1-1 / n)^{n}$, and

$$
\begin{aligned}
\mathbf{E}\left\{|S|^{2}\right\} & =\mathbf{E}\left\{\sum_{i=1}^{n} I_{\left\{N_{i}=0\right\}}+\sum_{i \neq j} I_{\left\{N_{i}=0, N_{j}=0\right\}}\right\} \\
& =\mathbf{E}|S|+n(n-1)\left(1-\frac{2}{n}\right)^{n}
\end{aligned}
$$

so that

$$
\begin{aligned}
\mathbf{E}\{\sqrt{|S|}\} & \geq \frac{\left(n\left(1-\frac{1}{n}\right)^{n}\right)^{3 / 2}}{\sqrt{n\left(1-\frac{1}{n}\right)^{n}+n(n-1)\left(1-\frac{2}{n}\right)^{n}}} \\
& \geq \sqrt{n} \frac{\left(1-\frac{1}{n}\right)^{n}}{\sqrt{\frac{1}{n}+\frac{(1-2 / n)^{n}}{(1-1 / n)^{n}}}} \\
& \geq \sqrt{n} \frac{\left(1-\frac{1}{n}\right)^{n}}{\sqrt{1-\frac{1}{n}-\frac{1}{n^{2}} \cdots \cdots}} \\
& \geq \sqrt{n}\left(1-\frac{1}{n}\right)^{n} \sqrt{e} \quad\left(\text { use }\left(1-\frac{1}{n}\right)^{n} \leq \frac{1}{e}\right) \\
& \geq \frac{1}{2} \sqrt{n} \quad(\text { for } n \geq 10) .
\end{aligned}
$$

The proof is now complete.

### 24.2 A General Upper Bound for Deleted Estimates

The following inequality is a general tool for obtaining distribution-free upper bounds for the difference between the deleted estimate and the true error probability $L_{n}$ :

Theorem 24.2. (Rogers and Wagner (1978); Devroye and Wagner (1976b)). Assume that $g_{n}$ is a symmetric classifier, that is, $g_{n}\left(x, D_{n}\right)=g_{n}\left(x, D_{n}^{\prime}\right)$, where $D_{n}^{\prime}$ is obtained by permuting the pairs of $D_{n}$ arbitrarily. Then

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{1}{n}+6 \mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq g_{n-1}\left(X, D_{n-1}\right)\right\} .
$$

Proof. First we express the three terms on the right-hand side of

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\}=\mathbf{E}\left\{L_{n}^{(D)^{2}}\right\}-2 \mathbf{E}\left\{L_{n}^{(D)} L_{n}\right\}+\mathbf{E}\left\{L_{n}^{2}\right\}
$$

The first term can be bounded, by using symmetry of $g_{n}$, by

$$
\begin{aligned}
& \mathbf{E}\left\{L_{n}^{(D)^{2}}\right\} \\
&= \mathbf{E}\left\{\left(\frac{1}{n} \sum_{i=1}^{n} I_{\left\{g_{n-1}\left(X_{i}, D_{n, i}\right) \neq Y_{i}\right\}}\right)^{2}\right\} \\
&= \mathbf{E}\left\{\frac{1}{n^{2}} \sum_{i=1}^{n} I_{\left\{g_{n-1}\left(X_{i}, D_{n, i}\right) \neq Y_{i}\right\}}\right\} \\
&+\mathbf{E}\left\{\frac{1}{n^{2}} \sum_{i \neq j} I_{\left\{g_{n-1}\left(X_{i}, D_{n, i}\right) \neq Y_{i}\right\}} I_{\left\{g_{n-1}\left(X_{j}, D_{n, j}\right) \neq Y_{j}\right\}}\right\} \\
&= \frac{\mathbf{E} L_{n-1}}{n}+\frac{n-1}{n} \mathbf{P}\left\{g_{n-1}\left(X_{1}, D_{n, 1}\right) \neq Y_{1}, g_{n-1}\left(X_{2}, D_{n, 2}\right) \neq Y_{2}\right\} \\
& \leq \frac{1}{n}+\mathbf{P}\left\{g_{n-1}\left(X_{1}, D_{n, 1}\right) \neq Y_{1}, g_{n-1}\left(X_{2}, D_{n, 2}\right) \neq Y_{2}\right\} .
\end{aligned}
$$

The second term is written as

$$
\begin{aligned}
\mathbf{E} & \left\{L_{n}^{(D)} L_{n}\right\} \\
& =\mathbf{E}\left\{L_{n} \frac{1}{n} \sum_{i=1}^{n} I_{\left\{g_{n-1}\left(X_{i}, D_{n, i}\right) \neq Y_{i}\right\}}\right\} \\
& =\frac{1}{n} \sum_{i=1}^{n} \mathbf{E}\left\{L_{n} I_{\left\{g_{n-1}\left(X_{i}, D_{n, i}\right) \neq Y_{i}\right\}}\right\} \\
& =\frac{1}{n} \sum_{i=1}^{n} \mathbf{E}\left\{\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y \mid D_{n}\right\} \mathbf{P}\left\{g_{n-1}\left(X_{i}, D_{n, i}\right) \neq Y_{i} \mid D_{n}\right\}\right\}
\end{aligned}
$$

$$
\begin{aligned}
& =\frac{1}{n} \sum_{i=1}^{n} \mathbf{E}\left\{\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y, g_{n-1}\left(X_{i}, D_{n, i}\right) \neq Y_{i} \mid D_{n}\right\}\right\} \\
& =\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y, g_{n-1}\left(X_{1}, D_{n, 1}\right) \neq Y_{1}\right\}
\end{aligned}
$$

For the third term, we introduce the pair $\left(X^{\prime}, Y^{\prime}\right)$, independent of $X, Y$, and $D_{n}$, having the same distribution as $(X, Y)$. Then

$$
\begin{aligned}
\mathbf{E}\left\{L_{n}^{2}\right\} & =\mathbf{E}\left\{\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y \mid D_{n}\right\}^{2}\right\} \\
& =\mathbf{E}\left\{\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y \mid D_{n}\right\} \mathbf{P}\left\{g_{n}\left(X^{\prime}, D_{n}\right) \neq Y^{\prime} \mid D_{n}\right\}\right\} \\
& =\mathbf{E}\left\{\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y, g_{n}\left(X^{\prime}, D_{n}\right) \neq Y^{\prime} \mid D_{n}\right\}\right\} \\
& =\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y, g_{n}\left(X^{\prime}, D_{n}\right) \neq Y^{\prime}\right\},
\end{aligned}
$$

where we used independence of $\left(X^{\prime}, Y^{\prime}\right)$.
We introduce the notation

$$
\begin{aligned}
D_{n}^{\prime} & =\left(X_{3}, Y_{3}\right), \ldots,\left(X_{n}, Y_{n}\right) \\
A_{k, i, j} & =\left\{g_{n}\left(X_{k} ;\left(X_{i}, Y_{i}\right),\left(X_{j}, Y_{j}\right), D_{n}^{\prime}\right) \neq Y_{k}\right\}, \\
B_{k, i} & =\left\{g_{n-1}\left(X_{k} ;\left(X_{i}, Y_{i}\right), D_{n}^{\prime}\right) \neq Y_{k}\right\}
\end{aligned}
$$

and we formally replace $(X, Y)$ and $\left(X^{\prime}, Y^{\prime}\right)$ by $\left(X_{\alpha}, Y_{\alpha}\right)$ and $\left(X_{\beta}, Y_{\beta}\right)$ so that we may work with the indices $\alpha$ and $\beta$. With this notation, we have shown thus far the following:

$$
\begin{aligned}
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{1}{n} & +\mathbf{P}\left\{A_{\alpha, 1,2}, A_{\beta, 1,2}\right\}-\mathbf{P}\left\{A_{\alpha, 1,2}, B_{1,2}\right\} \\
& +\mathbf{P}\left\{B_{1,2}, B_{2,1}\right\}-\mathbf{P}\left\{A_{\alpha, 1,2}, B_{2,1}\right\}
\end{aligned}
$$

Note that

$$
\begin{aligned}
& \mathbf{P}\left\{A_{\alpha, 1,2}, A_{\beta, 1,2}\right\}-\mathbf{P}\left\{A_{\alpha, 1,2}, B_{1,2}\right\} \\
&= \mathbf{P}\left\{A_{\alpha, 1,2}, A_{\beta, 1,2}\right\}-\mathbf{P}\left\{A_{\alpha, \beta, 2}, B_{\beta, 2}\right\} \\
&(\text { by symmetry }) \\
&= \mathbf{P}\left\{A_{\alpha, 1,2}, A_{\beta, 1,2}\right\}-\mathbf{P}\left\{A_{\alpha, \beta, 2}, A_{\beta, 1,2}\right\} \\
&+\mathbf{P}\left\{A_{\alpha, \beta, 2}, A_{\beta, 1,2}\right\}-\mathbf{P}\left\{A_{\alpha, \beta, 2}, B_{\beta, 2}\right\} \\
&= I+I I .
\end{aligned}
$$

Also,

$$
\begin{aligned}
& \mathbf{P}\left\{B_{1,2}, B_{2,1}\right\}-\mathbf{P}\left\{A_{\alpha, 1,2}, B_{1,2}\right\} \\
& = \\
& \quad \mathbf{P}\left\{B_{\alpha, \beta}, B_{\beta, \alpha}\right\}-\mathbf{P}\left\{A_{\alpha, \beta, 2}, B_{\beta, 2}\right\} \\
& \quad \text { (by symmetry) }
\end{aligned}
$$

$$
\begin{aligned}
= & \mathbf{P}\left\{B_{\alpha, \beta}, B_{\beta, \alpha}\right\}-\mathbf{P}\left\{A_{\alpha, \beta, 2}, B_{\beta, \alpha}\right\} \\
& +\mathbf{P}\left\{A_{\alpha, \beta, 2}, B_{\beta, \alpha}\right\}-\mathbf{P}\left\{A_{\alpha, \beta, 2}, B_{\beta, 2}\right\} \\
= & I I I+I V .
\end{aligned}
$$

Using the fact that for events $\left\{C_{i}\right\},\left|\mathbf{P}\left\{C_{i}, C_{j}\right\}-\mathbf{P}\left\{C_{i}, C_{k}\right\}\right| \leq \mathbf{P}\left\{C_{j} \triangle C_{k}\right\}$, we bound

$$
\begin{aligned}
I & \leq \mathbf{P}\left\{A_{\alpha, 1,2} \triangle A_{\alpha, \beta, 2}\right\} \\
I I & \leq \mathbf{P}\left\{A_{\beta, 1,2} \triangle B_{\beta, 2}\right\} \stackrel{\text { def }}{=} V \\
I I I & \leq \mathbf{P}\left\{B_{\alpha, \beta} \triangle A_{\alpha, \beta, 2}\right\}=V \\
I V & \leq \mathbf{P}\left\{B_{\beta, \alpha} \triangle B_{\beta, 2}\right\}
\end{aligned}
$$

The upper bounds for II and III are identical by symmetry. Also,

$$
I \leq \mathbf{P}\left\{A_{\alpha, 1,2} \triangle B_{\alpha, 2}\right\}+\mathbf{P}\left\{B_{\alpha, 2} \triangle A_{\alpha, \beta, 2}\right\}=2 V
$$

and

$$
I V \leq \mathbf{P}\left\{B_{\beta, \alpha} \triangle A_{\beta, \alpha, 2}\right\}+\mathbf{P}\left\{A_{\beta, \alpha, 2} \triangle B_{\beta, 2}\right\}=2 V
$$

for a grand total of 6 V . This concludes the proof.

### 24.3 Nearest Neighbor Rules

Theorem 24.2 can be used to obtain distribution-free upper bounds for specific rules. Here is the most important example.

Theorem 24.3. (ROGERS AND WAGNER (1978)). Let $g_{n}$ be the $k$-nearest neighbor rule with randomized tie-breaking. If $L_{n}^{(D)}$ is the deleted estimate with $g_{n-1}$ chosen as the $k-\mathrm{NN}$ rule (with the same $k$ and with the same randomizing random variables), then

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{6 k+1}{n}
$$

Proof. Because of the randomized tie-breaking, the $k$-nN rule is symmetric, and Theorem 24.2 is applicable. We only have to show that

$$
\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq g_{n-1}\left(X, D_{n-1}\right)\right\} \leq \frac{k}{n} .
$$

Clearly, $g_{n}\left(X, D_{n}\right) \neq g_{n-1}\left(X, D_{n-1}\right)$ can happen only if $X_{n}$ is among the $k$ nearest neighbors of $X$. But the probability of this event is just $k / n$, since by symmetry, all points are equally likely to be among the $k$ nearest neighbors.

REMARK. If $g_{n}$ is the $k$-NN rule such that distance ties are broken by comparing indices, then $g_{n}$ is not symmetric, and Theorem 24.2 is no longer applicable (unless, e.g., $X$ has a density). Another nonsymmetric classifier is the lazy histogram rule.

Remark. Applying Theorem 24.3 to the $1-$ nn rule, the Cauchy-Schwarz inequality implies $\mathbf{E}\left|L_{n}^{(D)}-L_{n}\right| \leq \sqrt{7 / n}$ for all distributions.

Remark. Clearly, the inequality of Theorem 24.3 holds for any rule that is some function of the $k$ nearest points. For the $k$-nearest neighbor rule, with a more careful analysis, Devroye and Wagner (1979a) improved Theorem 24.3 to

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{1}{n}+\frac{24 \sqrt{k}}{n \sqrt{2 \pi}}
$$

(see Problem 24.8).

Probability inequalities for $\left|L_{n}^{(D)}-L_{n}\right|$ can also be obtained with further work. By Chebyshev's inequality we immediately get

$$
\mathbf{P}\left\{\left|L_{n}^{(D)}-L_{n}\right|>\epsilon\right\} \leq \frac{\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\}}{\epsilon^{2}},
$$

so that the above bounds on the expected squared error can be used. Sharper distribution-free inequalities were obtained by Devroye and Wagner (1979a; 1979b) for several nonparametric rules. Here we present a result that follows immediately from what we have already seen:

Theorem 24.4. Consider the $k$-nearest neighbor rule with randomized tie-breaking. If $L_{n}^{(D)}$ is the deleted estimate with $g_{n-1}$ chosen as the $k$-NN rule with the same tie-breaking, then

$$
\mathbf{P}\left\{\left|L_{n}^{(D)}-\mathbf{E} L_{n-1}\right|>\epsilon\right\} \leq 2 e^{-n \epsilon^{2} /\left(k^{2} \gamma_{d}^{2}\right)}
$$

Proof. The result follows immediately from McDiarmid's inequality by the following argument: from Lemma 11.1, given $n$ points in $\mathcal{R}^{d}$, a particular point can be among the $k$ nearest neighbors of at most $k \gamma_{d}$ points. To see this, just set $\mu$ equal to the empirical measure of the $n$ points in Lemma 11.1. Therefore, changing the value of one pair from the training data can change the value of the estimate by at most $2 k \gamma_{d}$. Now, since $\mathbf{E} L_{n}^{(D)}=\mathbf{E} L_{n-1}$, Theorem 9.1 yields the result.

Exponential upper bounds for the probability $\mathbf{P}\left\{\left|L_{n}^{(D)}-L_{n}\right|>\epsilon\right\}$ are typically much harder to obtain. We mention one result without proof.

Theorem 24.5. (Devroye and Wagner (1979a)). For the $k$-nearest neighbor rule,

$$
\mathbf{P}\left\{\left|L_{n}^{(D)}-L_{n}\right|>\epsilon\right\} \leq 2 e^{-n \epsilon^{2} / 18}+6 e^{-n \epsilon^{3} /\left(108 k\left(\gamma_{d}+2\right)\right)}
$$

One of the drawbacks of the deleted estimate is that it requires much more computation than the resubstitution estimate. If conditional on $Y=0$ and $Y=1$, $X$ is gaussian, and the classification rule is the appropriate parametric rule, then the estimate can be computed quickly. See Lachenbruch and Mickey (1968), Fukunaga and Kessel (1971), and McLachlan (1992) for further references.

Another, and probably more serious, disadvantage of the deleted estimate is its large variance. This fact can be illustrated by the following example from Devroye and Wagner (1979b): let $n$ be even, and let the distribution of $(X, Y)$ be such that $Y$ is independent of $X$ with $\mathbf{P}\{Y=0\}=\mathbf{P}\{Y=1\}=1 / 2$. Consider the $k$-nearest neighbor rule with $k=n-1$. Then obviously, $L_{n}=1 / 2$. Clearly, if the number of zeros and ones among the labels $Y_{1}, \ldots, Y_{n}$ are equal, then $L_{n}^{(D)}=1$. Thus, for $0<\epsilon<1 / 2$,

$$
\mathbf{P}\left\{\left|L_{n}^{(D)}-L_{n}\right|>\epsilon\right\} \geq \mathbf{P}\left\{\sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}}=\frac{n}{2}\right\}=\frac{1}{2^{n}}\binom{n}{n / 2} .
$$

By Stirling's formula (Lemma A.3), we have

$$
\mathbf{P}\left\{\left|L_{n}^{(D)}-L_{n}\right|>\epsilon\right\} \geq \frac{1}{\sqrt{2 \pi n}} \frac{1}{e^{1 / 12}}
$$

Therefore, for this simple rule and certain distributions, the probability above can not decrease to zero faster than $1 / \sqrt{n}$. Note that in the example above, $\mathbf{E} L_{n}^{(D)}=$ $\mathbf{E} L_{n-1}=1 / 2$, so the lower bound holds for $\mathbf{P}\left\{\left|L_{n}^{(D)}-\mathbf{E} L_{n}^{(D)}\right|>\epsilon\right\}$ as well. Also, in this example, we have

$$
\mathbb{E}\left\{\left(L_{n}^{(D)}-\mathbf{E} L_{n}^{(D)}\right)^{2}\right\} \geq \frac{1}{4} \mathbf{P}\left\{\sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}}=\frac{n}{2}\right\} \geq \frac{1}{4 \sqrt{2 \pi n}} \frac{1}{e^{1 / 12}}
$$

In Chapter 31 we describe other estimates with much smaller variances.

### 24.4 Kernel Rules

Theorem 24.2 may also be used to obtain tight distribution-free upper bounds for the performance of the deleted estimate of the error probability of kernel rules. We have the following bound:

Theorem 24.6. Assume that $K \geq 0$ is a regular kernel of bounded support, that is, it is a function satisfying

$$
\begin{align*}
\text { (i) } & K(x) \geq \beta, \quad\|x\| \leq \rho  \tag{i}\\
\text { (ii) } & K(x) \leq B \\
\text { (iii) } & K(x)=0, \quad\|x\|>R
\end{align*}
$$

(ii) $K(x) \leq B$,
for some positive finite constants $\beta, \rho, B$ and $R$. Let the kernel rule be defined by

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} K\left(x-X_{i}\right) \geq \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} K\left(x-X_{i}\right) \\ 1 & \text { otherwise }\end{cases}
$$

and define $g_{n-1}$ similarly. Then there exist constants $C_{1}(d)$ depending upon $d$ only and $C_{2}(K)$ depending upon $K$ only such that for all $n$,

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{1}{n}+\frac{C_{1}(d) C_{2}(K)}{\sqrt{n}}
$$

One may take $C_{2}(K)=6(1+R / \rho)^{d / 2} \min (2, B / \beta)$.

REMARK. Since $C_{2}(K)$ is a scale-invariant factor, the theorem applies to the rule with $K(u)$ replaced by $K_{h}(u)=\frac{1}{h} K\left(\frac{u}{h}\right)$ for any smoothing factor. As it is, $C_{2}(K)$ is minimal and equal to 12 if we let $K$ be the uniform kernel on the unit ball ( $R=\rho, B=\beta$ ). The assumptions of the theorem require that $g_{n-1}$ is defined with the same kernel and smoothing factor as $g_{n}$.

REMARK. The theorem applies to virtually any kernel of compact support that is of interest to the practitioners. Note, however, that the gaussian kernel is not covered by the result. The theorem generalizes an earlier result of Devroye and Wagner (1979b), in which a more restricted class of kernels was considered. They showed that if $K$ is the uniform kernel then

$$
\mathbb{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{1}{2 n}+\frac{24}{\sqrt{n}}
$$

See Problem 24.4.

We need the following auxiliary inequality, which we quote without proof:
Lemma 24.1. (PETROV (1975), P.44). Let $Z_{1}, \ldots, Z_{n}$ be real-valued i.i.d. random variables. For $\epsilon>0$,

$$
\sup _{x} \mathbb{P}\left\{\sum_{i=1}^{n} Z_{i} \in[x, x+\epsilon]\right\} \leq \frac{C \epsilon}{\sqrt{n}} \frac{1}{\sqrt{\sup _{0<\lambda \leq \epsilon} \lambda^{2} \mathbf{P}\left\{\left|Z_{1}\right| \geq \lambda / 2\right\}}}
$$

where $C$ is a universal constant.
Corollary 24.1. Let $Z_{1}, \ldots, Z_{n}$ be real-valued i.i.d. random variables. For $\epsilon \geq$ $\lambda>0$,

$$
\mathbb{P}\left\{\left|\sum_{i=1}^{n} Z_{i}\right|<\frac{\epsilon}{2}\right\} \leq \frac{C}{\sqrt{n}} \frac{\epsilon}{\lambda} \frac{1}{\sqrt{\mathbb{P}\left\{\left|Z_{1}\right| \geq \lambda / 2\right\}}}
$$

where $C$ is a universal constant.

Proof of Theorem 24.6. We apply Theorem 24.2, by finding an upper bound for

$$
\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq g_{n-1}\left(X, D_{n-1}\right)\right\}
$$

For the kernel rule with kernel $K \geq 0$, in which $h$ is absorbed,

$$
\begin{aligned}
& \mathbb{P}\left\{g_{n}\left(X, D_{n}\right) \neq g_{n-1}\left(X, D_{n-1}\right)\right\} \\
& \quad \leq \mathbb{P}\left\{\left|\sum_{i=1}^{n-1}\left(2 Y_{i}-1\right) K\left(X-X_{i}\right)\right| \leq K\left(X-X_{n}\right), K\left(X-X_{n}\right)>0\right\}
\end{aligned}
$$

Define $B^{\prime}=\max (2 \beta, B)$. We have

$$
\begin{aligned}
& \mathbb{P}\left\{\left|\sum_{i=1}^{n-1}\left(2 Y_{i}-1\right) K\left(X-X_{i}\right)\right| \leq K\left(X-X_{n}\right), K\left(X-X_{n}\right)>0\right\} \\
& \quad \leq \mathbb{P}\left\{\left|\sum_{i=1}^{n-1}\left(2 Y_{i}-1\right) K\left(X-X_{i}\right)\right| \leq B^{\prime}, X_{n} \in S_{X, R}\right\}
\end{aligned}
$$

(where $S_{X, R}$ is the ball of radius $R$ centered at $X$ )

$$
=\mathbf{E}\left\{I_{\left\{X_{n} \in S_{X, R}\right\}} \frac{2 C B^{\prime}}{2 \beta \sqrt{n} \sqrt{\mathbf{P}\left\{\left|\left(2 Y_{1}-1\right) K\left(X-X_{1}\right)\right| \geq \beta \mid X\right\}}}\right\}
$$

(by Corollary 24.1 , since $2 \beta \leq B^{\prime}$ )
$\leq \mathbb{E}\left\{I_{\left\{X_{n} \in S_{X, R}\right\}} \frac{C B^{\prime}}{\beta \sqrt{n \mu\left(S_{X, \rho}\right)}}\right\}$
(recall that $K(u) \geq \beta$ for $\|u\| \leq \rho$,
so that $\left.\mathbf{P}\left\{\left|\left(2 Y_{1}-1\right) K\left(x-X_{1}\right)\right| \geq \beta\right\} \geq \mathbf{P}\left\{X_{1} \geq S_{x, \rho}\right\}\right)$
$=\frac{C B^{\prime}}{\beta \sqrt{n}} \iint_{S_{x, R}} \frac{\mu(d y)}{\sqrt{\mu\left(S_{x, \rho}\right)}} \mu(d x)$
$\leq \frac{C c_{d} B^{\prime}}{\beta \sqrt{n}}\left(1+\frac{R}{\rho}\right)^{d / 2}$,
where we used Lemma 10.2. The constant $c_{d}$ depends upon the dimension only.

### 24.5 Histogram Rules

In this section we discuss properties of the deleted estimate of the error probability of histogram rules. Let $\mathcal{P}=\left\{A_{1}, A_{2}, \ldots\right\}$ be a partition of $\mathcal{R}^{d}$, and let $g_{n}$ be the corresponding histogram classifier (see Chapters 6 and 9). To get a performance bound for the deleted estimate, we can simply apply Theorem 24.2.

Theorem 24.7. For the histogram rule $g_{n}$ corresponding to any partition $\mathcal{P}$, and for all $n$,

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{1}{n}+6 \sum_{i} \frac{\mu\left(A_{i}\right)^{3 / 2}}{\sqrt{\pi(n-1)}}+6 \sum_{i} \mu^{2}\left(A_{i}\right) e^{-n \mu\left(A_{i}\right)},
$$

and in particular,

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{1+6 / e}{n}+\frac{6}{\sqrt{\pi(n-1)}} .
$$

Proof. The first inequality follows from Theorem 24.2 if we can find an upper bound for $\mathbf{P}\left\{g_{n}(X) \neq g_{n-1}(X)\right\}$. We introduce the notation

$$
\nu_{0, n}(A)=\frac{1}{n} \sum_{j=1}^{n} I_{\left\{Y_{j}=0\right\}} I_{\left\{X_{j} \in A\right\}} .
$$

Clearly, $g_{n-1}(X)$ can differ from $g_{n}(X)$ only if both $X_{n}$ and $X$ fall in the same cell of the partition, and if the number of zeros in the cell is either equal, or less by one than the number of ones. Therefore, by independence, we have

$$
\begin{aligned}
& \mathbf{P}\left\{g_{n}(X) \neq g_{n-1}(X)\right\} \\
& \quad=\sum_{i} \mathbf{P}\left\{g_{n}(X) \neq g_{n-1}(X) \mid X \in A_{i}, X_{n} \in A_{i}\right\} \mu\left(A_{i}\right)^{2} \\
& \quad \leq \sum_{i} \mathbf{P}\left\{\left.v_{0, n-1}\left(A_{i}\right)=\left\lfloor\frac{\mu_{n-1}\left(A_{i}\right)}{2}\right\rfloor \right\rvert\, X \in A_{i}, X_{n} \in A_{i}\right\} \mu\left(A_{i}\right)^{2} .
\end{aligned}
$$

The terms in the sum above may be bounded as follows:

$$
\begin{aligned}
& \mathbf{P}\left\{\left.\nu_{0, n-1}(A)=\left\lfloor\frac{\mu_{n-1}(A)}{2}\right\rfloor \right\rvert\, X \in A, X_{n} \in A\right\} \\
&= \mathbf{P}\left\{\nu_{0, n-1}(A)=\left\lfloor\frac{\mu_{n-1}(A)}{2}\right\rfloor\right\} \quad \text { (by independence) } \\
& \leq \mathbf{P}\left\{\mu_{n-1}(A)=0\right\} \\
&+\mathbf{E}\left\{\mathbf{P}\left\{\left.\nu_{0, n-1}(A)=\left\lfloor\frac{\mu_{n-1}(A)}{2}\right\rfloor \right\rvert\, X_{1}, \ldots, X_{n}\right\} I_{\left\{\mu_{n-1}(A)>0\right\}}\right\} \\
& \leq(1-\mu(A))^{n}+\mathbf{E}\left\{\frac{1}{\sqrt{2 \pi(n-1) \mu_{n-1}(A)}} I_{\left\{\mu_{n-1}(A)>0\right\}}\right\}
\end{aligned}
$$

(by Lemma A.3)
$\leq(1-\mu(A))^{n}+\sqrt{\mathbf{E}\left\{\frac{I_{\left\{\mu_{n-1}(A)>0\right\}}}{2 \pi(n-1) \mu_{n-1}(A)}\right\}}$ (by Jensen's inequality)
$\leq e^{-n \mu(A)}+\sqrt{\frac{1}{\pi(n-1) \mu(A)}}$,
where in the last step we use Lemma A.2. This concludes the proof of the first inequality. The second one follows trivially by noting that $x e^{-x} \leq 1 / e$ for all $x$.

REMARK. It is easy to see that the inequalities of Theorem 24.7 are tight, up to a constant factor, in the sense that for any partition $\mathcal{P}$, there exists a distribution such that

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \geq \frac{1}{4 \sqrt{2 \pi n}} \frac{1}{e^{1 / 12}}
$$

(see Problem 24.5).

REMARK. The second inequality in Theorem 23.3 points out an important difference between the behavior of the resubstitution and the deleted estimates for histogram rules. As mentioned above, for some distributions the variance of $L_{n}^{(D)}$ can be of the order $1 / \sqrt{n}$. This should be contrasted with the much smaller variance of the resubstitution estimate. The small variance of $L_{n}^{(R)}$ comes often with a larger bias. Other types of error estimates with small variance are discussed in Chapter 31.

REMARK. Theorem 24.3 shows that for any partition,

$$
\sup _{(X, Y)} \mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\}=O\left(\frac{1}{\sqrt{n}}\right) .
$$

On the other hand, if $k=o(\sqrt{n})$, where $k$ is the number of cells in the partition, then for the resubstitution estimate we have a better guaranteed distribution-free performance:

$$
\sup _{(X, Y)} \mathbf{E}\left\{\left(L_{n}^{(R)}-L_{n}\right)^{2}\right\}=o\left(\frac{1}{\sqrt{n}}\right) .
$$

At first sight, the resubstitution estimate seems preferable to the deleted estimate. However, if the partition has a large number of cells, $L_{n}^{(R)}$ may be off the mark; see Theorem 23.4.

## Problems and Exercises

Problem 24.1. Show the following variant of Theorem 24.2: For all symmetric classifiers

$$
\begin{aligned}
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{1}{n} & +2 \mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq g_{n-1}\left(X, D_{n-1}\right)\right\} \\
& +\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq g_{n}\left(X, D_{n}^{*}\right)\right\} \\
& +\mathbf{P}\left\{g_{n-1}\left(X, D_{n-1}\right) \neq g_{n-1}\left(X, D_{n-1}^{*}\right)\right\}
\end{aligned}
$$

where $D_{n}^{*}$ and $D_{n-1}^{*}$ are just $D_{n}$ and $D_{n-1}$ with $\left(X_{1}, Y_{1}\right)$ replaced by an independent copy ( $X_{0}, Y_{0}$ ).

Problem 24.2. Let $g_{n}$ be the relabeling nN rule with the $k$-NN classifier as ancestral rule, as defined in Chapter 11. Provide an upper bound for the squared error $\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\}$ of the deleted estimate.

Problem 24.3. Let $g_{n}$ be the rule obtained by choosing the best $k \leq k_{0}$ in the $k$-nN rule ( $k_{0}$ is a constant) by minimizing the standard deleted estimate $L_{n}^{(D)}$ with respect to $k$. How would you estimate the probability of error for this rule? Give the best possible distribution-free performance guarantees you can find.

Problem 24.4. Consider the kernel rule with the window kernel $K=S_{0,1}$. Show that

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{1+6 / e}{n}+\frac{6}{\sqrt{\pi(n-1)}}
$$

Hint: Follow the line of the proof of Theorem 24.7.
Problem 24.5. Show that for any partition $\mathcal{P}$, there exists a distribution such that for the deleted estimate of the error probability of the corresponding histogram rule,

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \geq \frac{1}{4 \sqrt{2 \pi n}} \frac{1}{e^{1 / 12}}
$$

Hint: Proceed as in the proof of the similar inequality for $k$-nearest neighbor rules.
Problem 24.6. Consider the $k$-spacings method (see Chapter 21). We estimate the probability of error $\left(L_{n}\right)$ by a modified deleted estimate $\widehat{L}_{n}$ as follows:

$$
\widehat{L}_{n}=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{Y_{i} \neq g_{n, i}\left(X_{i}, D_{n}\right)\right\}},
$$

where $g_{n, i}$ is a histogram rule based upon the same $k$-spacings partition used for $g_{n}$-that is, the partition determined by $k$-spacings of the data points $X_{1}, \ldots, X_{n}$-but in which a majority vote is based upon the $Y_{j}$ 's in the same cell of the partition with $Y_{i}$ deleted. Show that

$$
\mathbf{E}\left\{\left(\widehat{L}_{n}-L_{n}\right)\right\} \leq \frac{6 k+1}{n}
$$

Hint: Condition on the $X_{i}$ 's, and verify that the inequality of Theorem 24.2 remains valid:
Problem 24.7. Consider a rule in which we rank the real-valued observations $X_{1}, \ldots, X_{n}$ from small to large, to obtain $X_{(1)}, \ldots, X_{(n)}$. Assume that $X_{1}$ has a density. Derive an inequality for the error $\left|L_{n}^{(D)}-L_{n}\right|$ for some deleted estimate $L_{n}^{(D)}$ (of your choice), when the rule is defined by a majority vote over the data-dependent partition

$$
\left(-\infty, X_{(1)}\right],\left(X_{(1)}, X_{(2)}\right],\left(X_{(3)}, X_{(5)}\right],\left(X_{(6)}, X_{(9)}\right], \ldots
$$

defined by $1,2,3,4, \ldots$ points, respectively.

Problem 24.8. Prove that for the $k$-nearest neighbor rule,

$$
\mathbf{E}\left\{\left(L_{n}^{(D)}-L_{n}\right)^{2}\right\} \leq \frac{1}{n}+\frac{24 \sqrt{k}}{n \sqrt{2 \pi}}
$$

(Devroye and Wagner (1979a)). Hint: Obtain a refined upper bound for

$$
\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq g_{n-1}\left(X, D_{n-1}\right)\right\},
$$

using techniques not unlike those of the proof of Theorem 24.7.
Problem 24.9. Open-ended problem. Investigate if Theorem 24.6 can be extended to kernels with unbounded support such as the gaussian kernel.

## 25

## Automatic Kernel Rules

We saw in Chapter 10 that for a large class of kernels, if the smoothing parameter $h$ converges to zero such that $n h^{d}$ goes to infinity as $n \rightarrow \infty$, then the kernel classification rule is universally consistent. For a particular $n$, asymptotic results provide little guidance in the selection of $h$. On the other hand, selecting the wrong value of $h$ may lead to catastrophic error rates - in fact, the crux of every nonparametric estimation problem is the choice of an appropriate smoothing factor. It tells us how far we generalize each data point $X_{i}$ in the space. Purely atomic distributions require little smoothing ( $h=0$ will generally be fine), while distributions with densities require a lot of smoothing. As there are no simple tests for verifying whether the data are drawn from an absolutely continuous distribution-let alone a distribution with a Lipschitz density-it is important to let the data $D_{n}$ determine $h$. A data-dependent smoothing factor is merely a mathematical function $H_{n}:\left(\mathcal{R}^{d} \times\{0,1\}\right)^{n} \rightarrow[0, \infty)$. For brevity, we will simply write $H_{n}$ to denote the random variable $H_{n}\left(D_{n}\right)$. This chapter develops results regarding such functions $H_{n}$.
This chapter is not a luxury but a necessity. Anybody developing software for pattern recognition must necessarily let the data do the talking-in fact, good universally applicable programs can have only data-dependent parameters.
Consider the family of kernel decision rules $g_{n}$ and let the smoothing factor $h$ play the role of parameter. The best parameter ( $H_{\mathrm{opr}}$ ) is the one that minimizes $L_{n}$. Unfortunately, it is unknown, as is $L_{\text {oft }}$, the corresponding minimal probability of error. The first goal of any data-dependent smoothing factor $H_{n}$ should be to approach the performance of $H_{\text {ort }}$. We are careful here to avoid saying that $H_{n}$ should be close to $H_{\text {ort }}$, as closeness of smoothing factors does not necessarily imply closeness of error probabilities and vice versa. Guarantees one might want
in this respect are

$$
\mathbf{E}\left\{L_{n}-L_{\mathrm{OPT}}\right\} \leq \alpha_{n}
$$

for some suitable sequence $\alpha_{n} \rightarrow 0$, or better still,

$$
\mathbf{E}\left\{L_{n}-L^{*}\right\} \leq\left(1+\beta_{n}\right) \mathbf{E}\left\{L_{\mathrm{OPT}}-L^{*}\right\}
$$

for another sequence $\beta_{n} \rightarrow 0$. But before one even attempts to develop such datadependent smoothing factors, one's first concern should be with consistency: is it true that with the given $H_{n}, L_{n} \rightarrow L^{*}$ in probability or with probability one? This question is dealt with in the next section. In subsequent sections, we give various examples of data-dependent smoothing factors.

### 25.1 Consistency

We start with consistency results that generalize Theorem 10.1. The first result assumes that the value of the smoothing parameter is picked from a discrete set.

Theorem 25.1. Assume that the random variable $H_{n}$ takes its values from the set of real numbers of the form $\frac{1}{\left(1+\delta_{n}\right)^{k}}$, where $k$ is a nonnegative integer and $\delta_{n}>0$. Let $K$ be a regular kernel function. (Recall Definition 10.1.) Define the kernel classification rule corresponding to the random smoothing parameter $H_{n}$ by

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{i=1}^{n} I_{\left\{Y_{i}=0\right\}} K\left(\frac{x-X_{i}}{H_{n}}\right) \geq \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} K\left(\frac{x-X_{i}}{H_{n}}\right) \\ 1 & \text { otherwise }\end{cases}
$$

If

$$
\frac{1}{\delta_{n}}=e^{o(n)}
$$

and

$$
H_{n} \rightarrow 0 \text { and } n H_{n}^{d} \rightarrow \infty \quad \text { with probability one as } n \rightarrow \infty
$$

then $L\left(g_{n}\right) \rightarrow L^{*}$ with probability one, that is, $g_{n}$ is strongly universally consistent.
Proof. The theorem is a straightforward extension of Theorem 10.1. Clearly, $L\left(g_{n}\right) \rightarrow L^{*}$ with probability one if and only if for every $\epsilon>0, I_{\left\{L\left(g_{n}\right)-L^{*}>\epsilon\right\}} \rightarrow 0$ with probability one. Now, for any $\beta>0$,

$$
I_{\left\{L\left(g_{n}\right)-L^{*}>\epsilon\right\}} \leq I_{\left\{1 / H_{n}>\beta, n H_{n}^{d}>\beta, L\left(g_{n}\right)-L^{*}>\epsilon\right\}}+I_{\left\{1 / H_{n} \leq \beta\right\}}+I_{\left\{n H_{n}^{d} \leq \beta\right\}} .
$$

We have to show that the random variables on the right-hand side converge to zero with probability one. The convergence of the second and third terms follows from the conditions on $H_{n}$. The convergence of the first term follows from Theorem 10.1 , since it states that for any $\epsilon>0$, there exist $\beta>0$ and $n_{0}$ such that for the error probability $L_{n, k}$ of the kernel rule with smoothing parameter $h=\frac{1}{\left(1+\delta_{n}\right)^{k}}$,

$$
\mathbf{P}\left\{L_{n, k}-L^{*}>\epsilon\right\} \leq 4 e^{-c n \epsilon^{2}}
$$

for some constant $c$ depending on the dimension only, provided that $n>n_{0}$, $h<1 / \beta$ and $n h^{d}>\beta$. Now clearly,

$$
\begin{aligned}
\mathbb{P} & \left\{L_{n}-L^{*}>\epsilon, 1 / H_{n}>\beta, n H_{n}^{d}>\beta\right\} \\
& \leq \mathbb{P}\left\{\begin{array}{l}
\left.\sup _{k:\left(1+\delta_{n}\right)^{k}>\beta, n /\left(1+\delta_{n}\right)^{k d}>\beta} L_{n, k}-L^{*}>\epsilon\right\} \\
\end{array} \quad \leq C_{n} \sup _{k:(1+\delta)^{k}>\beta, n /(1+\delta)^{k d}>\beta} \mathbf{P}\left\{L_{n, k}-L^{*}>\epsilon\right\},\right.
\end{aligned}
$$

by the union bound, where $C_{n}$ is the number of possible values of $H_{n}$ in the given range. As

$$
\beta<\left(1+\delta_{n}\right)^{k}<\left(\frac{n}{\beta}\right)^{1 / d}
$$

we note that

$$
C_{n} \leq 2+\frac{\frac{1}{d} \log \left(\frac{n}{\beta}\right)-\log \beta}{\log \left(1+\delta_{n}\right)}=O\left(\log n / \delta_{n}^{2}\right)=e^{o(n)}
$$

by the condition on the sequence $\left\{\delta_{n}\right\}$. Combining this with Theorem 10.1 , for $n>n_{0}$, we get

$$
\mathbf{P}\left\{L_{n}-L^{*}>\epsilon, 1 / H_{n}>\beta, n H_{n}^{d}>\beta\right\} \leq 4 C_{n} e^{-c n \epsilon^{2}},
$$

which is summable in $n$. The Borel-Cantelli lemma implies that

$$
I_{\left\{1 / H_{n}>\beta, n H_{n}^{d}>\beta, L\left(g_{n}\right)-L^{*}>\epsilon\right\}} \rightarrow 0
$$

with probability one, and the theorem is proved.

For weak consistency, it suffices to require convergence of $H_{n}$ and $n H_{n}^{d}$ in probability (Problem 25.1):

Theorem 25.2. Assume that the random variable $H_{n}$ takes its values from the set of real numbers of the form $\frac{1}{\left(1+\delta_{n}\right)^{k}}$, where $k$ is a nonnegative integer and $\delta_{n}>0$. Let $K$ be a regular kernel. If

$$
\frac{1}{\delta_{n}}=e^{o(n)}
$$

and

$$
H_{n} \rightarrow 0 \text { and } n H_{n}^{d} \rightarrow \infty \quad \text { in probability as } n \rightarrow \infty
$$

then the kernel classification rule corresponding to the random smoothing parameter $H_{n}$ is universally consistent, that is, $L\left(g_{n}\right) \rightarrow L^{*}$ in probability.

We are now prepared to prove a result similar to Theorem 25.1 without restricting the possible values of the random smoothing parameter $H_{n}$. For technical reasons,
we need to assume some additional regularity conditions on the kernel function: $K$ must be decreasing along rays starting from the origin, but it should not decrease too rapidly. Rapidly decreasing functions such as the Gaussian kernel, or functions of bounded support, such as the window kernel are excluded.

Theorem 25.3. Let $K$ be a regular kernel that is monotone decreasing along rays, that is, for any $x \in \mathcal{R}^{d}$ and $a>1, K(a x) \leq K(x)$. Assume in addition that there exists a constant $c>0$ such that for every sufficiently small $\delta>0$, and $x \in \mathcal{R}^{d}$, $K((1+\delta) x) \geq(1-c \delta) K(x)$. Let $\left\{H_{n}\right\}$ be a sequence of random variables satisfying

$$
H_{n} \rightarrow 0 \text { and } n H_{n}^{d} \rightarrow \infty \text { with probability one, as } n \rightarrow \infty
$$

Then the error probability $L\left(g_{n}\right)$ of the kernel classification rule with kernel $K$ and smoothing parameter $H_{n}$ converges to $L^{*}$ with probability one, that is, the rule is strongly universally consistent.

Remark. The technical condition on $K$ is needed to ensure that small changes in $h$ do not cause dramatic changes in $L\left(g_{n}\right)$. We expect some smooth behavior of $L\left(g_{n}\right)$ as a function of $h$. The conditions are rather restrictive, as the kernels must have infinite support and decrease slower than at a polynomial rate. An example satisfying the conditions is

$$
K(x)= \begin{cases}1 & \text { if }\|x\| \leq 1 \\ 1 /\|x\|^{r} & \text { otherwise }\end{cases}
$$

where $r>0$ (see Problem 25.2). The conditions on $H_{n}$ are by no means necessary. We have already seen that consistency occurs for atomic distributions if $K(0)>0$ and $H_{n} \equiv 0$, or for distributions with $L^{*}=1 / 2$ when $H_{n}$ takes any value. However, Theorem 25.3 provides us with a simple collection of sufficient conditions.

Proof of Theorem 25.3. First we discretize $H_{n}$. Define a sequence $\delta_{n} \rightarrow 0$ satisfying the condition in Theorem 25.1, and introduce the random variables $\underline{H}_{n}$ and $\bar{H}_{n}$ as follows: $\underline{H}_{n}=\frac{1}{\left(1+\delta_{n}\right)^{K_{n}}}$, where $K_{n}$ is the smallest integer such that $H_{n}>\frac{1}{\left(1+\delta_{n}\right)^{K_{n}}}$, and let $\bar{H}_{n}=\left(1+\delta_{n}\right) \underline{H}_{n}$. Thus, $\underline{H}_{n}<H_{n} \leq \bar{H}_{n}$. Note that both $\underline{H}_{n}$ and $\bar{H}_{n}$ satisfy the conditions of Theorem 25.1. As usual, the consistency proof is based on Theorem 2.3. Here, however, we need a somewhat tricky choice of the denominator of the functions that approximate $\eta(x)$. Introduce

$$
\widehat{\eta}_{n, H_{n}}(x)=\frac{\frac{1}{n} \sum_{i=1}^{n} I_{\left\{Y_{i}=1\right\}} K\left(\frac{x-X_{i}}{H_{n}}\right)}{\int K\left(\frac{x-z}{\bar{H}_{n}}\right) \mu(d z)} .
$$

Clearly, the value of the classification rule $g_{n}(x)$ equals one if and only if $\widehat{\eta}_{n, H_{n}}(x)$ is greater than the function defined similarly, with the $I_{\left\{Y_{i}=1\right\}}$ 's replaced with $I_{\left\{Y_{i}=0\right\}}$ Then by Theorem 2.3 it suffices to show that

$$
\int\left|\widehat{\eta}_{n, H_{n}}(x)-\eta(x)\right| \mu(d x) \rightarrow 0
$$

with probability one. We use the following decomposition:

$$
\begin{align*}
& \int\left|\widehat{\eta}_{n, H_{n}}(x)-\eta(x)\right| \mu(d x) \\
& \quad \leq \int\left|\widehat{\eta}_{n, H_{n}}(x)-\widehat{\eta}_{n, \bar{H}_{n}}(x)\right| \mu(d x)+\int\left|\widehat{\eta}_{n, \bar{H}_{n}}(x)-\eta(x)\right| \mu(d x) \tag{25.1}
\end{align*}
$$

The second term on the right-hand side converges to zero with probability one, which can be seen by repeating the argument of the proof of Theorem 25.1, using the observation that in the proof of Theorem 10.1 we proved consistency via an exponential probability inequality for

$$
\int\left|\eta_{n, h_{n}}(x)-\eta(x)\right| \mu(d x)
$$

The first term may be bounded as the following simple chain of inequalities indicates:

$$
\begin{aligned}
& \int\left|\widehat{\eta}_{n, H_{n}}(x)-\widehat{\eta}_{n, \bar{H}_{n}}(x)\right| \mu(d x) \\
& \left.\quad=\int \frac{\frac{1}{n} \sum_{i=1}^{n} Y_{i} K\left(\frac{x-X_{i}}{H_{n}}\right)-\frac{1}{n} \sum_{i=1}^{n} Y_{i} K\left(\frac{x-X_{i}}{\bar{H}_{n}}\right)}{\int K\left(\frac{x-z}{\bar{H}_{n}}\right) \mu(d z)} \right\rvert\, \mu(d x) \\
& \quad \leq \int \frac{\frac{1}{n} \sum_{i=1}^{n}\left(K\left(\frac{x-X_{i}}{\bar{H}_{n}}\right)-K\left(\frac{x-X_{i}}{\underline{H}_{n}}\right)\right)}{\int K\left(\frac{x-z}{\bar{H}_{n}}\right) \mu(d z)} \mu(d x)
\end{aligned}
$$

(by the monotonicity of $K$ )

$$
\leq \int \frac{\frac{1}{n} \sum_{i=1}^{n}\left(1-\left(1-c \delta_{n}\right)\right) K\left(\frac{x-X_{i}}{\bar{H}_{n}}\right)}{\int K\left(\frac{x-z}{\bar{H}_{n}}\right) \mu(d z)} \mu(d x)
$$

(from $\bar{H}_{n}=\left(1+\delta_{n}\right) \underline{H}_{n}$, and the condition on $K$, if $n$ is large enough)
$=c \delta_{n} \int \frac{\frac{1}{n} \sum_{i=1}^{n} K\left(\frac{x-X_{i}}{\bar{H}_{n}}\right)}{\int K\left(\frac{x-z}{\bar{H}_{n}}\right) \mu(d z)} \mu(d x)$.
Since $\bar{H}_{n}$ satisfies the conditions of Theorem 25.1, the integral on the right-hand side converges to one with probability one, just as we argued for the second term on the right-hand side of (25.1). But $\delta_{n}$ converges to zero. Therefore, the first term on the right-hand side of (25.1) tends to zero with probability one.

Remark. A quick inspection of the proof above shows that if $\left\{a_{n}\right\}$ and $\left\{b_{n}\right\}$ are deterministic sequences with the property that $a_{n}<b_{n}, b_{n} \rightarrow 0$, and $n a_{n}^{d} \rightarrow \infty$,
then for the kernel estimate with kernel as in Theorem 25.3, we have

$$
\sup _{a_{n} \leq h \leq b_{n}} L_{n}(h) \rightarrow L^{*} \quad \text { with probability one }
$$

for all distributions. One would never use the worst smoothing factor over the range $\left[a_{n}, b_{n}\right]$, but this corollary points out just how powerful Theorem 25.3 is.

### 25.2 Data Splitting

Our first example of a data-dependent $H_{n}$ is based upon the minimization of a suitable error estimate. You should have read Chapter 22 on data splitting if you want to understand the remainder of this section.

The data sequence $D_{n}=\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ is divided into two parts. The first part $D_{m}=\left(X_{1}, Y_{1}\right), \ldots,\left(X_{m}, Y_{m}\right)$ is used for training, while the remaining $l=n-m$ pairs constitute the testing sequence:

$$
T_{l}=\left(X_{m+1}, Y_{m+1}\right), \ldots,\left(X_{m+l}, Y_{m+l}\right)
$$

The training sequence $D_{m}$ is used to design a class of classifiers $\mathcal{C}_{m}$, which, in our case is the class of kernel rules based on $D_{m}$, with all possible values of $h>0$, for fixed kernel $K$. Note that the value of the kernel rule $g_{m}(x)$ with smoothing parameter $h$ is zero if and only if

$$
f_{m}(x)=\sum_{i=1}^{m}\left(Y_{i}-1 / 2\right) K\left(\frac{x-X_{i}}{h}\right) \leq 0 .
$$

Classifiers in $\mathcal{C}_{m}$ are denoted by $\phi_{m}$. A classifier is selected from $\mathcal{C}_{m}$ that minimizes the holdout estimate of the error probability:

$$
\widehat{L}_{m, l}\left(\phi_{m}\right)=\frac{1}{l} \sum_{i=1}^{l} I_{\left\{\phi_{m}\left(X_{m+i}\right) \neq Y_{m+i}\right\}}
$$

The particular rule selected in this manner is called $g_{n}$. The question is how far the error probability $L\left(g_{n}\right)$ of the obtained rule is from that of the optimal rule in $\mathcal{C}_{m}$.

Finite collections. It is computationally attractive to restrict the possible values of $h$ to a finite set of real numbers. For example, $\mathcal{C}_{m}$ could consist of all kernel rules with $h \in\left\{2^{-k_{m}}, 2^{-k_{m}+1}, \ldots, 1 / 2,1,2, \ldots, 2^{k_{m}}\right\}$, for some positive integer $k_{m}$. The advantage of this choice of $\mathcal{C}_{m}$ is that the best $h$ in this class is within a factor of two of the best $h$ among all possible real smoothing factors, unless the best $h$ is smaller than $2^{-k_{m}-1}$ or larger than $2^{k_{m}+1}$. Clearly, $\left|\mathcal{C}_{m}\right|=2 k_{m}+1$, and as pointed out in Chapter 22, for the selected rule $g_{n}$ Hoeffding's inequality and the union bound imply that

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)>\epsilon \mid D_{m}\right\} \leq\left(4 k_{m}+2\right) e^{-l \epsilon^{2} / 2} .
$$

If $k_{m}=e^{o(l)}$, then the upper bound decreases exponentially in $l$, and, in fact,

$$
\mathbb{E}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)\right\}=O\left(\sqrt{\frac{\log (l)}{l}}\right)
$$

By Theorem 10.1, $\mathcal{C}_{m}$ contains a subsequence of consistent rules if $m \rightarrow \infty$, and $k_{m} \rightarrow \infty$ as $n \rightarrow \infty$. To make sure that $g_{n}$ is strongly universally consistent as well, we only need that $\lim _{n \rightarrow \infty} l=\infty$, and $k_{m}=e^{o(l)}$ (see Theorem 22.1). Under these conditions, the rule is $\sqrt{\log \left(k_{m}\right) / l}$-optimal (see Chapter 22).

The discussion above does little to help us with the selection of $m, l$, and $k_{m}$. Safe, but possibly suboptimal, choices might be $l=n / 10, m=n-l, k_{m}=2 \log _{2} n$. Note that the argument above is valid for any regular kernel $K$.

InFinite collections. If we do not want to exclude any value of the smoothing parameter, and pick $h$ from $[0, \infty)$, then $\mathcal{C}_{m}$ is of infinite cardinality. Here, we need something stronger, like the Vapnik-Chervonenkis theory. For example, from Chapter 22, we have

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)>\epsilon \mid D_{m}\right\} \leq 4 e^{8} \mathcal{S}\left(\mathcal{C}_{m}, l^{2}\right) e^{-l \epsilon^{2} / 2}
$$

where $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ is the $l$-th shatter coefficient corresponding to the class of classifiers $\mathcal{C}_{m}$. We now obtain upper bounds for $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ for different choices of $K$.

Define the function

$$
f_{m}(x)=f_{m}\left(x, D_{m}\right)=\sum_{i=1}^{m}\left(Y_{i}-\frac{1}{2}\right) K\left(\frac{x-X_{i}}{h}\right) .
$$

Recall that for the kernel rule based on $D_{m}$,

$$
g_{m}(x)= \begin{cases}0 & \text { if } f_{m}\left(x, D_{m}\right) \leq 0 \\ 1 & \text { otherwise }\end{cases}
$$

We introduce the kernel complexity $\kappa_{m}$ :

$$
\begin{aligned}
\kappa_{m}= & \sup _{\substack{x,\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right)}}\{\text { Number of sign changes of } \\
& \left.f_{m}\left(x,\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right)\right) \text { as } h \text { varies from } 0 \text { to infinity }\right\} .
\end{aligned}
$$

Suppose we have a kernel with kernel complexity $\kappa_{m}$. Then, as $h$ varies from 0 to infinity, the binary $l$-vector

$$
\left(\operatorname{sign}\left(f_{m}\left(X_{m+1}\right)\right), \ldots, \operatorname{sign}\left(f_{m}\left(X_{m+l}\right)\right)\right.
$$

changes at most $l \kappa_{m}$ times. It can thus take at most $l \kappa_{m}+1$ different values. Therefore,

$$
\mathcal{S}\left(\mathcal{C}_{m}, l\right) \leq l \kappa_{m}+1
$$

We postpone the issue of computing kernel complexities until the next section. It suffices to note that if $g_{n}$ is obtained by minimizing the holdout error estimate $\widehat{L}_{m, l}\left(\phi_{m}\right)$ by varying $h$, then

$$
\begin{align*}
& \mathbf{E}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right) \mid D_{m}\right\}  \tag{25.2}\\
& \quad \leq 16 \sqrt{\frac{\log \left(8 e \mathcal{S}\left(\mathcal{C}_{m}, l\right)\right)}{2 l}}(\text { Corollary 12.1) } \\
& \quad \leq 16 \sqrt{\frac{\log \left(8 e\left(l \kappa_{m}+1\right)\right)}{2 l}} \tag{25.3}
\end{align*}
$$

Various probability bounds may also be derived from the results of Chapter 12. For example, we have

$$
\begin{align*}
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)>\epsilon \mid D_{m}\right\} & \leq 4 e^{8} \mathcal{S}\left(\mathcal{C}_{m}, l^{2}\right) e^{-l \epsilon^{2} / 2} \\
& \leq 4 e^{8}\left(l^{2} \kappa_{m}+1\right) e^{-l \epsilon^{2} / 2} \tag{25.4}
\end{align*}
$$

Theorem 25.4. Assume that $g_{n}$ minimizes the holdout estimate $\widehat{L}_{m, l}\left(\phi_{m}\right)$ over all kernel rules with fixed kernel $K$ of kernel complexity $\kappa_{m}$, and over all (unrestricted) smoothing factors $h>0$. Then $g_{n}$ is strongly universally consistent if
(i) $\lim _{n \rightarrow \infty} m=\infty$;
(ii) $\lim _{n \rightarrow \infty} \frac{\log \kappa_{m}}{l}=0$;
(iii) $\lim _{n \rightarrow \infty} \frac{l}{\log n}=\infty$;
(iv) $\quad K$ is a regular kernel.

For weak universal consistency, (iii) may be replaced by (v): $\lim _{n \rightarrow \infty} l=\infty$.

Proof. Note that $\mathcal{C}_{m}$ contains a strongly universally consistent subsequencetake $h=m^{-1 /(2 d)}$ for example, and apply Theorem 10.1, noting that $h \rightarrow 0$, yet $m h^{d} \rightarrow \infty$. Thus,

$$
\lim _{n \rightarrow \infty} \inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)=L^{*} \quad \text { with probability one. }
$$

It suffices to apply Theorem 22.1 and to note that the bound in (25.4) is summable in $n$ when $l / \log n \rightarrow \infty$ and $\log \kappa_{m}=o(l)$. For weak universal consistency, a simple application of (25.2) suffices to note that we only need $l \rightarrow \infty$ instead of $l / \log n \rightarrow \infty$.

Approximation errors decrease with $m$. For example, if class densities exist, we may combine the inequality of Problem 2.10 with bounds from Devroye and Györfi (1985) and Holmström and Klemelä (1992) to conclude that $\mathbf{E}\left\{\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)\right\}$ $L^{*}$ is of the order of $m^{-2 \alpha /(4+d)}$, with $\alpha \in[1,2]$, under suitable conditions on $K$
and the densities. By (25.2), the estimation error is $O\left(\sqrt{\log \left(l \kappa_{m}\right) / l}\right)$, requiring instead large values for $l$. Clearly, some sort of balance is called for. Ignoring the logarithmic term for now, we see that $l$ should be roughly $m^{4 \alpha /(4+d)}$ if we are to balance errors of both kinds. Unfortunately, all of this is ad hoc and based upon unverifiable distributional conditions. Ideally, one should let the data select $l$ and $m$. See Problem 25.3, and Problem 22.4 on optimal data splitting.

For some distributions, the estimation error is just too large to obtain asymptotic optimality as defined in Chapter 22 . For example, the best bound on the estimation error is $O(\sqrt{\log n / n})$, attained when $\kappa_{m}=1, l=n$. If the distribution of $X$ is atomic with finitely many atoms, then the expected approximation error is $O(1 / \sqrt{m})$. Hence the error introduced by the selection process smothers the approximation error when $m$ is linear in $n$. Similar conclusions may even be drawn when $X$ has a density: consider the uniform distribution on $[0,1] \cup[3,4]$ with $\eta(x)=1$ if $x \in[0,1]$ and $\eta(x)=0$ if $x \in[3,4]$. For the kernel rule with $h=1, K=I_{[-1,1]}$, the expected approximation error tends to $L^{*}=0$ exponentially quickly in $m$, and this is always far better than the estimation error which at best is $O(\sqrt{\log n / n})$.

### 25.3 Kernel Complexity

We now turn to the kernel complexity $\kappa_{m}$. The following lemmas are useful in our computations. If $l / \log n \rightarrow \infty$, we note that for strong consistency it suffices that $\kappa_{m}=O\left(m^{\gamma}\right)$ for some finite $\gamma$ (just verify the proof again). This, as it turns out, is satisfied for nearly all practical kernels.

Lemma 25.1. Let $0 \leq b_{1}<\cdots<b_{m}$ be fixed numbers, and let $a_{i} \in \mathcal{R}, 1 \leq i \leq$ $m$, be fixed, with the restriction that $a_{m} \neq 0$. Then the function $f(x)=\sum_{i=1}^{m} a_{i} x^{b_{i}}$ has at most $m-1$ nonzero positive roots.

Proof. Note first that $f$ cannot be identically zero on any interval of nonzero length. Let $Z(g)$ denote the number of nonzero positive roots of a function $g$. We have

$$
\begin{aligned}
& Z\left(\sum_{i=1}^{m} a_{i} x^{b_{i}}\right) \\
& \quad=Z\left(\sum_{i=1}^{m} a_{i} x^{c_{i}}\right) \quad\left(\text { where } c_{i}=b_{i}-b_{1}, \text { all } i \text {; thus, } c_{1}=0\right)
\end{aligned}
$$

$$
\leq Z\left(\sum_{i=2}^{m} a_{i} c_{i} x^{c_{i}-1}\right)+1
$$

(for a continuously differentiable $g$, we have $Z(g) \leq 1+Z\left(g^{\prime}\right)$ )

$$
=Z\left(\sum_{i=2}^{m} a_{i}^{\prime} x^{b_{i}^{\prime}}\right)+1
$$

(where $a_{i}^{\prime}=a_{i} c_{i}, b_{i}^{\prime}=c_{i}-c_{2}$, all $i \geq 2$ ).
Note that the $b_{i}^{\prime}$ are increasing, $b_{2}^{\prime}=0$, and $a_{i}^{\prime} \neq 0$. As $Z\left(a_{m} x^{b_{m}}\right)=0$ for $a_{m} \neq 0$, we derive our claim by simple induction on $m$.

Lemma 25.2. Let $a_{1}, \ldots, a_{m}$ be fixed real numbers, and let $b_{1}, \ldots, b_{m}$ be different nonnegative reals. Then if $\alpha \neq 0$, the function

$$
f(x)=\sum_{i=1}^{m} a_{i} e^{-b_{i} x^{\alpha}}, \quad x \geq 0
$$

is either identically zero, or takes the value 0 at most m times.

Proof. Define $y=e^{-x^{\alpha}}$. If $\alpha \neq 0, y$ ranges from 0 to 1 . By Lemma 25.1, $g(y)=$ $\sum_{i=1}^{m} a_{i} y^{b_{i}}$ takes the value 0 at at most $m-1$ positive $y$-values, unless it is identically zero everywhere. This concludes the proof of the lemma.

A star-shaped kernel is one of the form $K(x)=I_{\{x \in A\}}$, where $A$ is a star-shaped set of unit Lebesgue measure, that is, $x \notin A$ implies $c x \notin A$ for all $c \geq 1$. It is clear that $\kappa_{m}=m-1$ by a simple thresholding argument. On the real line, the kernel $K(x)=\sum_{i=-\infty}^{\infty} a_{i} I_{\{x \in[2 i, 2 i+1]\}}$ for $a_{i}>0$ oscillates infinitely often, and has $\kappa_{m}=\infty$ for all values of $m \geq 2$. We must therefore disallow such kernels. For the same reason, kernels such as $K(x)=(\sin x / x)^{r}$ on the real line are not good (see Problem 25.4).

If $K=\sum_{i=1}^{k} a_{i} I_{A_{i}}$ for some finite $k$, some numbers $a_{i}$ and some star-shaped sets $A_{i}$, then $\kappa_{m} \leq k(m-1)$.

Consider next kernels of the form

$$
K(x)=\|x\|^{-r} I_{\{x \in A\}},
$$

where $A$ is star-shaped, and $r \geq 0$ is a constant (see Sebestyen (1962)). We see that

$$
\begin{aligned}
f_{m}(x) & =\sum_{i=1}^{m}\left(Y_{i}-\frac{1}{2}\right) K\left(\frac{x-X_{i}}{h}\right) \\
& =h^{r} \sum_{i=1}^{m}\left(Y_{i}-\frac{1}{2}\right)\left\|x-X_{i}\right\|^{-r} I_{\left\{\left(x-X_{i}\right) / h \in A\right\}},
\end{aligned}
$$

which changes sign at most as often as $f_{m}(x) / h^{r}$. From our earlier remarks, it is easy to see that $\kappa_{m}=m-1$, as $\kappa_{m}$ is the same as for the kernel $K=I_{A}$. If $A$ is replaced by $\mathcal{R}^{d}$, then the kernel is not integrable, but clearly, $\kappa_{m}=0$. Assume next that we have

$$
K(x)= \begin{cases}1 & \text { if }\|x\| \leq 1 \\ 1 /\|x\|^{r} & \text { otherwise }\end{cases}
$$

where $r>0$. For $r>d$, these kernels are integrable and thus regular. Note that

$$
K\left(\frac{x}{h}\right)=I_{\{\|x\| \leq h\}}+I_{\{\|x\|>h\}} \frac{h^{r}}{\|x\|^{r}}
$$

As $h$ increases, $f_{m}(x)$, which is of the form

$$
\sum_{i:\left\|x-X_{i}\right\| \leq h}\left(Y_{i}-\frac{1}{2}\right)+h^{r} \sum_{i:\left\|x-X_{i}\right\|>h}\left(Y_{i}-\frac{1}{2}\right) \frac{1}{\left\|x-X_{i}\right\|^{r}}
$$

transfers an $X_{i}$ from one sum to the other at most $m$ times. On an interval on which no such transfer occurs, $f_{m}$ varies as $\alpha+\beta h^{r}$ and has at most one sign change. Therefore, $\kappa_{m}$ cannot be more than $m+1$ (one for each $h$-interval) plus $m$ (one for each transfer), so that $\kappa_{m} \leq 2 m+1$. For more practice with such computations, we refer to the exercises. We now continue with a few important classes of kernels.

Consider next exponential kernels such as

$$
K(x)=e^{-\|x\|^{\alpha}}
$$

for some $\alpha>0$, where $\|\cdot\|$ is any norm on $\mathcal{R}^{d}$. These kernels include the popular gaussian kernels. As the decision rule based on $D_{m}$ is of the form

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{m}\left(2 Y_{i}-1\right) e^{-\frac{\left\|x-X_{i}\right\|^{\alpha}}{h^{\alpha}}}>0 \\ 0 & \text { otherwise }\end{cases}
$$

a simple application of Lemma 25.2 shows that $\kappa_{m} \leq m$. The entire class of kernels behaves nicely.

Among compact support kernels, kernels of the form

$$
K(x)=\left(\sum_{i=1}^{k} a_{i}\|x\|^{b_{i}}\right) I_{\{\|x\| \leq 1\}}
$$

for real numbers $a_{i}$, and $b_{i} \geq 1$, are important. A particularly popular kernel in $d$-dimensional density estimation is Deheuvels' (1977) kernel

$$
K(x)=\left(1-\|x\|^{2}\right) I_{\{\|x\| \leq 1\}}
$$

If the kernel was $K(x)=\left(\sum_{i=1}^{k} a_{i}\|x\|^{b_{i}}\right)$, without the indicator function, then Lemma 25.1 would immediately yield the estimate $\kappa_{m} \leq k$, uniformly over all $m$. Such kernels would be particularly interesting. With the indicator function
multiplied in, we have $\kappa_{m} \leq k m$, simply because $f_{m}(x)$ at each $h$ is based upen a subset of the $X_{j}$ 's, $1 \leq j \leq m$, with the subset growing monotonically with $h$. For each subset, the function $f_{m}(x)$ is a polynomial in $\|x\|$ with powers $b_{1}, \ldots, b_{k}$, and changes sign at most $k$ times. Therefore, polynomial kernels of compact support also have small complexities. Observe that the " $k$ " in the bound $\kappa_{m} \leq k m$ refers to the number of terms in the polynomial and not the maximal power.

A large class of kernels of finite complexity may be obtained by applying the rich theory of total positivity. See Karlin (1968) for a thorough treatment. A realvalued function $L$ of two real variables is said to be totally positive on $A \times B \subset \mathcal{R}^{2}$ if for all $n$, and all $s_{1}<\cdots<s_{n}, s_{i} \in A, t_{1}<\cdots<t_{n}, t_{i} \in B$, the determinant of the matrix with elements $L\left(s_{i}, t_{j}\right)$ is nonnegative. A key property of such functions is the following result, which we cite without proof:

Theorem 25.5. (Schoenberg (1950)). Let L be a totally positive function on $A \times B \in \mathcal{R}^{2}$, and let $\alpha: B \rightarrow \mathcal{R}$ be a bounded function. Define the function

$$
\beta(s)=\int_{B} L(s, t) \alpha(t) \sigma(d t)
$$

on $A$, where $\sigma$ is a finite measure on $B$. Then $\beta(s)$ changes sign at most as many times as $\alpha(s)$ does. (The number of sign changes of a function $\beta$ is defined as the supremum of sign changes of sequences of the form $\beta\left(s_{1}\right), \ldots, \beta\left(s_{n}\right)$, where $n$ is arbitrary, and $s_{1}<\cdots<s_{n}$.)

Corollary 25.1. Assume that the kernel $K$ is such that the function $L(s, t)=$ $K(s t)$ is totally positive for $s>0$ and $t \in \mathcal{R}^{d}$. Then the kernel complexity of $K$ satisfies $\kappa_{m} \leq m-1$.

Proof. We apply Theorem 25.5. We are interested in the number of sign changes of the function

$$
\beta(s)=\sum_{i=1}^{m}\left(2 Y_{i}-1\right) K\left(\left(X_{i}-x\right) s\right)
$$

on $s \in(0, \infty)$ ( $s$ plays the role of $1 / h$ ). But $\beta(s)$ may be written as

$$
\beta(s)=\int_{B} L(s, t) \alpha(t) \sigma(d t)
$$

where $L(s, t)=K(s t)$, the measure $\sigma$ puts mass 1 on each $t=X_{i}-x, i=1, \ldots, m$, and $\alpha(t)$ is defined at these points as

$$
\alpha(t)=2 Y_{i}-1 \text { if } t=X_{i}-x
$$

Other values of $\alpha(t)$ are irrelevant for the integral above. Clearly, $\alpha(t)$ can be defined such that it changes sign at most $m-1$ times. Then Theorem 25.5 implies that $\beta(s)$ changes sign at most $m-1$ times, as desired.

This corollary equips us with a whole army of kernels with small complexity. For examples, refer to the monograph of Karlin (1968).

### 25.4 Multiparameter Kernel Rules

Assume that in the kernel rules considered in $\mathcal{C}_{m}$, we perform an optimization with respect to more than one parameter. Collect these parameters in $\theta$, and write the discrimination function as

$$
f_{m}(x)=\sum_{i=1}^{m}\left(Y_{i}-\frac{1}{2}\right) K_{\theta}\left(x-X_{i}\right)
$$

## EXAMPLES:

(i) Product kernels. Take

$$
K_{\theta}(x)=\prod_{j=1}^{d} K\left(\frac{x}{h^{(j)}}\right)
$$

where $\theta=\left(h^{(1)}, \ldots, h^{(d)}\right)$ is a vector of smoothing factors-one per dimen-sion-and $K$ is a fixed one-dimensional kernel.
(ii) Kernels of variable form. Define

$$
K_{\theta}(x)=e^{-\|x\|^{\alpha} / h^{\alpha}}
$$

where $\alpha>0$ is a shape parameter, and $h>0$ is the standard smoothing parameter. Here $\theta=(\alpha, h)$ is two-dimensional.
(iii) Define

$$
K_{\theta}(x)=\frac{1}{1+x^{T} R R^{T} x}=\frac{1}{1+\|R x\|^{2}}
$$

where $x^{T}$ is the transpose of $x$, and $R$ is an orthogonal transformation matrix, all of whose free components taken together are collected in $\theta$. Kernels of this kind may be used to adjust automatically to a certain variance-covariance structure in the data.
We will not spend a lot of time on these cases. Clearly, one route is to properly generalize the definition of kernel complexity. In some cases, it is more convenient to directly find upper bounds for $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$. In the product kernel case, with onedimensional kernel $I_{[-1,1]}$, we claim for example that

$$
\mathcal{S}\left(\mathcal{C}_{m}, l\right) \leq(l m)^{d}+1
$$

The corresponding rule takes a majority vote over centered rectangles with sides equal to $2 h^{(1)}, 2 h^{(2)}, \ldots, 2 h^{(d)}$. To see why the inequality is true, consider the $d$ dimensional quadrant of $l m$ points obtained by taking the absolute values of the vectors $X_{j}-X_{i}, m<j \leq m+l=n, 1 \leq i \leq m$, where the absolute value of a vector is a vector whose components are the absolute values of the components of the vector. To compute $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$, it suffices to count how many different subsets can be obtained from these $l m$ points by considering all possible rectangles with one vertex at the origin, and the diagonally opposite vertex in the quadrant. This is $1+(l m)^{d}$. The strong universal consistency of the latter family $\mathcal{C}_{m}$ is insured when $m \rightarrow \infty$ and $l / \log n \rightarrow \infty$.

### 25.5 Kernels of Infinite Complexity

In this section we demonstrate that not every kernel function supports smoothing factor selection based on data splitting. These kernels have infinite kernel complexity, or even worse, infinite VC dimension. Some of the examples may appear rather artificial, but some "nice" kernels will surprise us by misbehaving.

We begin with a kernel $K$ having the property that the class of sets

$$
\mathcal{A}=\left\{\left\{x: K\left(\frac{x-X_{1}}{h}\right)>0\right\}: h>0\right\}
$$

has vc dimension $V_{\mathcal{A}}=\infty$ for any fixed value of $X_{1}$ (i.e., using the notation of the previous sections, $V_{\mathcal{C}_{1}}=\infty$ ). Hence, with one sample point, all hope is lost to use the Vapnik-Chervonenkis inequality in any meaningful way. Unfortunately, the kernel $K$ takes alternately positive and negative values. In the second part of this section, a kernel is constructed that is unimodal and symmetric and has $V_{\mathcal{C}_{m}}=\infty$ for $m=4$ when $D_{4}=\left(\left(X_{1}, Y_{1}\right),\left(X_{2}, Y_{2}\right),\left(X_{3}, Y_{3}\right),\left(X_{4}, Y_{4}\right)\right)$ takes certain values. Finally, in the last part, we construct a positive kernel with the property that for any $m$ and any nondegenerate nonatomic distribution, $\lim _{m \rightarrow \infty} \mathbf{P}\left\{V_{\mathcal{C}_{m}}=\infty\right\}=1$.

We return to $\mathcal{A}$. Our function is picked as follows:

$$
K(x)=\alpha(x) g(i), \quad 2^{i} \leq x<2^{i+1}, \quad i \geq 1,
$$

where $g(i) \in\{-1,1\}$ for all $i, \alpha(x)>0$ for all $x, \alpha(x) \downarrow 0$ as $x \uparrow \infty$, $x>0$, and $\alpha(x) \downarrow 0$ as $x \downarrow-\infty, x<0$ (the monotonicity is not essential and may be dropped but the resulting class of kernels will be even less interesting). We enumerate all binary strings in lexicographical order, replace all 0's by -1 's, and map the bit sequence to $g(1), g(2), \ldots$. Hence, the bit sequence $0,1,00,01,10,11,000,001,010, \ldots$ becomes

$$
-1,1,-1,-1,-1,1,1,-1,1,1,-1,-1,-1,-1,-1,1,-1,1,-1, \ldots
$$

Call this sequence $S$. For every $n$, we can find a set $\left\{x_{1}, \ldots, x_{n}\right\}$ that can be shattered by sets from $\mathcal{A}$. Take $\left\{x_{1}, \ldots, x_{n}\right\}=X_{1}+\left\{2^{1}, \ldots, 2^{n}\right\}$. A subset of this may be characterized by a string $s_{n}$ from $\{-1,1\}^{n}, 1$ 's denoting membership, and -1 's denoting absence. We find the first occurrence of $s_{n}$ in $S$ and let the starting point be $g(k)$. Take $h=2^{1-k}$. Observe that

$$
\begin{aligned}
& \left(K\left(\frac{x_{1}-X_{1}}{h}\right), \ldots, K\left(\frac{x_{n}-X_{1}}{h}\right)\right) \\
& \quad=\left(K\left(2^{k}\right), \ldots, K\left(2^{k+n-1}\right)\right) \\
& \quad=\left(\alpha\left(2^{k}\right) g(k), \alpha\left(2^{k+1}\right) g(k+1), \ldots, \alpha\left(2^{k+n-1}\right) g(k+n-1)\right),
\end{aligned}
$$

which agrees in sign with $s_{n}$ as desired. Hence, the vC dimension is infinite.
The next kernel is symmetric, unimodal, and piecewise quadratic. The intervals into which $[0, \infty)$ is divided are denoted by $A_{0}, A_{1}, A_{2}, \ldots$, from left to right,
where

$$
A_{i}= \begin{cases}{\left[\frac{3}{4} 2^{i}, \frac{3}{2} 2^{i}\right)} & i \geq 1 \\ {\left[0, \frac{3}{2}\right)^{2}} & i=0\end{cases}
$$

On each $A_{i}, K$ is of the form $a x^{2}+b x+c$. Observe that $K^{\prime \prime}=2 a$ takes the sign of $a$. Also, any finite symmetric difference of order 2 has the sign of $a$, as

$$
K(x+\delta)+K(x-\delta)-2 K(x)=2 a \delta^{2}, x-\delta, x, x+\delta \in A_{i}
$$

We take four points and construct the class

$$
\mathcal{A}_{4}=\left\{\left\{x: \sum_{i=1}^{4} K\left(\frac{x-X_{i}}{h}\right)\left(2 Y_{i}-1\right)>0\right\}: h>0\right\},
$$

where $X_{1}=-\delta, X_{2}=\delta, X_{3}=X_{4}=0, Y_{1}=Y_{2}=0, Y_{3}=Y_{4}=1$ are fixed for now. On $A_{i}$, we let the " $a$ " coefficient have the same sign as $g(i)$, known from the previous example. All three quadratic coefficients are picked so that $K \geq 0$ and $K$ is unimodal. For each $n$, we show that the set $\left\{2^{1}, \ldots, 2^{n}\right\}$ can be shattered by intersecting with sets from $\mathcal{A}_{4}$. A subset is again identified by a $\{-1,1\}^{n}$-valued string $s_{n}$, and its first match in $S$ is $g(k), \ldots, g(k+n-1)$. We take $h=2^{1-k}$. Note that for $1 \leq i \leq n$,

$$
\begin{aligned}
\operatorname{sign} & \left(K\left(\frac{2^{i}-\delta}{h}\right)+K\left(\frac{2^{i}+\delta}{h}\right)-2 K\left(\frac{2^{i}}{h}\right)\right) \\
= & \operatorname{sign}\left(K\left(2^{k+i-1}-\delta 2^{k-1}\right)+K\left(2^{k+i-1}+\delta 2^{k-1}\right)-2 K\left(2^{k+i-1}\right)\right) \\
= & \operatorname{sign}\left(K^{\prime \prime}\left(2^{k+i-1}\right)\right) \\
& (\text { if } \delta \leq 1 / 4, \text { by the finite-difference property of quadratics }) \\
= & g(k+i-1)
\end{aligned}
$$

as desired. Hence, any subset of $\left\{2^{1}, \ldots, 2^{n}\right\}$ can be picked out.
The previous example works whenever $\delta \leq 1 / 4$. It takes just a little thought to see that if $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{4}, Y_{4}\right)$ are i.i.d. and drawn from the distribution of $(X, Y)$, then

$$
\mathbf{P}\left\{V_{\mathcal{A}_{4}}=\infty\right\} \geq \mathbf{P}\left\{X_{1}=-X_{2}, X_{3}=X_{4}=0, Y_{1}=Y_{2}=0, Y_{3}=Y_{4}=1\right\}
$$

and this is positive if given $Y=0, X$ has atoms at $\delta$ and $-\delta$ for some $\delta>0$; and given $Y=1, X$ has an atom at 0 . However, with some work, we may even remove these restrictions (see Problem 25.12).

We also draw the reader's attention to an 8-point example in which $K$ is symmetric, unimodal, and convex on $[0, \infty)$, yet $V_{\mathcal{C}_{8}}=\infty$. (See Problem 25.11.)

Next we turn to our general $m$-point example. Let the class of rules be given by

$$
g_{m, h}(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{m}\left(2 Y_{i}-1\right) K\left(\frac{x-X_{i}}{h}\right)>0 \\ 0 & \text { otherwise }\end{cases}
$$

$h>0$, where the data $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{m}, Y_{m}\right)$ are fixed for now. We exhibit a nonnegative kernel such that if the $X_{i}$ 's are different and not all the $Y_{i}$ 's are the same, $V_{\mathcal{C}_{m}}=\infty$. This situation occurs with probability tending to one whenever $X$ is nonatomic and $\mathbf{P}\{Y=1\} \in(0,1)$. It is stressed here that the same kernel $K$ is used regardless of the data.

The kernel is of the form $K(x)=K_{0}(x) I_{A}(x)$, where $K_{0}(x)=e^{-x^{2}}$, and $A \subset \mathcal{R}$ will be specially picked. Order $X_{1}, \ldots, X_{m}$ into $X_{(1)}<\cdots<X_{(m)}$ and find the first pair $X_{(i)}, X_{(i+1)}$ with opposite values for $Y_{(i)}$. Without loss of generality, assume that $X_{(i)}=-1, X_{(i+1)}=1,2 Y_{(i)}-1=-1$, and $2 Y_{(i+1)}-1=1\left(Y_{(i)}\right.$ is the $Y$-value that corresponds to $\left.X_{(i)}\right)$. Let the smallest value of $\left|X_{(j)}\right|, j \neq i, j \neq i+1$ be denoted by $\delta>1$. The contribution of all such $X_{(j)}$ 's for $x \in[-\delta / 3, \delta / 3]$ is not more than $(m-2) K_{0}((1+2 \delta / 3) / h) \leq m K_{0}((1+2 \delta / 3) / h)$. The contribution of either $X_{(i)}$ or $X_{(i+1)}$ is at least $K_{0}((1+\delta / 3) / h)$. For fixed $\delta$ and $m$ (after all, they are given), we first find $h^{*}$ such that $h \leq h^{*}$ implies

$$
K_{0}\left(\frac{1+\delta / 3}{h}\right)>m K_{0}\left(\frac{1+2 \delta / 3}{h}\right) .
$$

For $h \leq h^{*}$, the rule, for $x \in[-\delta / 3, \delta / 3]$, is equivalent to a 2 -point rule based on

$$
g_{m, h}(x)= \begin{cases}1 & \text { if }-K\left(\frac{x+1}{h}\right)+K\left(\frac{x-1}{h}\right)>0 \\ 0 & \text { otherwise. }\end{cases}
$$

We define the set $A$ by

$$
A=\bigcup_{k=1}^{\infty} \bigcup_{l=0}^{2^{k}-1} 3^{2^{k}+l} A_{k, l},
$$

where the sets $A_{k, l}$ are specified later. (For $c \neq 0$ and $B \subset \mathcal{R}, c B=\{x \in \mathcal{R}$ : $x / c \in B\}$.) Here $k$ represents the length of a bit string, and $l$ cycles through all $2^{k}$ bit strings of length $k$. For a particular such bit string, $b_{1}, \ldots, b_{k}$, represented by $l$, we define $A_{k, l}$ as follows. First of all, $A_{k, l} \subseteq[1 / 2,3 / 2]$, so that all sets $3^{2^{k}+l} A_{k, l}$ are nonoverlapping. $A_{k, l}$ consists of the sets

$$
\begin{aligned}
& \left(\bigcup_{i: b_{i}=0}\left[1-\frac{1}{(i+1) 2^{k}}, 1-\frac{1}{(i+2) 2^{k}}\right)\right) \\
& \quad \cup\left(\bigcup_{i: b_{i}=1}\left(1+\frac{1}{(i+2) 2^{k}}, 1+\frac{1}{(i+1) 2^{k}}\right]\right) .
\end{aligned}
$$

$A_{k, l}$ is completed by symmetry.
We now exhibit for each integer $k$ a set of that size that can be shattered. These sets will be positioned in $[-\delta / 3,0]$ at coordinate values $-1 /\left(2 \cdot 2^{\rho}\right),-1 /(3$. $\left.2^{\rho}\right), \ldots,-1 /\left((k+1) \cdot 2^{\rho}\right)$, where $\rho$ is a suitable large integer such that $1 /\left(2^{\rho+1}\right)<$ $\delta / 3$. Assume that we wish to extract the set indexed by the bit vector $\left(b_{1}, \ldots, b_{k}\right)$ ( $b_{i}=1$ means that $-1 /(i+1) 2^{\rho}$ must be extracted). To do so, we are only allowed to vary $h$. First, we find a pair $\left(k^{\prime}, l\right)$, where $k^{\prime}$ is at least $k$ and $1 /\left(2^{k^{\prime}+1}\right)<\delta / 3$.
$3^{-2^{k^{\prime}}} \leq h^{*}$ is needed too. Also, $l$ is such that it matches $b_{1}, \ldots, b_{k}$ in its first $k \leq k^{\prime}$ bits. Take $h=3^{-2^{k^{\prime}}-l}$ and $\rho=k^{\prime}$ and observe that

$$
\begin{aligned}
& g_{n, h}\left(-\frac{1}{(i+1) 2^{k^{\prime}}}\right) \\
&= \begin{cases}1 & \text { if } K\left(\frac{-\frac{1}{(i+1))^{k^{\prime}}}}{3^{2^{k^{\prime}+1}}}\right)>K\left(\frac{-\frac{1}{(i+1) 2^{k^{\prime}}}+1}{3^{2^{k^{\prime}}+1}}\right) \\
0 & \text { otherwise }\end{cases} \\
&= \begin{cases}1 & \text { if } 1+\frac{1}{(i+1) 2^{k^{\prime}}} \in A_{k, l}, 1-\frac{1}{(i+1) 2^{k^{\prime}}} \notin A_{k, l} \\
0 & \text { if } 1+\frac{1}{(i+1) 2^{k^{\prime}}} \notin A_{k, l}, 1-\frac{1}{(i+1) 2^{k^{\prime}}} \in A_{k, l} \\
\text { unimportant } & \text { otherwise }\end{cases} \\
& \quad= \begin{cases}1 & \text { if } b_{i}=1 \\
0 & \text { if } b_{i}=0,\end{cases}
\end{aligned}
$$

$1 \leq i \leq k$. Thus, we pick the desired set. This construction may be repeated for all values of $l$ of course. We have shown the following:

Theorem 25.6. If $X$ is nonatomic, $\mathbf{P}\{Y=1\} \in(0,1)$, and if $V_{\mathcal{C}_{m}}$ denotes the vc dimension of the class of kernel rules based on $m$ i.i.d. data drawn from the distribution of $(X, Y)$, and with the kernel specified above, then

$$
\lim _{m \rightarrow \infty} \mathbf{P}\left\{V_{\mathcal{C}_{m}}=\infty\right\}=1
$$

### 25.6 On Minimizing the Apparent Error Rate

In this section, we look more closely at kernel rules that are picked by minimizing the resubstitution estimate $L_{n}^{(R)}$ over kernel rules with smoothing factor $h>0$. We make two remarks in this respect:
(i) The procedure is generally inconsistent if $X$ is nonatomic.
(ii) The method is consistent if $X$ is purely atomic.

To see (i), take $K=I_{\left\{S_{0,1}\right\}}$. We note that $L_{n}^{(R)}=0$ if $h<\min _{i \neq j}\left\|X_{i}-X_{j}\right\|$ as $g_{n}\left(X_{i}\right)=Y_{i}$ for such $h$. In fact, if $H_{n}$ is the minimizing $h$, it may take any value on

$$
\left[0, \min _{i, j: Y_{i}=1, Y_{j}=0}\left\|X_{i}-X_{j}\right\|\right] .
$$

If $X$ is independent of $Y, \mathbf{P}\{Y=1\}=2 / 3$, and $X$ has a density, then,

$$
\lim _{c \rightarrow \infty} \liminf _{n \rightarrow \infty} \mathbf{P}\left\{\min _{i, j: Y_{i}=1, Y_{j}=0}\left\|X_{i}-X_{j}\right\|^{d} \leq \frac{c}{n^{2}}\right\}=1
$$

(Problem 25.13). Thus, $n H^{d} \rightarrow 0$ in probability, and therefore, in this case, $\mathbf{E}\left\{L\left(g_{n}\right)\right\} \rightarrow 2 / 3$ (since ties are broken by favoring the " 0 " class). Hence the inconsistency.

Consider next $X$ purely atomic. Fix $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, the data; and consider the class $\mathcal{C}_{n}$ of kernel rules in which $h>0$ is a free parameter. If a typical rule is $g_{n, h}$, let $\mathcal{A}_{n}$ be the class of sets $\left\{x: g_{n, h}(x)=1\right\}, h>0$. If $g_{n}$ is the rule that minimizes $L_{n}^{(R)}\left(g_{n, h}\right)$, we have

$$
\begin{aligned}
& L\left(g_{n}\right)-\inf _{g_{n, h} \in \mathcal{C}_{n}} L\left(g_{n, h}\right) \\
& \leq \\
& \leq \\
& \leq \sup _{g_{n, h} \in \mathcal{C}_{n}}\left|L_{n}^{(R)}\left(g_{n, h}\right)-L\left(g_{n, h}\right)\right| \quad \text { (Theorem 8.4) } \\
& \leq \\
& \quad \sup _{A \in \mathcal{A}_{n}}\left|\mu_{n}(A)-\mu(A)\right|+2 \sup _{A \in \mathcal{A}_{n}^{c}}\left|\mu_{n}(A)-\mu(A)\right| \\
& \\
& \quad\left(\mathcal{A}_{n}^{c} \text { denotes the collection of sets }\left\{x: g_{n, h}(x)=0\right\}\right) \\
& \leq
\end{aligned} \sup _{B \in \mathcal{B}}\left|\mu_{n}(B)-\mu(B)\right|, \quad .
$$

where $\mathcal{B}$ is the collection of all Borel sets. However, the latter quantity tends to zero with probability one-as denoting the set of the atoms of $X$ by $T$, we have

$$
\begin{aligned}
& \sup _{B \in \mathcal{B}}\left|\mu_{n}(B)-\mu(B)\right| \\
& \quad=\frac{1}{2} \sum_{x \in T}\left|\mu_{n}(\{x\})-\mu(\{x\})\right| \\
& \quad \leq \frac{1}{2} \sum_{x \in A}\left|\mu_{n}(\{x\})-\mu(\{x\})\right|+\mu_{n}\left(A^{c}\right)+\mu\left(A^{c}\right)
\end{aligned}
$$

(where $A$ is an arbitrary finite subset of $T$ )
$\leq 2 \mu\left(A^{c}\right)+\frac{1}{2} \sum_{x \in A}\left|\mu_{n}(\{x\})-\mu(\{x\})\right|+\left|\mu_{n}\left(A^{c}\right)-\mu\left(A^{c}\right)\right|$.
The first term is small by choice of $A$ and the last two terms are small by applying Hoeffding's inequality to each of the $|A|+1$ terms. For yet a different proof, the reader is referred to Problems 25.16 and 25.17.

Note next that

$$
L\left(g_{n}\right)-L^{*} \leq\left(L\left(g_{n}\right)-\inf _{g_{n, h} \in \mathcal{C}_{n}} L\left(g_{n, h}\right)\right)+\left(\inf _{g_{n, h} \in \mathcal{C}_{n}} L\left(g_{n, h}\right)-L^{*}\right)
$$

But if $K(0)>0$ and $K(x) \rightarrow 0$ as $\|x\| \rightarrow \infty$ along any ray, then we have $\inf _{g_{n, h} \in \mathcal{C}_{n}} L\left(g_{n, h}\right) \leq L\left(g_{n}^{\prime}\right)$, where $g_{n}^{\prime}$ is the fundamental rule discussed in Chapter 27 (in which $h=0$ ). Theorem 27.1 shows that $L\left(g_{n}^{\prime}\right) \rightarrow L^{*}$ with probability one, and therefore, $L\left(g_{n}\right) \rightarrow 0$ with probability one, proving

Theorem 25.7. Take any kernel $K$ with $K(0)>0$ and $K(x) \rightarrow 0$ as $\|x\| \rightarrow \infty$ along any ray, and let $g_{n}$ be selected by minimizing $L_{n}^{(R)}$ over all $h \geq 0$. Then $L\left(g_{n}\right) \rightarrow L^{*}$ almost surely whenever the distribution of $X$ is purely atomic.

REMARK. The above theorem is all the more surprising since we may take just about any kernel, such as $K(x)=e^{-\|x\|^{2}}, K(x)=\sin (\|x\|) /\|x\|$, or $K(x)=(1+\|x\|)^{-1 / 3}$. Also, if $X$ puts its mass on a dense subset of $\mathcal{R}^{d}$, the data will look and feel like data from an absolutely continuous distribution (well, very roughly speaking); yet there is a dramatic difference with rules that minimize $L_{n}^{(R)}$ when the $X_{i}$ 's are indeed drawn from a distribution with a density! $\square$

### 25.7 Minimizing the Deleted Estimate

We have seen in Chapter 24 that the deleted estimate of the error probability of a kernel rule is generally very reliable. This suggests using the estimate as a basis of selecting a smoothing factor for the kernel rule. Let $L_{n}^{(D)}\left(g_{n, h}\right)$ be the deleted estimate of $L\left(g_{n, h}\right)$, obtained by using in $g_{n-1, h}$ the same kernel $K$ and smoothing factor $h$. Define the set of $h$ 's for which

$$
L_{n}^{(D)}\left(g_{n, h}\right)=\inf _{h^{\prime} \in[0, \infty)} L_{n}^{(D)}\left(g_{n, h^{\prime}}\right)
$$

by $A$. Set

$$
H_{n}=\inf \{h: h \in A\}
$$

Two fundamental questions regarding $H_{n}$ must be asked:
(a) If we use $H_{n}$ in the kernel estimate, is the rule universally consistent? Note in this respect that Theorem 25.3 cannot be used because it is not true that $H_{n} \rightarrow 0$ in probability in all cases.
(b) If $H_{n}$ is used as smoothing factor, how does $\mathbb{E}\left\{L\left(g_{n}\right)\right\}-L^{*}$ compare to $\mathbf{E}\left\{\inf _{h} L\left(g_{n, h}\right)\right\}-L^{*} ?$

To our knowledge, both questions have been unanswered so far. We believe that this way of selecting $h$ is very effective. Below, generalizing a result by Tutz (1986) (Theorem 27.6), we show that the kernel rule obtained by minimizing the deleted estimate is consistent when the distribution of $X$ is atomic with finitely many atoms.

REMARK. If $\eta \equiv 1$ everywhere, so that $Y_{i} \equiv 1$ for all $i$, and if $K$ has support equal to the unit ball in $\mathcal{R}^{d}$, and $K>0$ on this ball, then $L_{n}^{(D)}\left(g_{n, h}\right)$ is minimal and zero if

$$
h \geq \max _{1 \leq j \leq n}\left\|X_{j}-X_{j}^{N N}\right\|
$$

where $X_{j}^{N N}$ denotes the nearest neighbor of $X_{j}$ among the data points $X_{1}, \ldots$, $X_{j-1}, X_{j+1}, \ldots, X_{n}$. But this shows that

$$
H_{n}=\max _{1 \leq j \leq n}\left\|X_{j}-X_{j}^{N N}\right\|
$$

We leave it as an exercise to show that for any nonatomic distribution of $X, n H_{n}^{d} \rightarrow$ $\infty$ with probability one. However, it is also true that for some distributions, $H_{n} \rightarrow$ $\infty$ with probability one (Problem 25.19). Nevertheless, for $\eta \equiv 1$, the rule is consistent!

If $\eta \equiv 0$ everywhere, then $Y_{i}=0$ for all $i$. Under the same condition on $K$ as above, with tie breaking in favor of class " 0 " (as in the entire book), it is clear that $H_{n} \equiv 0$. Interestingly, here too the rule is consistent despite the strange value of $H_{n} . \square$

We state the following theorem for window kernels, though it can be extended to include kernels taking finitely many different values (see Problem 25.18):

Theorem 25.8. Let $g_{n}$ be the kernel rule with kernel $K=I_{A}$ for some bounded set $A \subset \mathcal{R}^{d}$ containing the origin, and with smoothing factor chosen by minimizing the deleted estimate as described above. Then $\mathbf{E}\left\{L\left(g_{n}\right)\right\} \rightarrow L^{*}$ whenever $X$ has a discrete distribution with finitely many atoms.

Proof. Let $U_{1}, \ldots, U_{n}$ be independent, uniform [0, 1] random variables, independent of the data $D_{n}$, and consider the augmented sample

$$
D_{n}^{\prime}=\left(\left(X_{1}^{\prime}, Y_{1}\right), \ldots,\left(X_{n}^{\prime}, Y_{n}\right)\right)
$$

where $X_{i}^{\prime}=\left(X_{i}, U_{i}\right)$. For each $h \geq 0$, introduce the rule

$$
\begin{aligned}
& g_{n, h}^{\prime}\left(x^{\prime}\right) \\
& \quad= \begin{cases}1 & \text { if } \sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h}\right)\left(2 Y_{i}-1\right)-\sum_{i=1}^{n} K(0)\left(2 Y_{i}-1\right) I_{\left\{u=U_{i}\right\}}>0 \\
0 & \text { otherwise },\end{cases}
\end{aligned}
$$

where $x^{\prime}=(x, u) \in \mathcal{R}^{d} \times[0,1]$. Observe that minimizing the resubstitution estimate $L_{n}^{(R)}\left(g_{n, h}^{\prime}\right)$ over these rules yields the same $h$ as minimizing $L_{n}^{(D)}\left(g_{n, h}\right)$ over the original kernel rules. Furthermore, $\mathbf{E}\left\{L\left(g_{n}\right)\right\}=\mathbf{E}\left\{L\left(g_{n}^{\prime}\right)\right\}$ if $g_{n}^{\prime}$ is obtained by minimizing $L_{n}^{(R)}\left(g_{n, h}^{\prime}\right)$. It follows from the universal consistency of kernel rules (Theorem 10.1) that

$$
\lim _{n \rightarrow \infty} \inf _{h} L\left(g_{n, h}^{\prime}\right) \rightarrow L^{*} \quad \text { with probability one. }
$$

Thus, it suffices to investigate $L\left(g_{n}^{\prime}\right)-\inf _{h} L\left(g_{n, h}^{\prime}\right)$. For a given $D_{n}^{\prime}$, define the subsets $A_{h}^{\prime} \subset \mathcal{R}^{d} \times[0,1]$ by

$$
A_{h}^{\prime}=\left\{x^{\prime}: g_{n, h}^{\prime}\left(x^{\prime}\right)=1\right\}, \quad h \in[0, \infty) .
$$

Arguing as in the previous section, we see that

$$
L\left(g_{n}^{\prime}\right)-\inf _{h} L\left(g_{n, h}^{\prime}\right) \leq 4 \sup _{h}\left|v_{n}\left(A_{h}^{\prime}\right)-v\left(A_{h}^{\prime}\right)\right|,
$$

where $v$ is the measure of $X^{\prime}$ on $\mathcal{R}^{d} \times[0,1]$, and $v_{n}$ is the empirical measure determined by $D_{n}^{\prime}$. Observe that each $A_{h}^{\prime}$ can be written as

$$
A_{h}^{\prime}=A_{h} \times\left\{[0,1]-\left\{U_{1}, \ldots, U_{n}\right\}\right\} \cup B_{h}^{\prime},
$$

where

$$
A_{h}=\left\{x \in \mathcal{R}^{d}: \sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h}\right)\left(2 Y_{i}-1\right)>0\right\}
$$

and

$$
\begin{aligned}
B_{h}^{\prime}= & \left\{(x, u) \in \mathcal{R}^{d} \times\left\{U_{1}, \ldots, U_{n}\right\}:\right. \\
& \left.\sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h}\right)\left(2 Y_{i}-1\right)-K(0) \sum_{i=1}^{n}\left(2 Y_{i}-1\right) I_{\left\{u=U_{i}\right\}}>0\right\} .
\end{aligned}
$$

Clearly, as

$$
v_{n}\left(A_{h}^{\prime}\right)=v_{n}\left(B_{h}^{\prime}\right) \quad \text { and } \quad v\left(A_{h}^{\prime}\right)=\mu\left(A_{h}\right)
$$

it suffices to prove that $\sup _{h}\left|\nu_{n}\left(B_{h}^{\prime}\right)-\mu\left(A_{h}\right)\right| \rightarrow 0$ in probability, as $n \rightarrow \infty$. Then

$$
\begin{aligned}
& \sup _{h}\left|v_{n}\left(B_{h}^{\prime}\right)-\mu\left(A_{h}\right)\right| \\
& \quad \leq \sup _{h}\left|v_{n}\left(B_{h}^{\prime}\right)-v\left(B_{h}^{\prime}\right)\right|+\sup _{h}\left|v\left(B_{h}^{\prime}\right)-\mu\left(A_{h}\right)\right| \\
& \quad \leq \sup _{B^{\prime} \subset \mathcal{R}^{d} \times\left\{U_{1}, \ldots, U_{n}\right\}}\left|v_{n}\left(B^{\prime}\right)-v\left(B^{\prime}\right)\right|+\sup _{h}\left|v\left(B_{h}^{\prime}\right)-\mu\left(A_{h}\right)\right| .
\end{aligned}
$$

The first term on the right-hand side tends to zero in probability, which may be seen by Problem 25.17 and the independence of the $U_{i}$ 's of $D_{n}$. To bound the second term, observe that for each $h$,

$$
\begin{aligned}
\left|v\left(B_{h}^{\prime}\right)-\mu\left(A_{h}\right)\right| & \leq \mu\left(\left\{x:\left|\sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h}\right)\left(2 Y_{i}-1\right)\right| \leq K(0)\right\}\right) \\
& =\sum_{j=1}^{N} \mathbf{P}\left\{X=x_{j}\right\} I_{\left\{\left|\sum_{i=1}^{n} K\left(\frac{x_{j}-X_{i}}{h}\right)\left(2 Y_{i}-1\right)\right| \leq K(0)\right\}},
\end{aligned}
$$

where $x_{1}, \ldots, x_{N}$ are the atoms of $X$. Thus, the proof is finished if we show that for each $x_{j}$,

$$
\mathbf{E}\left\{\sup _{h} I_{\left\{\left|\sum_{i=1}^{n} K\left(\frac{x_{j}-x_{i}}{h}\right)\left(2 Y_{i}-1\right)\right| \leq K(0)\right\}}\right\} \rightarrow 0 .
$$

Now, we exploit the special form $K=I_{A}$ of the kernel. Observe that

$$
\begin{aligned}
\sup _{h} & I_{\left\{\left|\sum_{i=1}^{n} K\left(\frac{x_{j}-x_{i}}{h}\right)\left(2 Y_{i}-1\right)\right| \leq K(0)\right\}} \\
& =\sup _{h} I_{\left\{\left|\sum_{i=1}^{n} I_{A}\left(\frac{x_{j}-X_{i}}{h}\right)\left(2 Y_{i}-1\right)\right| \leq 1\right\}} \\
& \leq \sum_{k=1}^{N} I_{\left\{\left|V_{1}+\cdots+V_{k}\right| \leq 1\right\}},
\end{aligned}
$$

where

$$
V_{k}=\sum_{i: X_{i}=x_{(k)}}\left(2 Y_{i}-1\right)
$$

and $x_{(1)}, x_{(2)}, \ldots, x_{(N)}$ is an ordering of the atoms of $X$ such that $x_{(1)}=x_{j}$ and $\left(x_{j}-x_{(k)}\right) / h \in A$ implies $\left(x_{j}-x_{(k-1)}\right) / h \in A$ for $1<k \leq N$. By properties of the binomial distribution, as $n \rightarrow \infty$,

$$
\mathbf{E}\left\{\sum_{k=1}^{N} I_{\left\{\left|V_{1}+\cdots+V_{k}\right| \leq 1\right\}}\right\}=O\left(\frac{1}{\sqrt{n}}\right) .
$$

This completes the proof.

Historical remark. In an early paper by Habbema, Hermans, and van den Broek (1974), the deleted estimate is used to select an appropriate subspace for the kernel rule. The kernel rules in turn have smoothing factors that are selected by maximum likelihood.

### 25.8 Sieve Methods

Sieve methods pick a best estimate or rule from a limited class of rules. For example, our sieve $\mathcal{C}_{k}$ might consist of rules of the form

$$
\phi(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{k} a_{i} K\left(\frac{x-x_{i}}{h_{i}}\right)>0 \\ 0 & \text { otherwise }\end{cases}
$$

where the $a_{i}$ 's are real numbers, the $x_{i}$ 's are points from $\mathcal{R}^{d}$, and the $h_{i}$ 's are positive numbers. There are formally $(d+2) k$ free scalar parameters. If we pick a rule $\phi_{n}^{*}$ that minimizes the empirical error on $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, we are governed by the theorems of Chapters 12 and 18 , and we will need to find the VC dimension of $\mathcal{C}_{k}$. For this, conditions on $K$ will be needed. We will return to this question in Chapter 30 on neural networks, as they are closely related to the sieves described here.

### 25.9 Squared Error Minimization

Many researchers have considered the problem of selecting $h$ in order to minimize the $L_{2}$ error (integrated squared error) of the kernel estimate

$$
\widehat{\eta}_{n, h}=\frac{\frac{1}{n} \sum_{i=1}^{n} Y_{i} K\left(\frac{X_{i}-x}{h}\right)}{\frac{1}{n} \sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right)}
$$

of the regression function $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$,

$$
\int\left(\widehat{\eta}_{n, h}(x)-\eta(x)\right)^{2} \mu(d x) .
$$

For example, Härdle and Marron (1985) proposed and studied a cross-validation method for choosing the optimal $h$ for the kernel regression estimate. They obtain asymptotic optimality for the integrated squared error. Although their method gives us a choice for $h$ if we consider $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$ as the regression function, it is not clear that the $h$ thus obtained is optimal for the probability of error. In fact, as the following theorem illustrates, for some distributions, the smoothing parameter that minimizes the $L_{2}$ error yields a rather poor error probability compared to that corresponding to the optimal $h$.

Theorem 25.9. Let $d=1$. Consider the kernel classification rule with the window kernel $K=I_{[-1,1]}$, and smoothing parameter $h>0$. Denote its error probability by $L_{n}(h)$. Let $h^{*}$ be the smoothing parameter that minimizes the mean integrated squared error

$$
\mathbf{E}\left\{\int\left(\widehat{\eta}_{n, h}(x)-\eta(x)\right)^{2} \mu(d x)\right\}
$$

Then for some distributions

$$
\lim _{n \rightarrow \infty} \frac{\mathbf{E}\left\{L_{n}\left(h^{*}\right)\right\}}{\inf _{h} \mathbf{E}\left\{L_{n}(h)\right\}}=\infty
$$

and the convergence is exponentially fast.
We leave the details of the proof to the reader (Problem 25.20). Only a rough sketch is given here. Consider $X$ uniform on $[0,1] \cup[3,4]$, and define

$$
\eta(x)= \begin{cases}1-x / 2 & \text { if } x \in[0,1] \\ 2-x / 2 & \text { otherwise }\end{cases}
$$

The optimal value of $h$ (i.e., the value minimizing the error probability) is one. It is constant, independent of $n$. This shows that we should not a priori exclude any values of $h$, as is commonly done in studies on regression and density estimation. The minimal error probability can be bounded from above (using Hoeffding's inequality) by $e^{-c_{1} n}$ for some constant $c_{1}>0$. On the other hand, straightforward calculations show that the smoothing factor $h^{*}$ that minimizes the mean integrated
squared error goes to zero as $n \rightarrow \infty$ as $O\left(n^{-1 / 4}\right)$. The corresponding error probability is larger than $e^{-c_{2} h^{*} n}$ for some constant $c_{2}$. The order-of-magnitude difference between the exponents explains the exponential speed of convergence to infinity.

## Problems and Exercises

Problem 25.1. Prove Theorem 25.2.
Problem 25.2. Show that for any $r>0$, the kernel

$$
K(x)= \begin{cases}1 & \text { if }\|x\| \leq 1 \\ 1 /\|x\|^{r} & \text { otherwise }\end{cases}
$$

satisfies the conditions of Theorem 25.3.
Problem 25.3. Assume that $K$ is a regular kernel with kernel complexity $\kappa_{m} \leq m^{\gamma}$ for some constant $\gamma>0$. Let $g_{n}$ be selected so as to minimize the holdout error estimate $\widehat{L}_{l, m}\left(\phi_{m}\right)$ over $\mathcal{C}_{m}$, the class of kernel rules based upon the first $m$ data points, with smoothing factor $h>0$. Assume furthermore that we vary $l$ over $\left[\log ^{2} n, n / 2\right]$, and that we pick the best $l$ (and $m=n-l$ ) by minimizing the holdout error estimate again. Show that the obtained rule is strongly universally consistent.

Problem 25.4. Prove that the kernel complexity $\kappa_{m}$ of the de la Vallée-Poussin kernel $K(x)=(\sin x / x)^{2}, x \in \mathcal{R}$, is infinite when $m \geq 2$.

Problem 25.5. Show that $\kappa_{m}=\infty$ for $m \geq 2$ when $K(x)=\cos ^{2}(x), x \in \mathcal{R}$.
Problem 25.6. Compute an upper bound for the kernel complexity $\kappa_{m}$ for the following kernels, where $x=\left(x^{(1)}, \ldots, x^{(d)}\right) \in \mathcal{R}^{d}$ :

$$
\begin{array}{ll}
\text { A. } & K(x)=\prod_{i=1}^{d}\left(1-x^{(i)^{2}}\right) I_{\left\{\left\{x^{(i)} \mid \leq 1\right\}\right.} \\
\text { B. } & K(x)=\frac{1}{\left(1+\|x\|^{2}\right)^{\alpha}}, \quad \alpha>0 \\
\text { C. } & K(x)=\prod_{i=1}^{d} \frac{1}{1+x^{(i)^{2}}} \\
\text { D. } & K(x)=\prod_{i=1}^{d} \cos \left(x^{(i)}\right) I_{\left\{\left|x^{(i)}\right| \leq \pi / 2\right\}}
\end{array}
$$

Problem 25.7. Can you construct a kernel on $\mathcal{R}$ with the property that its complexity satisfies $2^{m} \leq \kappa_{m}<\infty$ for all $m$ ? Prove your claim.

Problem 25.8. Show that for kernel classes $\mathcal{C}_{m}$ with kernel rules having a fixed training sequence $D_{m}$ but variable $h>0$, we have $V_{\mathcal{C}_{m}} \leq \log _{2} \kappa_{m}$.

Problem 25.9. Calculate upper bounds for $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ when $\mathcal{C}_{m}$ is the class of kernel rules based on fixed training data but with variable parameter $\theta$ in the following cases (for definitions, see Section 25.4):
(1) $K_{\theta}$ is a product kernel of $\mathcal{R}^{d}$, where the unidimensional kernel $K$ has kernel complexity $\kappa_{m}$.
(2) $K_{\theta}=\frac{1}{1+\|x\|^{\alpha} / h^{\alpha}}$, where $\theta=(h, \alpha), h>0, \alpha>0$ are two parameters.
(3) $K_{\theta}(x)=I_{\{x \in A\}}$, where $A$ is any ellipsoid of $\mathcal{R}^{d}$ centered at the origin.

Problem 25.10. Prove or disprove: if $D_{m}$ is fixed and $\mathcal{C}_{m}$ is the class of all kernel rules based on $D_{m}$ with $K=I_{A}, A$ being any convex set of $\mathcal{R}^{d}$ containing the origin, is it possible that $V_{\mathcal{C}_{m}}=\infty$, or is $V_{\mathcal{C}_{m}}<\infty$ for all possible configurations of $D_{m}$ ?

Problem 25.11. Let

$$
\mathcal{A}_{8}=\left\{\left\{x: \sum_{i=1}^{8} K\left(\frac{x-X_{i}}{h}\right)\left(2 Y_{i}-1\right)>0\right\}, h>0\right\}
$$

with $X_{1}=-3 \delta, X_{2}=X_{3}=X_{4}=-\delta, X_{5}=X_{6}=X_{7}=\delta, X_{8}=3 \delta, Y_{1}=1, Y_{2}=Y_{3}=Y_{4}=$ $-1, Y_{5}=Y_{6}=Y_{7}=1, Y_{8}=-1$, and let $K$ be piecewise cubic. Extending the quadratic example in the text, show that the vc dimension of $\mathcal{A}_{8}$ is infinite for some $K$ in this class that is symmetric, unimodal, positive, and convex on $[0, \infty)$.

Problem 25.12. Draw $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{4}, Y_{4}\right)$ from the distribution of $(X, Y)$ on $\mathcal{R}^{d} \times\{0,1\}$, where $X$ has a density and $\eta(x) \in(0,1)$ at all $x$. Find a symmetric unimodal $K$ such that

$$
\mathcal{A}_{4}=\left\{\left\{x: \sum_{i=1}^{4} K\left(\frac{x-X_{i}}{h}\right)\left(2 Y_{i}-1\right)>0\right\}, h>0\right\}
$$

has vc dimension satisfying $\mathbf{P}\left\{V_{\mathcal{A}_{4}}=\infty\right\}>0$. Can you find such a kernel $K$ with a bounded support?

Problem 25.13. Let $X$ have a density $f$ on $\mathcal{R}^{d}$ and let $X_{1}, \ldots, X_{n}$ be i.i.d., drawn from $f$. Show that

$$
\lim _{c \rightarrow \infty} \liminf _{n \rightarrow \infty} \mathbf{P}\left\{\min _{i \neq j}\left\|X_{i}-X_{j}\right\|^{d} \leq \frac{c}{n^{2}}\right\}=1
$$

Apply this result in the following situation: define $H_{n}=\min _{i, j: Y_{i}=1, Y_{j}=0}\left\|X_{i}-X_{j}\right\|$, where $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ are i.i.d. $\mathcal{R}^{d} \times\{0,1\}$-valued random variables distributed as $(X, Y)$, with $X$ absolutely continuous, $Y$ independent of $X$, and $\mathbf{P}\{Y=1\} \in(0,1)$. Show that

$$
\lim _{c \rightarrow \infty} \liminf _{n \rightarrow \infty} \mathbf{P}\left\{H_{n}^{d} \leq \frac{c}{n^{2}}\right\}=1
$$

Conclude that $n H^{d} \rightarrow 0$ in probability. If you have a kernel rule with kernel $S_{0,1}$ on $\mathcal{R}^{d}$, and if the smoothing factor $H_{n}$ is random but satisfies $n H_{n}^{d} \rightarrow 0$ in probability, then

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{L_{n}\right\}=\mathbf{P}\{Y=1\}
$$

whenever $X$ has a density. Show this.
Problem 25.14. Consider the variable kernel rule based upon the variable kernel density estimate of Breiman, Meisel, and Purcell (1977) and studied by Krzyzak (1983)

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i: Y_{i}=1} K_{H_{i}}\left(x-X_{i}\right)>\sum_{i: Y_{i}=0} K_{H_{i}}\left(x-X_{i}\right) \\ 0 & \text { otherwise }\end{cases}
$$

Here $K$ is a positive-valued kernel, $K_{u}=\left(1 / u^{d}\right) K(x / u)$ for $u>0$, and $H_{i}$ is the distance between $X_{i}$ and the $k$-th nearest neighbor of $X_{i}$ among $X_{j}, j \neq i, 1 \leq j \leq n$. Investigate the consistency of this rule when $k / n \rightarrow 0$ and $k \rightarrow \infty$ and $X_{1}$ has a density.

Problem 25.15. Continuation. Fix $k=1$, and let $K$ be the normal density in $\mathcal{R}^{d}$. If $X_{1}$ has a density, what can you say about the asymptotic probability of error of the variable kernel rule? Is the inequality of Cover and Hart still valid? Repeat the exercise for the uniform kernel on the unit ball of $\mathcal{R}^{d}$.

Problem 25.16. If $X_{1}, \ldots, X_{n}$ are discrete i.i.d. random variables and $N$ denotes the number of different values taken by $X_{1}, \ldots, X_{n}$, then

$$
\lim _{n \rightarrow \infty} \frac{\mathbf{E}\{N\}}{n}=0,
$$

and $N / n \rightarrow 0$ with probability one. Hint: For the weak convergence, assume without loss of generality that the probabilities are monotone. For the strong convergence, use McDiarmid's inequality.

Problem 25.17. Let $\mathcal{B}$ be the class of all Borel subsets of $\mathcal{R}^{d}$. Using the previous exercise, show that for any discrete distribution,

$$
\sup _{A \in \mathcal{B}}\left|\mu_{n}(A)-\mu(A)\right| \rightarrow 0 \quad \text { with probability one. }
$$

Hint: Recall the necessary and sufficient condition $\mathbb{E}\left\{\log N_{A}\left(X_{1}, \ldots, X_{n}\right)\right\} / n \rightarrow 0$ from Chapter 12.

Problem 25.18. Prove Theorem 25.8 allowing kernel functions taking finitely many different values.

Problem 25.19. Let $X_{1}, \ldots, X_{n}$ be an i.i.d. sample drawn from the distribution of $X$. Let $X_{j}^{N N}$ denote the nearest neighbor of $X_{j}$ among $X_{1}, \ldots, X_{j-1}, X_{j+1}, \ldots, X_{n}$. Define

$$
B_{n}=\max _{j}\left\|X_{j}-X_{j}^{N N}\right\|
$$

(1) Show that for all nonatomic distributions of $X, n B_{n}^{d} \rightarrow \infty$ with probability one.
(2) Is it true that for every $X$ with a density, there exists a constant $c>0$ such that with probability one, $n B_{n}^{d} \geq c \log n$ for all $n$ large enough?
(3) Exhibit distributions on the real line for which $B_{n} \rightarrow \infty$ with probability one.

Hint: Look at the difference between the first and second order statistics.
Problem 25.20. Prove Theorem 25.9. Hint: Use the example given in the text. Get an upper bound for the error probability corresponding to $h=1$ by Hoeffding's inequality. The mean integrated squared error can be computed for every $h$ in a straightforward way by observing that

$$
\mathbf{E}\left\{\int\left(\widehat{\eta}_{n, h}(x)-\eta(x)\right)^{2} \mu(d x)\right\}=\int \operatorname{Var}\left\{\widehat{\eta}_{n, h}(x)\right\}+\left(\mathbf{E}\left\{\widehat{\eta}_{n, h}(x)\right\}-\eta(x)\right)^{2} \mu(d x)
$$

Split the integral between 0 and 1 in three parts, 0 to $h, h$ to $1-h$, and $1-h$ to 1 . Setting the derivative of the obtained expression with respect to $h$ equal to zero leads to a third-order
equation in $h$, whose roots are $O\left(n^{-1 / 4}\right)$. To get a lower bound for the corresponding error probability, use the crude bound

$$
\mathbf{E}\left\{L_{n}(h)\right\} \geq \frac{1}{2} \int_{1-h}^{1} \mathbf{P}\left\{g_{n}(X) \neq Y \mid X=x\right\} d x .
$$

Now, estimate the tail of a binomial distribution from below; and use Stirling's formula to show that, modulo polynomial factors, the error probability is larger than $2^{-n^{3 / 4}}$.

## 26

## Automatic Nearest Neighbor Rules

The error probability of the $k$-nearest neighbor rule converges to the Bayes risk for all distributions when $k \rightarrow \infty$, and $k / n \rightarrow 0$ as $n \rightarrow \infty$. The convergence result is extended here to include data-dependent choices of $k$. We also look at the data-based selection of a metric and of weights in weighted nearest neighbor rules.
To keep the notation consistent with that of earlier chapters, random (data-based) values of $k$ are denoted by $K_{n}$. In most instances, $K_{n}$ is merely a function of $D_{n}$, the data sequence $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$. The reader should not confuse $K_{n}$ with the kernel $K$ in other chapters.

### 26.1 Consistency

We start with a general theorem assessing strong consistency of the $k$-nearest neighbor rule with data-dependent choices of $k$. For the sake of simplicity, we assume the existence of the density of $X$. The general case can be taken care of by introducing an appropriate tie-breaking method as in Chapter 11.

Theorem 26.1. Let $K_{1}, K_{2}, \ldots$ be integer valued random variables, and let $g_{n}$ be the $K_{n}$-nearest neighbor rule. If $X$ has a density, and

$$
K_{n} \rightarrow \infty \text { and } K_{n} / n \rightarrow 0 \text { with probability one as } n \rightarrow \infty \text {, }
$$

then $L\left(g_{n}\right) \rightarrow L^{*}$ with probability one.

PROOF. $L\left(g_{n}\right) \rightarrow L^{*}$ with probability one if and only if for every $\epsilon>0, I_{\left\{L\left(g_{n}\right)-L^{*}>\epsilon\right\}}$ $\rightarrow 0$ with probability one. Clearly, for any $\beta>0$,

$$
I_{\left\{L\left(g_{n}\right)-L^{*}>\epsilon\right\}} \leq I_{\left\{1 / K_{n}+K_{n} / n<2 \beta, L\left(g_{n}\right)-L^{*}>\epsilon\right\}}+I_{\left\{1 / K_{n}+K_{n} / n \geq 2 \beta\right\}} .
$$

We are done if both random variables on the right-hand side converge to zero with probability one. The convergence of the second term for all $\beta>0$ follows trivially from the conditions of the theorem. The convergence of the first term follows from the remark following Theorem 11.1, which states that for any $\epsilon>0$, there exist $\beta>0$ and $n_{0}$ such that for the error probability $L_{n, k}$ of the $k$-nearest neighbor rule,

$$
\mathbf{P}\left\{L_{n, k}-L^{*}>\epsilon\right\} \leq 4 e^{-c n \epsilon^{2}}
$$

for some constant $c$ depending on the dimension only, provided that $n>n_{0}$, $k>1 / \beta$ and $k / n<\beta$. Now clearly,

$$
\begin{aligned}
\mathbf{P}\left\{L\left(g_{n}\right)-L^{*}>\epsilon, 1 / K_{n}+K_{n} / n<2 \beta\right\} & \leq \mathbf{P}\left\{\sup _{1 / \beta \leq k \leq n \beta} L_{n, k}-L^{*}>\epsilon\right\} \\
& \leq n \sup _{1 / \beta \leq k \leq n \beta} \mathbf{P}\left\{L_{n, k}-L^{*}>\epsilon\right\}
\end{aligned}
$$

by the union bound. Combining this with Theorem 11.1, we get

$$
\mathbf{P}\left\{L\left(g_{n}\right)-L^{*}>\epsilon, 1 / K_{n}+K_{n} / n<2 \beta\right\} \leq 4 n e^{-c n \epsilon^{2}}, \quad n \geq n_{0}
$$

The Borel-Cantelli lemma implies that

$$
I_{\left\{1 / K_{n}+K_{n} / n<2 \beta, L\left(g_{n}\right)-L^{*}>\epsilon\right\}} \rightarrow 0
$$

with probability one, and the theorem is proved.

Sometimes we only know that $K_{n} \rightarrow \infty$ and $K_{n} / n \rightarrow 0$ in probability. In such cases weak consistency is guaranteed. The proof is left as an exercise (Problem 26.1).

Theorem 26.2. Let $K_{1}, K_{2}, \ldots$ be integer valued random variables, and let $g_{n}$ be the $K_{n}$-nearest neighbor rule. If $X$ has a density, and

$$
K_{n} \rightarrow \infty \quad \text { and } \quad K_{n} / n \rightarrow 0 \quad \text { in probability as } n \rightarrow \infty
$$

then $\lim _{n \rightarrow \infty} \mathbf{E} L\left(g_{n}\right)=L^{*}$, that is, $g_{n}$ is weakly consistent.

### 26.2 Data Splitting

Consistency by itself may be obtained by choosing $k=\lfloor\sqrt{n}\rfloor$, but few-if anyusers will want to blindly use such recipes. Instead, a healthy dose of feedback from the data is preferable. If we proceed as in Chapter 22, we may split
the data sequence $D_{n}=\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ into a training sequence $D_{m}=$ $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{m}, Y_{m}\right)$, and a testing sequence $T_{l}=\left(X_{m+1}, Y_{m+1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, where $m+l=n$. The training sequence $D_{m}$ is used to design a class of classifiers $\mathcal{C}_{m}$. The testing sequence is used to select a classifier from $\mathcal{C}_{m}$ that minimizes the holdout estimate of the error probability,

$$
\widehat{L}_{m, l}\left(\phi_{m}\right)=\frac{1}{l} \sum_{i=1}^{l} I_{\left\{\phi_{m}\left(X_{n+i}\right) \neq Y_{n+i}\right\}} .
$$

If $\mathcal{C}_{m}$ contains all $k$-nearest neighbor rules with $1 \leq k \leq m$, then $\left|\mathcal{C}_{m}\right|=m$. Therefore, we have

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)>\epsilon \mid D_{m}\right\} \leq 2 m e^{-l \epsilon^{2} / 2},
$$

where $g_{n}$ is the selected $k$-nearest neighbor rule. By combining Theorems 11.1 and 22.1 , we immediately deduce that $g_{n}$ is universally consistent if

$$
\lim _{n \rightarrow \infty} m=\infty, \quad \lim _{n \rightarrow \infty} \frac{l}{\log m}=\infty
$$

It is strongly universally consistent if $\lim _{n \rightarrow \infty} l / \log n=\infty$ also. Note too that

$$
\mathbf{E}\left\{\left|L\left(g_{n}\right)-\inf _{\phi_{m} \in C_{m}} L\left(\phi_{m}\right)\right|\right\} \leq 2 \sqrt{\frac{\log (2 m)+1}{2 l}}
$$

(by Problem 12.1), so that it is indeed important to pick $l$ much larger than $\log m$.

### 26.3 Data Splitting for Weighted NN Rules

Royall (1966) introduced the weighted nN rule in which the $i$-th nearest neighbor receives weight $w_{i}$, where $w_{1} \geq w_{2} \geq \ldots \geq w_{k} \geq 0$ and the $w_{i}$ 's sum to one. We assume that $w_{k+1}=\cdots=w_{n}=0$ if there are $n$ data points. Besides the natural appeal of attaching more weight to nearer neighbors, there is also a practical byproduct: if the $w_{i}$ 's are all of the form $1 / z_{i}$, where $z_{i}$ is a prime integer, then no two subsums of $w_{i}$ 's are equal, and therefore voting ties are avoided altogether.
Consider now data splitting in which $\mathcal{C}_{m}$ consists of all weighted $k$-NN rules as described above-clearly, $k \leq m$ now. As $\left|\mathcal{C}_{m}\right|=\infty$, we compute the shatter coefficients $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$. We claim that

$$
\mathcal{S}\left(\mathcal{C}_{m}, l\right) \leq \begin{cases}\sum_{2^{l}}^{k}  \tag{26.1}\\ l_{i=1}^{l}\binom{l}{i} \leq l^{k} & \text { if } l \geq k \\ \text { if } l<k\end{cases}
$$

This result is true even if we do not insist that $w_{1} \geq w_{2} \geq \ldots \geq w_{k} \geq 0$.
Proof of (26.1). Each $X_{j}$ in the testing sequence is classified based upon the sign of $\sum_{i=1}^{k} a_{i j} w_{i}$, where $a_{i j} \in\{-1,1\}$ depends upon the class of the $i$-th nearest
neighbor of $X_{j}$ among $X_{1}, \ldots, X_{m}$ (and does not depend upon the $w_{i}$ 's). Consider the $l$-vector of signs of $\sum_{i=1}^{k} a_{i j} w_{i}, m<j \leq n$. In the computation of $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$, we consider the $a_{i j}$ 's as fixed numbers, and vary the $w_{i}$ 's subject to the condition laid out above. Here is the crucial step in the argument: the collection of all vectors ( $w_{1}, \ldots, w_{k}$ ) for which $X_{j}$ is assigned to class 1 is a linear halfspace of $\mathcal{R}^{k}$. Therefore, $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ is bounded from above by the number of cells in the partition of $\mathcal{R}^{k}$ defined by $l$ linear halfspaces. This is bounded by $\sum_{i=1}^{k}\binom{l}{i}$ (see Problem 22.1 ) if $k \leq l$.

Let $g_{n}$ be the rule in $\left|\mathcal{C}_{m}\right|$ that minimizes the empirical error committed on the test sequence $\left(X_{m+1}, Y_{m+1}\right), \ldots,\left(X_{n}, Y_{n}\right)$. Then by (22.1), if $l=n-m \geq k$, we have

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)>\epsilon \mid D_{m}\right\} \leq 4 e^{8} l^{2 k} e^{-l \epsilon^{2} / 2} .
$$

If $k \rightarrow \infty$ (which implies $m \rightarrow \infty$ ), we have universal consistency when $k \log (l) / l \rightarrow 0$. The estimation error is of the order of $\sqrt{k \log l / l}$-in the terminology of Chapter 22, the rule is $\sqrt{k \log l / l}$-optimal. This error must be weighed against the unknown approximation error. Let us present a quick heuristic argument. On the real line, there is compelling evidence to suggest that when $X$ has a smooth density, $k=\mathrm{cm}^{4 / 5}$ is nearly optimal. With this choice, if both $l$ and $m$ grow linearly in $n$, the estimation error is of the order of $\sqrt{\log n} / n^{1 / 10}$. This is painfully large--to reduce this error by a factor of two, sample sizes must rise by a factor of about 1000 . The reason for this disappointing result is that $\mathcal{C}_{m}$ is just too rich for the values of $k$ that interest us. Automatic selection may lead to rules that overfit the data.

If we restrict $\mathcal{C}_{m}$ by making $m$ very small, the following rough argument may be used to glean information about the size of $m$ and $l=n-m$. We will take $k=m \ll n$. For smooth regression function $\eta$, the estimation error may be anywhere between $m^{-2 / 5}$ and $m^{-4 / 5}$ on the real line. As the estimation error is of the order of $\sqrt{m \log n / n}$, equating the errors leads to the rough recipe that $m \approx(n / \log n)^{5 / 9}$, and $m \approx(n / \log n)^{5 / 13}$, respectively. Both errors are then about $(n / \log n)^{-2 / 9}$ and $(n / \log n)^{-4 / 9}$, respectively. This is better than with the previous example with $m$ linear in $n$. Unfortunately, it is difficult to test whether the conditions on $\eta$ that guarantee certain errors are satisfied. The above procedure is thus doomed to remain heuristic.

### 26.4 Reference Data and Data Splitting

Split the data into $D_{m}$ and $T_{l}$ as is done in the previous section. Let $\mathcal{C}_{m}$ contain all $1-\mathrm{NN}$ rules that are based upon the data $\left(x_{1}, y_{1}\right), \ldots,\left(x_{k}, y_{k}\right)$, where $k \leq m$ is to be picked, $\left\{x_{1}, \ldots, x_{k}\right\} \subset\left\{X_{1}, \ldots X_{m}\right\}$, and $\left\{y_{1}, \ldots, y_{k}\right\} \in\{0,1\}^{k}$. Note that because the $y_{i}$ 's are free parameters, $\left\{\left(x_{1}, y_{1}\right), \ldots,\left(x_{k}, y_{k}\right)\right\}$ is not necessarily a subset of $\left\{\left(X_{1}, Y_{1}\right), \ldots,\left(X_{m}, Y_{m}\right)\right\}$-this allows us to flip certain $y$-values at some
data points. Trivially, $\left|\mathcal{C}_{m}\right|=\binom{m}{k} 2^{k}$. Hence,

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi_{m} \in \mathcal{C}_{m}} L\left(\phi_{m}\right)>\epsilon \mid D_{m}\right\} \leq 2^{k+1}\binom{m}{k} e^{-l \epsilon^{2} / 2},
$$

where $l=n-m$, and $g_{n} \in \mathcal{C}_{m}$ minimizes the empirical error on the test sequence $T_{l}$. The best rule in $\mathcal{C}_{m}$ is universally consistent when $k \rightarrow \infty$ (see Theorem 19.4). Therefore, $g_{n}$ is universally consistent when the above bound converges to zero. Sufficient conditions are
(i) $\quad \lim _{n \rightarrow \infty} l=\infty$;
(ii) $\lim _{n \rightarrow \infty} \frac{k \log m}{l}=0$.

As the estimation error is $O(\sqrt{k \log m / l})$, it is important to make $l$ large, while keeping $k$ small.

The selected sequence $\left(x_{1}, y_{1}\right), \ldots,\left(x_{k}, y_{k}\right)$ may be called reference data, as it captures the information in the larger data set. If $k$ is sufficiently small, the computation of $g_{n}(x)$ is extremely fast. The idea of selecting reference data or throwing out useless or "bad" data points has been proposed and studied by many researchers under names such as condensed nn rules, edited nN rules, and selective nN rules. See Hart (1968), Gates (1972), Wilson (1972), Wagner (1973), Ullmann (1974), Ritter et al. (1975), Tomek (1976b), and Devijver and Kittler (1980). See also Section 19.1.

### 26.5 Variable Metric NN Rules

The data may also be used to select a suitable metric for use with the $k$-NN rule. The metric adapts itself for certain scale information gleaned from the data. For example, we may compute the distance between $x_{1}$ and $x_{2}$ by the formula

$$
\begin{aligned}
\left\|A^{T}\left(x_{1}-x_{2}\right)\right\| & =\left(\left(x_{1}-x_{2}\right)^{T} A A^{T}\left(x_{1}-x_{2}\right)\right)^{1 / 2} \\
& =\left(\left(x_{1}-x_{2}\right)^{T} \Sigma\left(x_{1}-x_{2}\right)\right)^{1 / 2}
\end{aligned}
$$

where $\left(x_{1}-x_{2}\right)$ is a column vector, $(\cdot)^{T}$ denotes its transpose, $A$ is a $d \times d$ transformation matrix, and $\Sigma=A A^{T}$ is a positive definite matrix. The elements of $A$ or $\Sigma$ may be determined from the data according to some heuristic formulas. We refer to Fukunaga and Hostetler (1973), Short and Fukunaga (1981), Fukunaga and Flick (1984), and Myles and Hand (1990) for more information.

For example, the object of principal component analysis is to find a transformation matrix $A$ such that the components of the vector $A^{T} X$ have unit variance and are uncorrelated. These methods are typically based on estimating the eigenvalues of the covariance matrix of $X$. For such situations, we prove the following general consistency result:

Theorem 26.3. Let the random metric $\rho_{n}$ be of the form

$$
\rho_{n}(x, y)=\left\|A_{n}^{T}(x-y)\right\|,
$$

where the matrix $A_{n}$ is a function of $X_{1}, \ldots, X_{n}$. Assume that distance ties occur with zero probability, and there are two sequences of nonnegative random variables $\left\{m_{n}\right\}$ and $\left\{M_{n}\right\}$ such that for any $n$ and $x, y \in \mathcal{R}^{d}$,

$$
m_{n}\|x-y\| \leq \rho_{n}(x, y) \leq M_{n}\|x-y\|
$$

and

$$
\mathbf{P}\left\{\liminf _{n \rightarrow \infty} \frac{m_{n}}{M_{n}}=0\right\}=0
$$

If

$$
\lim _{n \rightarrow \infty} k_{n}=\infty \quad \text { and } \quad \lim _{n \rightarrow \infty} \frac{k_{n}}{n}=0
$$

then the $k_{n}$-nearest neighbor rule based on the metric $\rho_{n}$ is consistent.

Proof. We verify the three conditions of Theorem 6.3. In this case, $W_{n i}(X)=1 / k_{n}$ if $X_{i}$ is one of the $k_{n}$ nearest neighbors of $X$ (according to $\rho_{n}$ ), and zero otherwise. Condition (iii) holds trivially. Just as in the proof of consistency of the ordinary $k_{n}$ nearest neighbor rule, for condition (i) we need the property that the number of data points that can be among the $k_{n}$ nearest neighbors of a particular point is at most $k_{n} \gamma_{d}$, where the constant $\gamma_{d}$ depends on the dimension only. This is a deterministic property, and it can be proven exactly the same way as for the standard nearest neighbor rule. The only condition of Theorem 6.3 whose justification needs extra work is condition (ii): we need to show that for any $a>0$,

$$
\lim _{n \rightarrow \infty} \mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X) I_{\left[\left\|X-X_{i}\right\|>a\right\}}\right\}=0
$$

Denote the $k$-th nearest neighbor of $X$ according to $\rho_{n}$ by $\tilde{X}_{(k)}$, and the $k$-th nearest neighbor of $X$ according to the Euclidean metric by $X_{(k)}$. Then

$$
\begin{aligned}
& \mathbf{E}\left\{\sum_{i=1}^{n} W_{n i}(X) I_{\left\{\left\|X-X_{i}\right\|>a\right\}}\right\} \\
& =\frac{1}{k_{n}} \mathbf{E}\left\{\sum_{i=1}^{n} I_{\left\{\rho_{n}\left(X, X_{i}\right) \leq \rho_{n}\left(X, \tilde{X}_{\left(k_{n}\right)}\right),\left\|X-X_{i}\right\|>a\right\}}\right\} \\
& \leq \frac{1}{k_{n}} \mathbf{E}\left\{\sum_{i=1}^{n} I_{\left\{\rho_{n}\left(X, X_{i}\right) \leq \rho_{n}\left(X, \tilde{X}_{\left(k_{n}\right)}\right), \rho_{n}\left(X, X_{i}\right)>m_{n} a\right\}}\right\} \\
& \quad\left(\text { since } m_{n}\|x-y\| \leq \rho_{n}(x, y)\right)
\end{aligned}
$$

$$
\begin{aligned}
& =\frac{1}{k_{n}} \mathbf{E}\left\{\sum_{i=1}^{n} I_{\left\{m_{n} a \leq \rho_{n}\left(X, X_{i}\right) \leq \rho_{n}\left(X, \tilde{X}_{\left(k_{n}\right)}\right)\right\}}\right\} \\
& =\frac{1}{k_{n}} \mathbf{E}\left\{\sum_{j=1}^{k_{n}} I_{\left\{m_{n} a \leq \rho_{n}\left(X, \tilde{X}_{(j)}\right)\right\}}\right\} \\
& \leq \mathbf{P}\left\{m_{n} a \leq \rho_{n}\left(X, \tilde{X}_{\left(k_{n}\right)}\right)\right\} \\
& =\mathbf{P}\left\{\sum_{i=1}^{n} I_{\left\{\rho_{n}\left(X, X_{i}\right) \leq m_{n} a\right\}}<k_{n}\right\} \\
& \leq \mathbf{P}\left\{\sum_{i=1}^{n} I_{\left\{M_{n}\left\|X-X_{i}\right\| \leq m_{n} a\right\}}<k_{n}\right\} \\
& =\left(\operatorname{since} M_{n}\|x-y\| \geq \rho_{n}(x, y)\right) \\
& =\mathbf{P}\left\{\left\|X-X_{\left(k_{n}\right)}\right\|>a \frac{m_{n}}{M_{n}}\right\} .
\end{aligned}
$$

But we know from the consistency proof of the ordinary $k_{n}$-nearest neighbor rule in Chapter 11 that for each $a>0$,

$$
\lim _{n \rightarrow \infty} \mathbb{P}\left\{\left\|X-X_{\left(k_{n}\right)}\right\|>a\right\}=0
$$

It follows from the condition on $m_{n} / M_{n}$ that the probability above converges to zero as well.

The conditions of the theorem hold if, for example, the elements of the matrix $A_{n}$ converge to the elements of an invertible matrix $A$ in probability. In that case, we may take $m_{n}$ as the smallest, and $M_{n}$ as the largest eigenvalues of $A_{n}^{T} A_{n}$. Then $m_{n} / M_{n}$ converges to the ratio of the smallest and largest eigenvalues of $A^{T} A$, a positive number.

If we pick the elements of $A$ by minimizing the empirical error of a test sequence $T_{l}$ over $\mathcal{C}_{m}$, where $\mathcal{C}_{m}$ contains all $k$-NN rules based upon a training sequence $D_{m}$ (thus, the elements of $A$ are the free parameters), the value of $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ is too large to be useful-see Problem 26.3. Furthermore, such minimization is not computationally feasible.

### 26.6 Selection of $k$ Based on the Deleted Estimate

If you wish to use all the available data in the training sequence, without splitting, then empirical selection based on minimization of other estimates of the error probability may be your solution. Unfortunately, performance guarantees for the selected rule are rarely available. If the class of rules is finite, as when $k$ is selected from $\{1, \ldots, n\}$, there are useful inequalities. We will show you how this works.

Let $\mathcal{C}_{n}$ be the class of all $k$-nearest neighbor rules based on a fixed training sequence $D_{n}$, but with $k$ variable. Clearly, $\left|C_{n}\right|=n$. Assume that the deleted estimate is used to pick a classifier $g_{n}$ from $\mathcal{C}_{n}$ :

$$
L_{n}^{(D)}\left(g_{n}\right) \leq L_{n}^{(D)}\left(\phi_{n}\right) \text { for all } \phi_{n} \in \mathcal{C}_{n}
$$

We can derive performance bounds for $g_{n}$ from Theorem 24.5. Since the result gives poor bounds for large values of $k$, the range of $k$ 's has to be restricted-see the discussion following Theorem 24.5. Let $k_{0}$ denote the value of the largest $k$ allowed, that is, $\mathcal{C}_{n}$ now contains all $k$-nearest neighbor rules with $k$ ranging from 1 to $k_{0}$.

Theorem 26.4. Let $g_{n}$ be the classifier minimizing the deleted estimate of the error probability over $\mathcal{C}_{n}$, the class of $k$-nearest neighbor rules with $1 \leq k \leq k_{0}$. Then

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi_{n} \in \mathcal{C}_{n}} L\left(\phi_{n}\right)>\epsilon\right\} \leq k_{0} e^{-c n \epsilon^{3} / k_{0}}
$$

where $c$ is a constant depending on the dimension only. If $k_{0} \rightarrow \infty$ and $k_{0} \log n / n \rightarrow 0$, then $g_{n}$ is strongly universally consistent.

Proof. From Theorem 8.4 we recall

$$
L\left(g_{n}\right)-\inf _{\phi_{n} \in \mathcal{C}_{n}} L\left(\phi_{n}\right) \leq 2 \sup _{\phi_{n} \in \mathcal{C}_{n}}\left|L\left(\phi_{n}\right)-L_{n}^{(D)}\left(\phi_{n}\right)\right|
$$

The inequality now follows from Theorem 24.5 via the union bound. Universal consistency follows from the previous inequality, and the fact that the $k_{0}-\mathrm{NN}$ rule is strongly universally consistent (see Theorem 11.1).

## Problems and Exercises

Problem 26.1. Let $K_{1}, K_{2}, \ldots$ be integer valued random variables, and let $g_{n}$ be the $K_{n}$ nearest neighbor rule. Show that if $X$ has a density, and $K_{n} \rightarrow \infty$ and $K_{n} / n \rightarrow 0$ in probability as $n \rightarrow \infty$, then $\mathbf{E}\left\{L\left(g_{n}\right)\right\} \rightarrow L^{*}$.

Problem 26.2. Let $\mathcal{C}$ be the class of all 1 -NN rules based upon pairs $\left(x_{1}, y_{1}\right), \ldots,\left(x_{k}, y_{k}\right)$, where $k$ is a fixed parameter (possibly varying with $n$ ), and the ( $x_{i}, y_{i}$ )'s are variable pairs from $\mathcal{R}^{d} \times\{0,1\}$. Let $g_{n}$ be the rule that minimizes the empirical error over $\mathcal{C}$, or, equivalently; let $g_{n}$ be the rule that minimizes the resubstitution estimate $L_{n}^{(R)}$ over $\mathcal{C}$.
(1) Compute a suitable upper bound for $\mathcal{S}(\mathcal{C}, n)$.
(2) Compute a good upper bound for $V_{\mathcal{C}}$ as a function of $k$ and $d$.
(3) If $k \rightarrow \infty$ with $n$, show that the sequence of classes $\mathcal{C}$ contains a strongly universally consistent subsequence of rules (you may assume for convenience that $X$ has a density to avoid distance ties).
(4) Under what condition on $k$ can you guarantee strong universal consistency of $g_{n}$ ?
(5) Give an upper bound for

$$
\mathbf{E}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}} L(\phi)\right\} .
$$

Problem 26.3. Let $\mathcal{C}_{m}$ contain all $k$-NN rules based upon data pairs $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{m}, Y_{m}\right)$. The metric used in computing the neighbors is derived from the norm

$$
\|x\|_{\Sigma}=\sum_{i=1}^{d} \sum_{j=1}^{d} \sigma_{i j} x^{(i)} x^{(j)}, x=\left(x^{(1)}, \ldots, x^{(d)}\right)
$$

where $\left\{\sigma_{i j}\right\}$ forms a positive definite matrix $\Sigma$. The elements of $\Sigma$ are the free parameters in $\mathcal{C}_{m}$. Compute upper and lower bounds for $\mathcal{S}\left(\mathcal{C}_{m}, l\right)$ as a function of $m, l, k$, and $d$.

Problem 26.4. Let $g_{n}$ be the rule obtained by minimizing $L_{n}^{(D)}$ over all $k$-NN rules with $1 \leq k \leq n$. Prove or disprove: $g_{n}$ is strongly universally consistent. Note that in view of Theorem 26.4, it suffices to consider $\epsilon n / \log n \leq k \leq n-1$ for all $\epsilon>0$.

## 27

## Hypercubes and Discrete Spaces

In many situations, the pair $(X, Y)$ is purely binary, taking values in $\{0,1\}^{d} \times\{0,1\}$. Examples include boolean settings (each component of $X$ represents "on" or "off"), representations of continuous variables through quantization (continuous variables are always represented by bit strings in computers), and ordinal data (a component of $X$ is 1 if and only if a certain item is present). The components of $X$ are denoted by $X^{(1)}, \ldots, X^{(d)}$. In this chapter, we review pattern recognition briefly in this setup.

Without any particular structure in the distribution of $(X, Y)$ or the function $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}, x \in\{0,1\}^{d}$, the pattern recognition problem might as well be cast in the space of the first $2^{d}$ positive integers: $(X, Y) \in\left\{1, \ldots, 2^{d}\right\} \times\{0,1\}$. This is dealt with in the first section. However, things become more interesting under certain structural assumptions, such as the assumption that the components of $X$ be independent. This is dealt with in the third section. General discrimination rules on hypercubes are treated in the rest of the chapter.

### 27.1 Multinomial Discrimination

At first sight, discrimination on a finite set $\{1, \ldots, k\}$-called multinomial discrimination -may seem utterly trivial. Let us call the following rule the fundamental rule, as it captures what most of us would do in the absence of any additional information.

$$
g_{n}^{*}(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{n} I_{\left\{X_{i}=x, Y_{i}=1\right\}}>\sum_{i=1}^{n} I_{\left\{X_{i}=x, Y_{i}=0\right\}} \\ 0 & \text { otherwise }\end{cases}
$$

The fundamental rule coincides with the standard kernel and histogram rules if the smoothing factor or bin width are taken small enough. If $p_{i}=\mathbf{P}\{X=i\}$ and $\eta$ is as usual, then it takes just a second to see that

$$
L_{n} \geq \sum_{x: \sum_{i=1}^{n} I_{\left\{X_{i}=x\right\}}=0} \eta(x) p_{x}
$$

and

$$
\mathbf{E} L_{n} \geq \sum_{x} \eta(x) p_{x}\left(1-p_{x}\right)^{n}
$$

where $L_{n}=L\left(g_{n}^{*}\right)$. Assume $\eta(x)=1$ at all $x$. Then $L^{*}=0$ and

$$
\mathbf{E} L_{n} \geq \sum_{x} p_{x}\left(1-p_{x}\right)^{n}
$$

If $p_{x}=1 / k$ for all $x$, we have $\mathbf{E} L_{n} \geq(1-1 / k)^{n} \geq 1 / 2$ if $k \geq 2 n$. This simple calculation shows that we cannot say anything useful about fundamental rules unless $k \leq 2 n$ at the very least. On the positive side, the following universal bound is useful.

Theorem 27.1. For the fundamental rule, we have $L_{n} \rightarrow L^{*}$ with probability one as $n \rightarrow \infty$, and, in fact, for all distributions,

$$
\mathbf{E} L_{n} \leq L^{*}+\sqrt{\frac{k}{2(n+1)}}+\frac{k}{e n}
$$

and

$$
\mathbf{E} L_{n} \leq L^{*}+1.075 \sqrt{\frac{k}{n}}
$$

Proof. The first statement follows trivially from the strong universal consistency of histogram rules (see Theorem 9.4). It is the universal inequality that is of interest here. If $B(n, p)$ denotes a binomial random variable with parameters $n$ and $p$, then we have

$$
\begin{aligned}
\mathbf{E} L_{n}= & \sum_{x=1}^{k} p_{x}(\eta(x)+(1-2 \eta(x)) \mathbf{P}\{B(N(x), \eta(x))>N(x) / 2\}) \\
& \text { (here } N(x) \text { is a binomial }(n, p(x)) \text { random variable) } \\
= & \sum_{x=1}^{k} p_{x}(1-\eta(x)+(2 \eta(x)-1) \mathbf{P}\{B(N(x), \eta(x)) \leq N(x) / 2\}) .
\end{aligned}
$$

From this, if $\xi(x)=\min (\eta(x), 1-\eta(x))$,
$\mathbf{E} L_{n}$

$$
\begin{aligned}
& \leq \sum_{x=1}^{k} p_{x}(\xi(x)+(1-2 \xi(x)) \mathbf{P}\{B(N(x), \xi(x)) \geq N(x) / 2\}) \\
& \leq \quad L^{*}+\sum_{x=1}^{k} p_{x}(1-2 \xi(x)) \mathbf{E}\left\{\frac{N(x) \xi(x)(1-\xi(x))}{N(x) \xi(x)(1-2 \xi(x))+\left(\frac{1}{2}-\xi(x)\right)^{2} N(x)^{2}}\right\}
\end{aligned}
$$

(by the Chebyshev-Cantelli inequality-Theorem A.17)

$$
\leq \quad L^{*}+\sum_{x=1}^{k} p_{x}(1-2 \xi(x)) \mathbf{E}\left\{\frac{1}{1+(1-2 \xi(x))^{2} N(x)}\right\}
$$

$$
(\text { since } \xi(x)(1-\xi(x)) \leq 1 / 4)
$$

$$
\leq \quad L^{*}+\sum_{x=1}^{k} p_{x} \mathbf{E}\left\{\frac{1}{2 \sqrt{N(x)}} I_{\{N(x)>0\}}+(1-2 \xi(x)) I_{\{N(x)=0\}}\right\}
$$

$$
\text { (since the function } \left.u /\left(1+N u^{2}\right) \leq 1 /(2 \sqrt{N}) \text { for } 0 \leq u \leq 1\right)
$$

$$
\leq L^{*}+\sum_{x=1}^{k} p_{x}\left(1-p_{x}\right)^{n}+\frac{1}{2} \sum_{x=1}^{k} p_{x} \sqrt{\mathbf{E}\left\{\frac{1}{N(x)} I_{\{N(x)=0\}}\right\}}
$$

(by Jensen's inequality)

$$
\leq L^{*}+\left(1-\frac{1}{k}\right)^{n}+\frac{1}{2} \sum_{x=1}^{k} p_{x} \sqrt{\frac{2}{(n+1) p_{x}}}
$$

(by Lemma A. 2 and the fact that the worst distribution has $p_{x}=1 / k$ for all $x$ )

$$
\begin{aligned}
& \leq L^{*}+e^{-n / k}+\frac{1}{\sqrt{2(n+1)}} \sum_{x=1}^{k} \sqrt{p_{x}} \\
& \leq L^{*}+\frac{k}{e n}+\sqrt{\frac{k}{2(n+1)}}\left(\text { since } e^{-u} \leq 1 /(e u) \text { for } u \geq 0,\right.
\end{aligned}
$$

and by the Cauchy-Schwarz inequality).

This concludes the proof, as $\sqrt{1 / 2}+1 / e \leq 1.075$.

For several key properties of the fundamental rule, the reader is referred to Glick (1973) and to Problem 27.1. Other references include Krzanowski (1987), Goldstein (1977), and Goldstein and Dillon (1978). Note also the following extension of Theorem 27.1, which shows that the fundamental rule can handle all discrete distributions.

Theorem 27.2. If $X$ is purely atomic (with possibly infinitely many atoms), then $L_{n} \rightarrow L^{*}$ with probability one for the fundamental rule.

Proof. Number the atoms $1,2, \ldots$ Define $X^{\prime}=\min (X, k)$ and replace $\left(X_{i}, Y_{i}\right)$ by $\left(X_{i}^{\prime}, Y_{i}\right)$, where $X_{i}^{\prime}=\min \left(X_{i}, k\right)$. Apply the fundamental rule to the new problem and note that by Theorem 27.1, if $k$ is fixed, $L_{n} \rightarrow L^{*}$ with probability one for the new rule (and new distribution). However, the difference in $L_{n}$ 's and in $L^{*}$ 's between the truncated and nontruncated versions cannot be more than $\mathbf{P}\{X \geq k\}$, which may be made as small as desired by choice of $k$.

### 27.2 Quantization

Consider a fixed partition of space $\mathcal{R}^{d}$ into $k$ sets $\left\{A_{1}, \ldots, A_{k}\right\}$, and let $g_{n}$ be the standard partitioning rule based upon majority votes in the $A_{i}$ 's (ties are broken by favoring the response " 0, ," as elsewhere in the book). We consider two rules:
(1) The rule $g_{n}$ considered above; the data are $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, with $(X, Y) \in \mathcal{R}^{d} \times\{0,1\}$. The probability of error is denoted by $L_{n}$, and the Bayes probability of error is $L^{*}=\mathbf{E}\{\min (\eta(X), 1-\eta(X))\}$ with $\eta(x)=$ $\mathbf{P}\{Y=1 \mid X=x\}, x \in \mathcal{R}^{d}$.
(2) The fundamental rule $g_{n}^{\prime}$ operating on the quantized data $\left(X_{1}^{\prime}, Y_{1}\right), \ldots$, $\left(X_{n}^{\prime}, Y_{n}\right)$, with $\left(X^{\prime}, Y\right) \in\{1, \ldots, k\} \times\{0,1\}, X_{i}^{\prime}=j$ if $X \in A_{j}$. The Bayes probability of error is $L^{\prime *}=\mathbf{E}\left\{\min \left(\eta^{\prime}\left(X^{\prime}\right), 1-\eta^{\prime}\left(X^{\prime}\right)\right)\right\}$ with $\eta^{\prime}\left(x^{\prime}\right)=\mathbf{P}\{Y=$ $\left.1 \mid X^{\prime}=x^{\prime}\right\}, x^{\prime} \in\{1, \ldots, k\}$. The probability of error is denoted by $L_{n}^{\prime}$.

Clearly, $g_{n}^{\prime}$ is nothing but the fundamental rule for the quantized data. As $g_{n}(x)=$ $g_{n}^{\prime}\left(x^{\prime}\right)$, where $x^{\prime}=j$ if $x \in A_{j}$, we see that

$$
\begin{aligned}
L_{n} & =\mathbf{P}\left\{Y \neq g_{n}(X) \mid X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right\} \\
& =\mathbf{P}\left\{Y \neq g_{n}^{\prime}\left(X^{\prime}\right) \mid X_{1}^{\prime}, Y_{1}, \ldots, X_{n}^{\prime}, Y_{n}\right\} \\
& =L_{n}^{\prime}
\end{aligned}
$$

However, the Bayes error probabilities are different in the two situations. We claim, however, the following:

Theorem 27.3. For the standard partitioning rule,

$$
\mathbf{E}\left\{L_{n}\right\} \leq L^{*}+1.075 \sqrt{\frac{k}{n}}+\delta
$$

where $\delta=\mathbf{E}\left\{\left|\eta(X)-\eta^{\prime}\left(X^{\prime}\right)\right|\right\}$. Furthermore,

$$
L^{*} \leq L^{*} \leq L^{*}+\delta
$$

Proof. Clearly, $L^{*} \leq L^{* *}$ (see Problem 2.1). Also,

$$
\begin{aligned}
L^{\prime *} & =\mathbf{E}\left\{\min \left(\eta^{\prime}\left(X^{\prime}\right), 1-\eta^{\prime}\left(X^{\prime}\right)\right)\right\} \\
& \leq \mathbf{E}\left\{\min \left(\eta(X)+\left|\eta(X)-\eta^{\prime}\left(X^{\prime}\right)\right|, 1-\eta(X)+\left|\eta(X)-\eta^{\prime}\left(X^{\prime}\right)\right|\right)\right\} \\
& \leq \delta+\mathbf{E}\{\min (\eta(X), 1-\eta(X))\} \\
& =\delta+L^{*}
\end{aligned}
$$

Furthermore,

$$
\mathbf{E}\left\{L_{n}\right\} \leq L^{*}+1.075 \sqrt{\frac{k}{n}} \leq L^{*}+\delta+1.075 \sqrt{\frac{k}{n}}
$$

by Theorem 27.1.

As an immediate corollary, we show how to use the last bound to get useful distribution-free performance bounds.

Theorem 27.4. Let $\mathcal{F}$ be a class of distributions of $(X, Y)$ for which $X \in[0,1]^{d}$ with probability one, and for some constants $c>0,1 \geq \alpha>0$,

$$
|\eta(x)-\eta(z)| \leq c\|x-z\|^{\alpha}, \quad x, z \in \mathcal{R}^{d}
$$

Then, if we consider all cubic histogram rules $g_{n}$ (see Chapter 6 for a definition), we have

$$
\inf _{\text {cubic histogram rule } g_{n}} \mathbf{E}\left\{L\left(g_{n}\right)-L^{*}\right\} \leq \frac{a}{\sqrt{n}}+\frac{b}{n^{\frac{\alpha}{d+2 \alpha}}},
$$

where $a$ and $b$ are constants depending upon $c, \alpha$, and $d$ only (see the proof for explicit expressions).

Theorem 27.4 establishes the existence of rules that perform uniformly at rate $O\left(n^{-\alpha /(d+2 \alpha)}\right)$ over $\mathcal{F}$. Results like this have an impact on the number of data points required to guarantee a given performance for any $(X, Y) \in \mathcal{F}$.
Proof. Consider a cubic grid with cells of volume $h^{d}$. As the number of cells covering $[0,1]^{d}$ does not exceed $(1 / h+2)^{d}$, we apply Theorem 27.3 to obtain

$$
\mathbf{E}\left\{L\left(g_{n}\right)\right\} \leq L^{*}+1.075\left(2+\frac{1}{h}\right)^{d / 2}+\delta
$$

where

$$
\delta \leq \sup _{x}\left|\eta(x)-\eta^{\prime}\left(x^{\prime}\right)\right| \leq c \sup _{A_{i}} \sup _{x, z \in A_{i}}\|z-x\|^{\alpha} \leq c(h \sqrt{d})^{\alpha}
$$

The right-hand side is approximately maximal when

$$
h=\left(\frac{1.075 d}{2 \alpha(c \sqrt{d})^{\alpha} \sqrt{n}}\right)^{\frac{2}{d+2 \alpha}} \stackrel{\text { def }}{=} \frac{c^{\prime}}{n^{\frac{1}{d+2 \alpha}}}
$$

Resubstitution yields the result.

### 27.3 Independent Components

Let $X=\left(X^{(1)}, \ldots, X^{(d)}\right)$ have components that are conditionally independent, given $\{Y=1\}$, and also, given $\{Y=0\}$. Introduce the notation

$$
\begin{aligned}
p(i) & =\mathbf{P}\left\{X^{(i)}=1 \mid Y=1\right\}, \\
q(i) & =\mathbf{P}\left\{X^{(i)}=1 \mid Y=0\right\} \\
p & =\mathbf{P}\{Y=1\} .
\end{aligned}
$$

With $x=\left(x^{(1)}, \ldots, x^{(d)}\right) \in\{0,1\}^{d}$, we see that

$$
\eta(x)=\frac{\mathbf{P}\{Y=1, X=x\}}{\mathbf{P}\{X=x\}}=\frac{p \mathbf{P}\{X=x \mid Y=1\}}{p \mathbf{P}\{X=x \mid Y=1\}+(1-p) \mathbf{P}\{X=x \mid Y=0\}},
$$

and

$$
\begin{aligned}
& \mathbf{P}\{X=x \mid Y=1\}=\prod_{i=1}^{d} p(i)^{x^{(i)}}(1-p(i))^{1-x^{(i)}} \\
& \mathbf{P}\{X=x \mid Y=0\}=\prod_{i=1}^{d} q(i)^{x^{(i)}}(1-q(i))^{1-x^{(i)}}
\end{aligned}
$$

Simple consideration shows that the Bayes rule is given by

$$
g^{*}(x)=\left\{\begin{array}{cc}
1 & \text { if } p \prod_{i=1}^{d} p(i)^{x^{(i)}}(1-p(i))^{1-x^{(i)}} \\
& >(1-p) \prod_{i=1}^{d} q(i)^{x^{(i)}}(1-q(i))^{1-x^{(i)}} \\
0 & \text { otherwise. }
\end{array}\right.
$$

Taking logarithms, it is easy to see that this is equivalent to the following rule:

$$
g^{*}(x)= \begin{cases}1 & \text { if } \alpha_{0}+\sum_{i=1}^{d} \alpha_{i} x^{(i)}>0 \\ 0 & \text { otherwise }\end{cases}
$$

where

$$
\begin{aligned}
\alpha_{0} & =\log \left(\frac{p}{1-p}\right)+\sum_{i=1}^{d} \log \left(\frac{1-p(i)}{1-q(i)}\right) \\
\alpha_{i} & =\log \left(\frac{p(i)}{q(i)} \cdot \frac{1-q(i)}{1-p(i)}\right) \quad i=1, \ldots, d
\end{aligned}
$$

In other words, the Bayes classifier is linear! This beautiful fact was noted by Minsky (1961), Winder (1963), and Chow (1965).

Having identified the Bayes rule as a linear discrimination rule, we may apply the full force of the Vapnik-Chervonenkis theory. Let $g_{n}$ be the rule that minimizes the empirical error over the class of all linear discrimination rules. As the class
of linear halfspaces of $\mathcal{R}^{d}$ has VC dimension $d+1$ (see Corollary 13.1), we recall from Theorem 13.11 that for $g_{n}$,

$$
\begin{aligned}
\mathbf{P}\left\{L\left(g_{n}\right)-L^{*}>\epsilon\right\} & \leq 8 n^{d+1} e^{-n \epsilon^{2} / 128} \\
\mathbf{E}\left\{L\left(g_{n}\right)-L^{*}\right\} & \leq 16 \sqrt{\frac{(d+1) \log n+4}{2 n}} \quad \text { (by Corollary 12.1), } \\
\mathbf{P}\left\{L\left(g_{n}\right)-L^{*}>\epsilon\right\} & \leq 16(\sqrt{n} \epsilon)^{4096(d+1)} e^{-n \epsilon^{2} / 2}, n \epsilon^{2} \geq 64
\end{aligned}
$$

and

$$
\mathbf{E}\left\{L\left(g_{n}\right)-L^{*}\right\} \leq \frac{c}{\sqrt{n}},
$$

where

$$
c=16+\sqrt{10^{13}(d+1) \log \left(10^{12}(d+1)\right)} \quad(\text { Problem 12.10 })
$$

These bounds are useful (i.e., they tend to zero with $n$ ) if $d=o(n / \log n)$. In contrast, without the independence assumption, we have pointed out that no nontrivial guarantee can be given about $\mathbf{E}\left\{L\left(g_{n}\right)\right\}$ unless $2^{d}<n$. The independence assumption has led us out of the high-dimensional quagmire.

One may wish to attempt to estimate the $p(i)$ 's and $q(i)$ 's by $\widehat{p}(i)$ and $\widehat{q}(i)$ and to use these in the plug-in rule $g_{n}(x)$ that decides 1 if

$$
\widehat{p} \prod_{i=1}^{d} \widehat{p}(i)^{x^{(i)}}(1-\widehat{p}(i))^{1-x^{(i)}}>(1-\widehat{p}) \prod_{i=1}^{d} \widehat{q}(i)^{x^{(i)}}(1-\widehat{q}(i))^{1-x^{(i)}}
$$

where $\widehat{p}$ is the standard sample-based estimate of $p$. The maximum-likelihood estimate of $p(i)$ is given by

$$
\widehat{p}(i)=\frac{\sum_{j=1}^{n} I_{\left\{X_{j}^{(i)}=1, Y_{j}=1\right\}}}{\sum_{j=1}^{n} I_{\left\{Y_{j}=1\right\}}},
$$

while

$$
\widehat{q}(i)=\frac{\sum_{j=1}^{n} I_{\left\{X_{j}^{(i)}=1, Y_{j}=0\right\}}}{\sum_{j=1}^{n} I_{\left\{Y_{j}=0\right\}}}
$$

with $0 / 0$ equal to 0 (see Problem 27.6). Note that the plug-in rule too is linear, with

$$
g_{n}(x)= \begin{cases}1 & \text { if } \alpha_{0}+\sum_{i=1}^{d} \alpha_{i} x^{(i)}>0 \\ 0 & \text { otherwise }\end{cases}
$$

where

$$
\begin{aligned}
\alpha_{0} & =\log \left(\frac{\widehat{p}}{1-\widehat{p}}\right)+\sum_{i=1}^{d} \log \left(\frac{N_{01}(i)}{N_{01}(i)+N_{11}(i)} \cdot \frac{N_{00}(i)+N_{10}(i)}{N_{00}(i)}\right) \\
\alpha_{i} & =\log \left(\frac{N_{11}(i)}{N_{01}(i)} \cdot \frac{N_{00}(i)}{N_{10}(i)}\right) \quad i=1, \ldots, d
\end{aligned}
$$

and

$$
\begin{gathered}
N_{00}(i)=\sum_{j=1}^{n} I_{\left\{X_{j}^{(i)}=0, Y_{j}=0\right\}},
\end{gathered} N_{01}(i)=\sum_{j=1}^{n} I_{\left\{X_{j}^{(i)}=0, Y_{j}=1\right\}}, ~\left(N_{11}(i)=\sum_{j=1}^{n} I_{\left\{X_{j}^{(i)}=1, Y_{j}=1\right\}},\right.
$$

For all this, we refer to Warner, Toronto, Veasey, and Stephenson (1961) (see also McLachlan (1992)).

Use the inequality

$$
\mathbf{E}\left\{L\left(g_{n}\right)-L^{*}\right\} \leq 2 \mathbf{E}\left\{\left|\eta_{n}(X)-\eta(X)\right|\right\},
$$

where $\eta$ is as given in the text, and $\eta_{n}$ is as $\eta$ but with $p, p(i), q(i)$ replaced by $\widehat{p}, \widehat{p}(i), \widehat{q}(i)$, to establish consistency:

Theorem 27.5. For the plug-in rule, $L\left(g_{n}\right) \rightarrow L^{*}$ with probability one and $\mathbf{E}\left\{L\left(g_{n}\right)\right\}-L^{*} \rightarrow 0$, whenever the components are independent.

We refer to the Problem section for an evaluation of an upper bound for $\mathbf{E}\left\{L\left(g_{n}\right)\right.$ $\left.L^{*}\right\}$ (see Problem 27.7).

Linear discrimination on the hypercube has of course limited value. The world is full of examples that are not linearly separable. For example, on $\{0,1\}^{2}$, if $Y=X^{(1)}\left(1-X^{(2)}\right)+X^{(2)}\left(1-X^{(1)}\right)$ (so that $Y$ implements the boolean "xor" or "exclusive or" function), the problem is not linearly separable if all four possible values of $X$ have positive probability. However, the exclusive or function may be dealt with very nicely if one considers quadratic discriminants dealt with in a later section (27.5) on series methods.

### 27.4 Boolean Classifiers

By a boolean classification problem, we mean a pattern recognition problem on the hypercube for which $L^{*}=0$. This setup relates to the fact that if we consider $Y=1$ as a circuit failure and $Y=0$ as an operable circuit, then $Y$ is a deterministic function of the $X^{(i)}$ 's, which may be considered as gates or switches. In that case, $Y$ may be written as a boolean function of $X^{(1)}, \ldots, X^{(d)}$. We may limit boolean classifiers in various ways by partially specifying this function. For example, following Natarajan (1991), we first consider all monomials, that is, all functions $g:\{0,1\}^{d} \rightarrow\{0,1\}$ of the form

$$
x^{\left(i_{1}\right)} x^{\left(i_{2}\right)} \cdots x^{\left(i_{k}\right)}
$$

for some $k \leq d$ and some indices $1 \leq i_{1}<\cdots<i_{k} \leq d$ (note that algebraic multiplication corresponds to a boolean "and"). In such situations, one might attempt to minimize the empirical error. As we know that $g^{*}$ is also a monomial, it is clear that the minimal empirical error is zero. One such minimizing monomial is given by

$$
g_{n}\left(x^{(1)}, \ldots, x^{(d)}\right)=x^{\left(i_{1}\right)} \cdots x^{\left(i_{k}\right)}
$$

where

$$
\min _{1 \leq i \leq n} X_{i}^{\left(i_{1}\right)}=\cdots=\min _{1 \leq i \leq n} X_{i}^{\left(i_{k}\right)}=1, \quad \text { and } \min _{1 \leq i \leq n} X_{i}^{(j)}=0, \quad j \notin\left\{i_{1}, \ldots, i_{k}\right\}
$$

Thus, $g_{n}$ picks those components for which every data point has a " 1 ". Clearly, the empirical error is zero. The number of possible functions is $2^{d}$. Therefore, by Theorem 12.1,

$$
\mathbf{P}\left\{L\left(g_{n}\right)>\epsilon\right\} \leq 2^{d} e^{-n \epsilon},
$$

and

$$
\mathbf{E}\left\{L\left(g_{n}\right)\right\} \leq \frac{d+1}{n} .
$$

Here again, we have avoided the curse of dimensionality. For good performance, it suffices that $n$ be a bit larger than $d$, regardless of the distribution of the data!

Assume that we limit the complexity of a boolean classifier $g$ by requiring that $g$ must be written as an expression having at most $k$ operations "not," "and," or "or," with the $x^{(i)}$ 's as inputs. To avoid problems with precedence rules, we assume that any number of parentheses is allowed in the expression. One may visualize each expression as an expression tree, that is, a tree in which internal nodes represent operations and leaves represent operands (inputs). The number of such binary trees with $k$ internal nodes (and thus $k+1$ leaves) is given by the Catalan number

$$
\frac{1}{k+1}\binom{2 k}{k}
$$

(see, e.g., Kemp (1984)). Furthermore, we may associate any of the $x^{(i)}$ 's with the leaves (possibly preceded by "not") and "and" or "or" with each binary internal node, thus obtaining a total of not more than

$$
2^{k}(2 d)^{k+1} \frac{1}{k+1}\binom{2 k}{k}
$$

possible boolean functions of this kind. As $k \rightarrow \infty$, this bound is not more than

$$
(4 d)^{k+1} \frac{4^{k}}{\sqrt{\pi k^{3}}} \leq(16 d)^{k}
$$

for $k$ large enough. Again, for $k$ large enough,

$$
\mathbb{P}\left\{L\left(g_{n}\right)>\epsilon\right\} \leq(16 d)^{k} e^{-n \epsilon}
$$

and

$$
\mathbf{E}\left\{L\left(g_{n}\right)\right\} \leq \frac{k \log (16 d)+1}{n}
$$

Note that $k$ is much more important than $d$ in determining the sample size. For historic reasons, we mention that if $g$ is any boolean expression consisting of at most $k$ "not," "and," or "or" operations, then the number of such functions was shown by Pippenger (1977) not to exceed

$$
\left(\frac{16(d+k)^{2}}{k}\right)^{k}
$$

Pearl (1979) used Pippenger's estimate to obtain performance bounds such as the ones given above.

### 27.5 Series Methods for the Hypercube

It is interesting to note that we may write any function $\eta$ on the hypercube as a linear combination of Rademacher-Walsh polynomials

$$
\psi_{i}(x)= \begin{cases}1 & i=0 \\ 2 x^{(1)}-1 & i=1 \\ \vdots & \vdots \\ 2 x^{(d)}-1 & i=d \\ \left(2 x^{(1)}-1\right)\left(2 x^{(2)}-1\right) & i=d+1 \\ \vdots & \vdots \\ \left(2 x^{(d-1)}-1\right)\left(2 x^{(d)}-1\right) & i=d+1+\binom{d}{2} \\ \left(2 x^{(1)}-1\right)\left(2 x^{(2)}-1\right)\left(2 x^{(3)}-1\right) & i=d+2+\binom{d}{2} \\ \vdots & \vdots \\ \left(2 x^{(1)}-1\right) \cdots\left(2 x^{(d)}-1\right) & i=2^{d}-1 .\end{cases}
$$

We verify easily that

$$
\sum_{x} \psi_{i}(x) \psi_{j}(x)= \begin{cases}2^{d} & \text { if } i=j \\ 0 & \text { if } i \neq j\end{cases}
$$

so that the $\psi_{i}$ 's form an orthogonal system. Therefore, we may write

$$
\mu(\{x\}) \eta(x)=\sum_{i=0}^{2^{d}-1} a_{i} \psi_{i}(x)
$$

where

$$
a_{i}=\frac{1}{2^{d}} \sum_{x} \psi_{i}(x) \eta(x) \mu(\{x\})=\mathbf{E}\left\{\frac{\psi_{i}(X)}{2^{d}} I_{\{Y=1\}}\right\}
$$

and $\mu(\{x\})=\mathbf{P}\{X=x\}$. Also,

$$
\mu(\{x\})(1-\eta(x))=\sum_{i=0}^{2^{d}-1} b_{i} \psi_{i}(x)
$$

with

$$
b_{i}=\frac{1}{2^{d}} \sum_{x} \psi_{i}(x)(1-\eta(x)) \mu(\{x\})=\mathbf{E}\left\{\frac{\psi_{i}(X)}{2^{d}} I_{\{Y=0\}}\right\} .
$$

Sample-based estimates of $a_{i}$ and $b_{i}$ are

$$
\begin{aligned}
& \widehat{a}_{i}=\frac{1}{n} \sum_{j=1}^{n} \frac{\psi_{i}\left(X_{j}\right)}{2^{d}} I_{\left\{Y_{j}=1\right\}}, \\
& \widehat{b}_{i}=\frac{1}{n} \sum_{j=1}^{n} \frac{\psi_{i}\left(X_{j}\right)}{2^{d}} I_{\left\{Y_{j}=0\right\}} .
\end{aligned}
$$

The Bayes rule is given by

$$
g^{*}(x)= \begin{cases}1 & \text { if } \sum_{i=0}^{2^{d}-1}\left(a_{i}-b_{i}\right) \psi_{i}(x)>0 \\ 0 & \text { otherwise }\end{cases}
$$

Replacing $a_{i}-b_{i}$ formally by $\widehat{a_{i}}-\widehat{b_{i}}$ yields the plug-in rule. Observe that this is just a discrete version of the Fourier series rules discussed in Chapter 17. This rule requires the estimation of $2^{d}$ differences $a_{i}-b_{i}$. Therefore, we might as well have used the fundamental rule.

When our hand is forced by the dimension, we may wish to consider only rules in the class $\mathcal{C}$ given by

$$
g(x)= \begin{cases}1 & \text { if } \sum_{i=0}^{\binom{d}{0}+\binom{d}{1}+\cdots+\binom{d}{k}}\left(a_{i}^{\prime}-b_{i}^{\prime}\right) \psi_{i}(x)>0 \\ 0 & \text { otherwise },\end{cases}
$$

where $k \leq d$ is a positive integer, and $a_{0}^{\prime}, b_{0}^{\prime}, a_{1}^{\prime}, b_{1}^{\prime}, \ldots$ are arbitrary constants. We have seen that the VC dimension of this class is not more than $\binom{d}{0}+\binom{d}{1}+$ $\cdots+\binom{d}{k} \leq d^{k}+1$ (Theorem 13.9). Within this class, estimation errors of the order of $O\left(\sqrt{d^{k} \log n / n}\right)$ are thus possible if we minimize the empirical error (Theorem 13.12). This, in effect, forces us to take $k \ll \log _{d} n$. For larger $k$, pattern recognition is all but impossible.

As an interesting side note, observe that for a given parameter $k$, each member of $\mathcal{C}$ is a $k$-th degree polynomial in $x$.

Remark. Performance of the plug-in rule. Define the plug-in rule by

$$
g_{n}(x)= \begin{cases}1 & \text { if } \sum_{i=0}^{\binom{d}{0}+\binom{d}{1}+\cdots+\binom{d}{k}}\left(\widehat{a}_{i}-\widehat{b}_{i}\right) \psi_{i}(x)>0 \\ 0 & \text { otherwise } .\end{cases}
$$

As $n \rightarrow \infty$, by the law of large numbers, $g_{n}$ approaches $g_{\infty}$ :

$$
g_{\infty}(x)= \begin{cases}1 & \text { if } \sum_{i=0}^{\binom{d}{0}+\binom{d}{1}+\cdots+\binom{d}{k}}\left(a_{i}-b_{i}\right) \psi_{i}(x)>0 \\ 0 & \text { otherwise },\end{cases}
$$

where $a_{i}, b_{i}$ are as defined above. Interestingly, $g_{\infty}$ may be much worse than the best rule in $\mathcal{C}$ ! Consider the simple example for $d=2, k=1$ :

$$
\text { Table of } \eta(x):
$$

|  | $x^{(1)}=0$ | $x^{(1)}=1$ |
| :---: | :---: | :---: |
| $x^{(2)}=1$ | $1 / 8$ | $6 / 8$ |
| $x^{(2)}=0$ | $5 / 8$ | $1 / 8$ |

Table of $\mu(\{x\})=\mathbf{P}\{X=x\}$ :

|  | $x^{(1)}=0$ | $x^{(1)}=1$ |
| :---: | :---: | :---: |
| $x^{(2)}=1$ | $p$ | $\frac{5-10 p}{9}$ |
| $x^{(2)}=0$ | $\frac{4-8 p}{9}$ | $p$ |

A simple calculation shows that for any choice $p \in[0,1 / 2]$, either $g_{\infty} \equiv 1$ or $g_{\infty} \equiv 0$. We have $g_{\infty} \equiv 1$ if $p<10 / 47$, and $g_{\infty} \equiv 0$ otherwise. However, in obvious notation, $L(g \equiv 1)=(41 p+11) / 36, L(g \equiv 0)=(50-82 p) / 72$. The best constant rule is $g \equiv 1$ when $p<7 / 41$ and $g \equiv 0$ otherwise. For $7 / 41<p<10 / 47$, the plug-in rule does not even pick the best constant rule, let alone the best rule in $\mathcal{C}$ with $k=1$, which it was intended to pick.

This example highlights the danger of parametric rules or plug-in rules when applied to incorrect or incomplete models.

Remark. Historical notes. The Rademacher-Walsh expansion occurs frequently in switching theory, and was given in Duda and Hart (1973). The Bahadur-Lazarsfeld expansion (Bahadur (1961)) is similar in nature. Ito (1969) presents error bounds for discrimination based upon a $k$-term truncation of the series. Ott and Kronmal (1976) provide further statistical properties. The rules described here with $k$ defining the number of interactions are also obtained if we model $\mathbb{P}\{X=x \mid Y=1\}$ and $\mathbb{P}\{X=x \mid Y=1\}$ by functions of the form

$$
\exp \left(\sum_{i=0}^{\binom{d}{0}+\binom{d}{1}+\cdots+\binom{d}{k}} a_{i} \psi_{i}(x)\right) \text { and } \exp \left(\sum_{i=0}^{\binom{d}{0}+\binom{d}{1}+\cdots+\binom{d}{k}} b_{i} \psi_{i}(x)\right) .
$$

The latter model is called the log-linear model (see McLachlan (1992, section 7.3)).

### 27.6 Maximum Likelihood

The maximum likelihood method (see Chapter 15) should not be used for picking the best rule from a class $\mathcal{C}$ that is not guaranteed to include the Bayes rule. Perhaps
a simple example will suffice to make the point. Consider the following hypercube setting with $d=2$ :

$$
\begin{array}{cc||c|c} 
& & x^{(1)}=0 & x^{(1)}=1 \\
\eta(x): & & \\
& & x^{(2)}=1 & 0 \\
x^{(2)}=0 & 1 & 1 \\
\mu(\{x\}): & & x^{(1)}=0 & x^{(1)}=1 \\
\hline x^{(2)}=1 & p & q \\
\hline x^{(2)}=0 & r & s
\end{array}
$$

In this example, $L^{*}=0$. Apply maximum likelihood to a class $\mathcal{F}$ with two members $\left\{\eta_{A}, \eta_{B}\right\}$, where

$$
\begin{aligned}
& \eta_{A}(x)= \begin{cases}4 / 5 & \text { if } x^{(1)}=0 \\
1 / 5 & \text { if } x^{(1)}=1\end{cases} \\
& \eta_{B}(x)= \begin{cases}0 & \text { if } x^{(1)}=0 \\
1 & \text { if } x^{(1)}=1\end{cases}
\end{aligned}
$$

Then maximum likelihood won't even pick the best member from $\mathcal{F}$ ! To verify this, with $g_{A}(x)=I_{\left\{\eta_{A}(x)>1 / 2\right\}}, g_{B}(x)=I_{\left\{\eta_{B}(x)>1 / 2\right\}}$, we see that $L\left(g_{A}\right)=p+q$, $L\left(g_{B}\right)=r+s$. However, if $\eta_{\mathrm{ML}}$ is given by

$$
\eta_{\mathrm{ML}}=\underset{\eta^{\prime} \in\left\{\eta_{A}, \eta_{B}\right\}}{\arg \max }\left(\prod_{i: Y_{i}=1} \eta^{\prime}\left(X_{i}\right) \cdot \prod_{i: Y_{i}=0}\left(1-\eta^{\prime}\left(X_{i}\right)\right)\right),
$$

and if we write $N_{i j}$ for the number of data pairs ( $X_{k}, Y_{k}$ ) having $X_{k}=i, Y_{k}=j$, then $\eta_{\mathrm{ML}}=\eta_{\mathrm{A}}$ if

$$
\left(\frac{4}{5}\right)^{N_{01}+N_{10}}\left(\frac{1}{5}\right)^{N_{11}+N_{00}}> \begin{cases}0 & \text { if } N_{01}+N_{10}>0 \\ 1 & \text { if } N_{01}+N_{10}=0\end{cases}
$$

and $\eta_{\mathrm{ML}}=\eta_{B}$ otherwise. Equivalently, $\eta_{\mathrm{ML}}=\eta_{A}$ if and only if $N_{01}+N_{10}>0$. Apply the strong law of large numbers to note that $N_{01} / n \rightarrow r, N_{10} / n \rightarrow s$, $N_{00} / n \rightarrow p$, and $N_{11} / n \rightarrow q$ with probability one, as $n \rightarrow \infty$. Thus,

$$
\lim _{n \rightarrow \infty} \mathbb{P}\left\{\eta_{\mathrm{ML}}=\eta_{A}\right\}= \begin{cases}1 & \text { if } r+s>0 \\ 0 & \text { otherwise }\end{cases}
$$

Take $r+s=\epsilon$ very small, $p+q=1-\epsilon$. Then, for the maximum likelihood rule,

$$
\mathbf{E} L_{n} \rightarrow L\left(g_{A}\right)=1-\epsilon>\epsilon=\min \left(L\left(g_{A}\right), L\left(g_{B}\right)\right)
$$

However, when $\mathcal{F}$ contains the Bayes rule, maximum likelihood is consistent (see Theorem 15.1).

### 27.7 Kernel Methods

Sometimes, $d$ is so large with respect to $n$ that the atoms in the hypercube are sparsely populated. Some amount of smoothing may help under some circumstances. Consider a kernel $K$, and define the kernel rule

$$
g_{n}(x)= \begin{cases}1 & \text { if } \frac{1}{n} \sum_{i=1}^{n}\left(2 Y_{i}-1\right) K\left(\frac{\left\|x-X_{i}\right\|}{h}\right)>0 \\ 0 & \text { otherwise }\end{cases}
$$

where $\|x-z\|$ is just the Hamming distance (i.e., the number of disagreements between components of $x$ and $z$ ). With $K(u)=e^{-u^{2}}$, the rule above reduces to a rule given in Aitchison and Aitken (1976). In their paper, different $h$ 's are considered for the two classes, but we won't consider that distinction here. Observe that at $h=0$, we obtain the fundamental rule. As $h \rightarrow \infty$, we obtain a (degenerate) majority rule over the entire sample. The weight given to an observation $X_{i}$ decreases exponentially in $\left\|x-X_{i}\right\|^{2}$. For consistency, we merely need $h \rightarrow 0$ (the condition $n h^{d} \rightarrow \infty$ of Theorem 10.1 is no longer needed). And in fact, we even have consistency with $h \equiv 0$ as this yields the fundamental rule.

The data-based choice of $h$ has been the object of several papers, including Hall (1981) and Hall and Wand (1988). In the latter paper, a mean squared error criterion is minimized. We only mention the work of Tutz (1986; 1988; 1989), who picks $h$ so as to minimize the deleted estimate $L_{n}^{(D)}$.

Theorem 27.6. (TuTz (1986)). Let $H_{n}$ be the smoothing factor in the AitchisonAitken rule that minimizes $L_{n}^{(D)}$. Then the rule is weakly consistent for all distributions of $(X, Y)$ on the hypercube.

Proof. See Theorem 25.8.

## Problems and Exercises

Problem 27.1. The fundamental rule. Let $g_{n}^{*}$ be the fundamental rule on a finite set $\{1, \ldots, k\}$, and define $L_{n}=L\left(g_{n}^{*}\right)$. Let $g^{*}$ be the Bayes rule (with error probability $L^{*}$ ), and let

$$
\Delta=\inf _{x: \eta(x) \neq 1 / 2}\left(\frac{1}{2}-\min (\eta(x), 1-\eta(x))\right) .
$$

Let $L_{n}^{(R)}$ be the resubstitution error estimate (or apparent error rate). Show the following:
(1) $\mathbf{E}\left\{L_{n}^{(R)}\right\} \leq \mathbf{E} L_{n}$ (the apparent error rate is always optimistic; Hills (1966)).
(2) $\mathbf{E}\left\{L_{n}^{(R)}\right\} \leq L^{*}$ (the apparent error rate is in fact very optimistic; Glick (1973)).
(3) $\mathbf{E} L_{n} \leq L^{*}+e^{-2 n \Delta^{2}}$ (for a similar inequality related to Hellinger distances, see Glick (1973)). This is an exponential but distribution-dependent error rate.

Problem 27.2. Discrete Lipschitz classes. Consider the class of regression functions $\eta \in[0,1]$ with $|\eta(x)-\eta(z)| \leq c \rho(x, z)^{\alpha}$, where $x, z \in\{0,1\}^{d}, \rho(\cdot, \cdot)$ denotes the Hamming
distance, and $c>0$ and $\alpha>0$ are constants (note that $\alpha$ is not bounded from above). The purpose is to design a discrimination rule for which uniformly over all distributions of $(X, Y)$ on $\{0,1\}^{d} \times\{0,1\}$ with such $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$, we have

$$
\mathbf{E}\left\{L_{n}-L^{*}\right\} \leq \frac{\Psi(c, \alpha, d)}{\sqrt{n}}
$$

where the function $\Psi(c, \alpha, d)$ is as small as possible. Note: for $c=1$, the class contains all regression functions on the hypercube, and thus $\Psi(1, \alpha, d)=1.075 \cdot 2^{d / 2}$ (Theorem 27.1). How small should $c$ be to make $\Psi$ polynomial in $d$ ?

Problem 27.3. With a cubic histogram partition of $[0,1]^{d}$ into $k^{d}$ cells (of volume $1 / k^{d}$ each), we have, for the Lipschitz ( $c, \alpha$ ) class $\mathcal{F}$ of Theorem 27.4,

$$
\sup _{(X, Y) \in \mathcal{F}} \delta=c\left(\frac{\sqrt{d}}{k}\right)^{\alpha} .
$$

This grows as $d^{\alpha / 2}$. Can you define a partition into $k^{d}$ cells for which $\sup _{(X, Y) \in \mathcal{F}} \delta$ is smaller? Hint: Consult Conway and Sloane (1993).

Problem 27.4. Consider the following randomized histogram rule: $X_{1}, \ldots, X_{k}$ partition $[0,1]^{d}$ into polyhedra based on the nearest neighbor rule. Within each cell, we employ a majority rule based upon $X_{k+1}, \ldots, X_{n}$. If $X$ is uniform on $[0,1]^{d}$ and $\eta$ is Lipschitz ( $c, \alpha$ ) (as in Theorem 27.4), then can you derive an upper bound for $\mathbf{E}\left\{L_{n}-L^{*}\right\}$ as a function of $k, n, c, \alpha$, and $d$ ? How does your bound compare with the cubic histogram rule that uses the same number ( $k$ ) of cells?

Problem 27.5. Let $\mathcal{F}$ be the class of all Lipschitz $(c, \alpha)$ functions $\eta^{\prime} \in \mathcal{R}^{d} \rightarrow[0,1]$. Let $(X, Y) \in \mathcal{F}$ denote the fact that $(X, Y)$ has regression function $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$ in $\mathcal{F}$. Then, for any cubic histogram rule, show that

$$
\sup _{(X, Y) \in \mathcal{F}} \mathbf{E}\left\{L_{n}-L^{*}\right\} \geq \frac{1}{2}-L^{*}
$$

Thus, the compactness condition on the space is essential for the distribution-free error bound given in Theorem 27.4.

Problem 27.6. Independent model. Show that in the independent model, the maximum likelihood estimate $\widehat{p}(i)$ of $p(i)$ is given by

$$
\frac{\sum_{j=1}^{n} I_{\left\{X_{j}^{(i)}=1, Y_{j}=1\right\}}}{\sum_{j=1}^{n} I_{\left\{Y_{j}=1\right\}}}
$$

Problem 27.7. Independent model. For the plug-in rule in the independent model, is it true that $\mathbf{E}\left\{L_{n}-L^{*}\right\}=O(1 / \sqrt{n})$ uniformly over all pairs $(X, Y)$ on $\{0,1\}^{d} \times\{0,1\}$ ? If so, find a constant $c$ depending upon $d$ only, such that $\mathbf{E}\left\{L_{n}-L^{*}\right\} \leq c / \sqrt{n}$. If not, provide a counterexample.

Problem 27.8. Consider a hypercube problem in which $X=\left(X^{(1)}, \ldots, X^{(d)}\right)$ and each $X^{(i)} \in\{-1,0,1\}$ (a ternary generalization). Assume that the $X^{(i)}$ 's are independent but not
necessarily identically distributed. Show that there exists a quadratic Bayes rule, i.e., $g^{*}(x)$ is 1 on the set

$$
\alpha_{0}+\sum_{i=1}^{d} \alpha_{i} x^{(i)}+\sum_{i, j=1}^{d} \alpha_{i j} x^{(i)} x^{(j)}>0
$$

where $\alpha_{0},\left\{\alpha_{i}\right\}$, and $\left\{\alpha_{i j}\right\}$ are some weights (Kazmierczak and Steinbuch (1963)).
Problem 27.9. Let $\mathcal{A}$ be the class of all sets on the hypercube $\{0,1\}^{d}$ of the form $x^{\left(i_{1}\right)} \ldots$ $x^{\left(i_{k}\right)}=1$, where $\left(x^{(1)}, \ldots, x^{(d)}\right) \in\{0,1\}^{d}, 1 \leq i_{1}<\cdots<i_{k} \leq d$. (Thus, $\mathcal{A}$ is the class of all sets carved out by the monomials.) Show that the vc dimension of $\mathcal{C}$ is $d$. Hint: The set $\{(0,1,1, \ldots, 1),(1,0,1, \ldots, 1), \ldots,(1,1,1, \ldots, 0)\}$ is shattered by $\mathcal{A}$. No set of size $d+1$ can be shattered by $\mathcal{A}$ by the pigeonhole principle.

Problem 27.10. Show that the Catalan number

$$
\frac{1}{n+1}\binom{2 n}{n} \sim \frac{4^{n}}{\sqrt{\pi n^{3}}} .
$$

(See, e.g., Kemp (1984).)
Problem 27.11. Provide an argument to show that the number of boolean functions with at most $k$ operations "and" or "or" and $d$ operands of the form $x^{(i)}$ or $1-x^{(i)}, x^{(i)} \in\{0,1\}$, is not more than

$$
2(2 d)^{k+1} \frac{1}{k+1}\binom{2 k}{k}
$$

(this is $2^{k}$ times less than the bound given in the text).
Problem 27.12. Provide upper and lower bounds on the vc dimension of the class of sets $\mathcal{A}$ on the hypercube $\{0,1\}^{d}$ that can be described by a boolean expression with the $x^{(i)}$ 's or $1-x^{(i)}$ 's as operands and with at most $k$ operations "and" or "or."

Problem 27.13. Linear discrimination on the hypercube. Let $g_{n}$ be the rule that minimizes the empirical error $\widehat{L}_{n}(\phi)$ over all linear rules $\phi$, when the data are drawn from any distribution on $\{0,1\}^{d} \times\{0,1\}$. Let $L_{n}$ be its probability of error, and let $L$ be the minimum error probability over all linear rules. Show that for $\epsilon>0$,

$$
\mathbf{P}\left\{L_{n}-L>\epsilon\right\} \leq 4\binom{2^{d}}{d} e^{-n \epsilon^{2} / 2}
$$

Deduce that

$$
\mathbf{E}\left\{L_{n}-L\right\} \leq 2 \frac{\left.\sqrt{1+\log 4\left({\left(2^{d}\right)}_{d}\right.}\right)}{\sqrt{2 n}} \leq 2 \frac{d+1}{\sqrt{2 n}} .
$$

Compare this result with the general Vapnik-Chervonenkis bound for linear rules (Theorem 13.11) and deduce when the bound given above is better. Hint: Count the number of possible linear rules.

Problem 27.14. On the hypercube $\{0,1\}^{d}$, show that the kernel rule of Aitchison and Aitken (1976) is strongly consistent when $\lim _{n \rightarrow \infty} h=0$.

Problem 27.15. Pick $h$ in the kernel estimate by minimizing the resubstitution estimate $L_{n}^{(R)}$, and call it $H_{n}^{(R)}$. For $L_{n}^{(D)}$, we call it $H_{n}^{(D)}$. Assume that the kernel function is of the form $K(\|\cdot\| / h)$ with $K \geq 0, K(u) \downarrow 0$ as $u \uparrow \infty$. Let $\widehat{L}_{n}$ be the error estimate for the
kernel rule with one of these two choices. Is it possible to find a constant $c$, depending upon $d$ and $K$ only, such that

$$
\mathbf{E}\left\{\widehat{L}_{n}-\inf _{h \geq 0} L_{n}(h)\right\} \leq \frac{c}{\sqrt{n}} ?
$$

If so, give a proof. If not, provide a counterexample. Note: if the answer is positive, a minor corollary of this result is Tutz's theorem. However, an explicit constant $c$ may aid in determining appropriate sample sizes. It may also be minimized with respect to $K$.

Problem 27.16. (Simon (1991).) Construct a partition of the hypercube $\{0,1\}^{d}$ in the following manner, based upon a binary classification tree with perpendicular splits: Every node at level $i$ splits the subset according to $x^{(i)}=0$ or $x^{(i)}=1$, so that there are at most $d$ levels of nodes. (In practice, the most important component should be $x^{(1)}$.) For example, all possible partitions of $\{0,1\}^{2}$ obtainable with 2 cuts are


Assign to each internal node (each region) a class. Define the Horton-Strahler number $\xi$ of a tree as follows: if a tree has one node, then $\xi=0$. If the root of the tree has left and right subtrees with Horton-Strahler numbers $\xi_{1}$ and $\xi_{2}$, then set

$$
\xi=\max \left(\xi_{1}, \xi_{2}\right)+I_{\left\{\xi_{1}=\xi_{2}\right\}} .
$$

Let $\mathcal{C}$ be the class of classifiers $g$ described above with Horton-Strahler number $\leq \xi$.
(1) Let $S=\left\{x \in\{0,1\}^{d}:\|x\| \leq \xi\right\}$, where $\|\cdot\|$ denotes Hamming distance from the all-zero vector. Show that $S$ is shattered by the class of sets $\{g=1: g \in \mathcal{C}\}$.
(2) Show that $|S|=\sum_{i=0}^{\xi}\binom{d}{i}$.
(3) Conclude that the vC dimension of $\mathcal{C}$ is at least $\sum_{i=0}^{\xi}\binom{d}{i}$. (Simon has shown that the vc dimension of $\mathcal{C}$ is exactly this, but that proof is more involved.)
(4) Assuming $L^{*}=0$, obtain an upper bound for $\mathbf{E}\left\{L_{n}\right\}$ as a function of $\xi$ and $d$, where $L_{n}$ is the probability of error for the rule picked by minimizing the empirical error over $\mathcal{C}$.
(5) Interpret $\xi$ as the height of the largest complete binary tree that can be embedded in the classification tree.

## 28

## Epsilon Entropy and Totally Bounded Sets

### 28.1 Definitions

This chapter deals with discrimination rules that are picked from a certain class of classifiers by minimizing the empirical probability of error over a finite set of carefully selected rules. We begin with a class $\mathcal{F}$ of regression functions (i.e., a posteriori probability functions) $\eta: \mathcal{R}^{d} \rightarrow[0,1]$ from which $\eta_{n}$ will be picked by the data. The massiveness of $\mathcal{F}$ can be measured in many ways-the route followed here is suggested in the work of Kolmogorov and Tikhomirov (1961). We will depart from their work only in details. We suggest comparing the results here with those from Chapters 12 and 15.

Let $\mathcal{F}_{\epsilon}=\left\{\eta^{(1)}, \ldots, \eta^{(N)}\right\}$ be a finite collection of functions $\mathcal{R}^{d} \rightarrow[0,1]$ such that

$$
\mathcal{F} \subset \bigcup_{\eta^{\prime} \in \mathcal{F}_{\epsilon}} S_{\eta^{\prime}, \epsilon},
$$

where $S_{\eta^{\prime}, \epsilon}$ is the ball of all functions $\xi: \mathcal{R}^{d} \rightarrow[0,1]$ with

$$
\left\|\xi-\eta^{\prime}\right\|_{\infty}=\sup _{x}\left|\xi(x)-\eta^{\prime}(x)\right|<\epsilon
$$

In other words, for each $\eta^{\prime} \in \mathcal{F}$, there exists an $\eta^{(i)} \in \mathcal{F}_{\epsilon}$ with $\sup _{x} \mid \eta^{\prime}(x)-$ $\eta^{(i)}(x) \mid<\epsilon$. The fewer $\eta^{(i)}$ 's needed to cover $\mathcal{F}$, the smaller $\mathcal{F}$ is, in a certain sense. $\mathcal{F}_{\epsilon}$ is called an $\epsilon$-cover of $\mathcal{F}$. The minimal value of $\left|\mathcal{F}_{\epsilon}\right|$ over all $\epsilon$-covers is called the $\epsilon$-covering number $\left(\mathcal{N}_{\epsilon}\right)$. Following Kolmogorov and Tikhomirov (1961),

$$
\log _{2} \mathcal{N}_{\epsilon}
$$

is called the $\epsilon$-entropy of $\mathcal{F}$. We will also call it the metric entropy. A collection $\mathcal{F}$ is totally bounded if $\mathcal{N}_{\epsilon}<\infty$ for all $\epsilon>0$. It is with such classes that we are concerned in this chapter. The next section gives a few examples. In the following section, we define the skeleton estimate based upon picking the empirically best member from $\mathcal{F}_{\epsilon}$.


Figure 28.1. An $\epsilon$-cover of the unit square.

### 28.2 Examples: Totally Bounded Classes

The simple scalar parametric class $\mathcal{F}=\left\{e^{-\theta|x|}, x \in \mathcal{R} ; \theta>0\right\}$ is not totally bounded (Problem 28.1). This is due simply to the presence of an unrestricted scale factor. It would still fail to be totally bounded if we restricted $\theta$ to $[1, \infty)$ or $[0,1]$. However, if we force $\theta \in[0,1]$ and change the class $\mathcal{F}$ to have functions

$$
\eta^{\prime}(x)= \begin{cases}e^{-\theta|x|} & \text { if }|x| \leq 1 \\ 0 & \text { otherwise }\end{cases}
$$

then the class is totally bounded. While it is usually difficult to compute $\mathcal{N}_{\epsilon}$ exactly, it is often simple to obtain matching upper and lower bounds. Here is a simple argument. Take $\theta_{i}=2 \epsilon i, 0 \leq i \leq\lfloor 1 /(2 \epsilon)\rfloor$ and define $\theta^{*}=1$. Each of these $\theta$-values defines a function $\eta^{\prime}$. Collect these in $\mathcal{F}_{\epsilon}$. Note that with $\eta^{\prime} \in \mathcal{F}$, with parameter $\theta$, if $\widehat{\theta}$ is the nearest value among $\left\{\theta_{i}, 0 \leq i \leq\lfloor 1 /(2 \epsilon)\rfloor\right\} \cup\left\{\theta^{*}\right\}$, then $|\widehat{\theta}-\theta| \leq \epsilon$. But then

$$
\sup _{|x| \leq 1}\left|e^{-\theta|x|}-e^{-\widehat{\theta}|x|}\right| \leq|\theta-\widehat{\theta}| \leq \epsilon
$$

Hence $\mathcal{F}_{\epsilon}$ is an $\epsilon$-cover of $\mathcal{F}$ of cardinality $\lfloor 1 /(2 \epsilon)\rfloor+2$. We conclude that $\mathcal{F}$ is totally bounded, and that

$$
\mathcal{N}_{\epsilon} \leq 2+\lfloor 1 /(2 \epsilon)\rfloor .
$$

(See Problem 28.2 for a $d$-dimensional generalization.)

For a lower bound, we use once again an idea from Kolmogorov and Tikhomirov (1961). Let $\mathcal{G}_{\epsilon}=\left\{\eta^{(1)}, \ldots, \eta^{(M)}\right\} \subset \mathcal{F}$ be a subset with the property that for every $i \neq j, \sup _{x}\left|\eta^{(i)}(x)-\eta^{(j)}(x)\right| \geq \epsilon$. The set $\mathcal{G}_{\epsilon}$ is thus $\epsilon$-separated. The maximal cardinality $\epsilon$-separated set is called the $\epsilon$-packing number (or $\epsilon$-separation number) $\mathcal{M}_{\epsilon}$. It is easy to see that

$$
\mathcal{M}_{2 \epsilon} \leq \mathcal{N}_{\epsilon} \leq \mathcal{M}_{\epsilon}
$$

(Problem 28.3). With this in hand, we see that $\mathcal{G}_{\epsilon}$ may be constructed as follows for our example: Begin with $\theta_{0}=0$. Then define $\theta_{1}$ by $e^{-\theta_{1}}=1-\epsilon, e^{-\theta_{2}}=1-2 \epsilon$, etcetera, until $\theta_{i}>1$. It is clear that this way an $\epsilon$-separated set $\mathcal{G}_{\epsilon}$ may be constructed with $\left|\mathcal{G}_{\epsilon}\right|=\lfloor(1-1 / e) / \epsilon\rfloor+1$. Thus,

$$
\mathcal{N}_{\epsilon} \geq \mathcal{M}_{2 \epsilon} \geq\left\lfloor\frac{1-1 / e}{2 \epsilon}\right\rfloor+1 .
$$

The $\epsilon$-entropy of $\mathcal{F}$ grows as $\log _{2}(1 / \epsilon)$ as $\epsilon \downarrow 0$.
Consider next a larger class, not of a parametric nature: let $\mathcal{F}$ be a class of functions $\eta$ on $[0, \Delta]$ satisfying the Lipschitz condition

$$
\left|\eta(x)-\eta\left(x^{\prime}\right)\right| \leq c\left|x-x^{\prime}\right|
$$

and taking values on [0, 1]. Kolmogorov and Tikhomirov (1961) showed that if

$$
\epsilon \leq \min \left(\frac{1}{4}, \frac{1}{16 \Delta c}\right)
$$

then

$$
\begin{aligned}
\frac{\Delta c}{\epsilon}+\log _{2} \frac{1}{\epsilon}-3 & \leq \log _{2} \mathcal{M}_{2 \epsilon} \\
& \leq \log _{2} \mathcal{N}_{\epsilon} \\
& \leq \frac{\Delta c}{\epsilon}+\log _{2} \frac{1}{\epsilon}+3
\end{aligned}
$$

(see Problem 28.5). Observe that the metric entropy is exponentially larger than for the parametric class considered above. This has a major impact on sample sizes needed for similar performances (see the next sections).

Another class of functions of interest is that containing functions $\eta:[0,1] \rightarrow$ $[0,1]$ that are $s$-times differentiable and for which the $s$-th derivative $\eta^{(s)}$ satisfies a Hölder condition of order $\alpha$,

$$
\left|\eta^{(s)}(x)-\eta^{(s)}\left(x^{\prime}\right)\right| \leq c\left|x-x^{\prime}\right|^{\alpha}, \quad x, x^{\prime} \in[0,1]
$$

where $c$ is a constant. In that case, $\log _{2} \mathcal{M}_{\epsilon}$ and $\log _{2} \mathcal{N}_{\epsilon}$ are both $\Theta\left(\epsilon^{-1 /(s+\alpha)}\right)$ as $\epsilon \downarrow 0$, where $a_{n}=\Theta\left(b_{n}\right)$ means that $a_{n}=O\left(b_{n}\right)$ and $b_{n}=O\left(a_{n}\right)$. This result, also due to Kolmogorov and Tikhomirov (1961), establishes a continuum of rates of increase of the $\epsilon$-entropy. In $\mathcal{R}^{d}$, with functions $\eta:[0,1]^{d} \rightarrow[0,1]$, if the Hölder condition holds for all derivatives of order $s$, then $\log _{2} \mathcal{N}_{\epsilon}=\Theta\left(\epsilon^{-d /(s+\alpha)}\right)$.

Here we have an interesting interpretation of dimension: doubling the dimension roughly offsets the effect of doubling the number of derivatives (or the degree of smoothness) of the $\eta$ 's. Working with Lipschitz functions on $\mathcal{R}^{1}$ is roughly equivalent to working with functions on $\mathcal{R}^{25}$ for which all 24-th order derivatives are Lipschitz! As there are $25^{24}$ such derivatives, we note immediately how much we must pay for certain performances in high dimensions.

Let $\mathcal{F}$ be the class of all entire analytic functions $\eta:[0,2 \pi] \rightarrow[0,1]$ whose periodic continuation satisfies

$$
|\eta(z)| \leq c e^{\sigma|\Im(z)|}
$$

for some constants $c, \sigma$ ( $z$ is a complex variable, $\mathfrak{J}(z)$ is its imaginary part). For this class, we know that

$$
\log _{2} \mathcal{N}_{\epsilon} \sim(4\lfloor\sigma\rfloor+2) \log \frac{1}{\epsilon} \quad \text { as } \epsilon \downarrow 0
$$

(Kolmogorov and Tikhomirov (1961)). The class appears to be as small as our parametric class. See also Vitushkin (1961).

### 28.3 Skeleton Estimates

The members of $\mathcal{F}_{\epsilon}$ form a representative skeleton of $\mathcal{F}$. We assume that $\mathcal{F}_{\epsilon} \subset \mathcal{F}$ (this condition was not imposed in the definition of an $\epsilon$-cover). For each $\eta^{\prime} \in \mathcal{F}$, we define its discrimination rule as

$$
g(x)= \begin{cases}1 & \text { if } \eta^{\prime}(x)>1 / 2 \\ 0 & \text { otherwise }\end{cases}
$$

Thus, we will take the liberty of referring to $\eta^{\prime}$ as a rule. For each such $\eta^{\prime}$, we define the probability of error as usual:

$$
L\left(\eta^{\prime}\right)=\mathbf{P}\{g(X) \neq Y\}
$$

The empirical probability of error of $\eta^{\prime}$ is denoted by

$$
\widehat{L}_{n}\left(\eta^{\prime}\right)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{g\left(X_{i}\right) \neq \neq Y_{i}\right\}}
$$

We define the skeleton estimate $\eta_{n}$ by

$$
\eta_{n}=\underset{\eta^{\prime} \in \mathcal{F}_{\epsilon}}{\arg \min } \widehat{L}_{n}\left(\eta^{\prime}\right) .
$$

One of the best rules in $\mathcal{F}$ is denoted by $\eta^{*}$ :

$$
L\left(\eta^{*}\right) \leq L\left(\eta^{\prime}\right), \quad \eta^{\prime} \in \mathcal{F}
$$

The first objective, as in standard empirical risk minimization, is to ensure that $L\left(\eta_{n}\right)$ is close to $L\left(\eta^{*}\right)$. If the true a posteriori probability function $\eta$ is in $\mathcal{F}$ (recall that $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$ ), then it is clear that $L^{*}=L\left(\eta^{*}\right)$. It will be seen from the theorem below that under this condition, the skeleton estimate has nice consistency and rate-of-convergence properties. The result is distribution-free in the sense that no structure on the distribution of $X$ is assumed. Problems 6.9 and 28.9 show that convergence of $L\left(\eta_{n}\right)$ to $L\left(\eta^{*}\right)$ for all $\eta$ 's-that is, not only for those in $\mathcal{F}$--holds if $X$ has a density. In any case, because the skeleton estimate is selected from a finite deterministic set (that may be constructed before data are collected!), probability bounding is trivial: for all $\epsilon>0, \delta>0$, we have

$$
\begin{aligned}
& \mathbf{P}\left\{L\left(\eta_{n}\right)-\inf _{\eta^{\prime} \in \mathcal{F}_{\epsilon}} L\left(\eta^{\prime}\right)>\delta\right\} \\
& \leq \quad\left|\mathcal{F}_{\epsilon}\right| \sup _{\eta^{\prime} \in \mathcal{F}_{\epsilon}} \mathbf{P}\left\{\left|\widehat{L}_{n}\left(\eta^{\prime}\right)-L\left(\eta^{\prime}\right)\right|>\delta / 2\right\} \quad \text { (by Lemma 8.2) } \\
& \quad \leq 2\left|\mathcal{F}_{\epsilon}\right| e^{-n \delta^{2} / 2} \quad \text { (Hoeffding's inequality). }
\end{aligned}
$$

Theorem 28.1. Let $\mathcal{F}$ be a totally bounded class of functions $\eta^{\prime}: \mathcal{R}^{d} \rightarrow[0,1]$. There is a sequence $\left\{\epsilon_{n}>0\right\}$ and a sequence of skeletons $\mathcal{F}_{\epsilon_{n}} \subset \mathcal{F}$ such that if $\eta_{n}$ is the skeleton estimate drawn from $\mathcal{F}_{\epsilon_{n}}$, then

$$
L\left(\eta_{n}\right) \rightarrow L^{*} \text { with probability one },
$$

whenever the true regression function $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$ is in $\mathcal{F}$.
It suffices to take $\mathcal{F}_{\epsilon}$ as an $\epsilon$-cover of $\mathcal{F}$ (note that $\left|\mathcal{F}_{\epsilon}\right|$ need not equal the $\epsilon$-covering number $\mathcal{N}_{\epsilon}$ ), and to define $\epsilon_{n}$ as the smallest positive number for which

$$
2 n \geq \frac{\log \left|\mathcal{F}_{\epsilon_{n}}\right|}{\epsilon_{n}^{2}}
$$

Finally, with $\epsilon_{n}$ picked in this manner,

$$
\mathbf{E}\left\{L\left(\eta_{n}\right)-L^{*}\right\} \leq(2+\sqrt{8}) \epsilon_{n}+\sqrt{\frac{\pi}{n}}
$$

Proof. We note that $\inf _{\eta^{\prime} \in \mathcal{F}_{\epsilon}} L\left(\eta^{\prime}\right) \leq L^{*}+2 \epsilon$, because if $\eta^{\prime} \in S_{\eta, \epsilon}$, then

$$
\mathbf{E}\left\{\left|\eta^{\prime}(X)-\eta(X)\right|\right\} \leq \sup _{x}\left|\eta^{\prime}(x)-\eta(x)\right| \leq \epsilon
$$

and thus, by Theorem $2.2, L\left(\eta^{\prime}\right)-L^{*} \leq 2 \epsilon$. Then for any $\delta \geq 2 \epsilon_{n}$,

$$
\begin{aligned}
\mathbf{P}\left\{L\left(\eta_{n}\right)-L^{*}>\delta\right\} & \leq \mathbf{P}\left\{L\left(\eta_{n}\right)-\inf _{\eta^{\prime} \in \mathcal{F}_{\epsilon_{n}}} L\left(\eta^{\prime}\right)>\delta-2 \epsilon_{n}\right\} \\
& \leq 2\left|\mathcal{F}_{\epsilon_{n}}\right| e^{-n\left(\delta-2 \epsilon_{n}\right)^{2} / 2} \quad \text { (see above) } \\
& \leq 2 e^{2 n \epsilon_{n}^{2}-n\left(\delta-2 \epsilon_{n}\right)^{2} / 2}
\end{aligned}
$$

which is summable in $n$, as $\epsilon_{n}=o(1)$. This shows the first part of the theorem. For the second part, we have

$$
\begin{aligned}
& \mathbf{E}\left\{L\left(\eta_{n}\right)\right\}-L^{*} \\
& \leq 2 \epsilon_{n}+\mathbf{E}\left\{L\left(\eta_{n}\right)\right\}-\inf _{\eta^{\prime} \in \mathcal{F}_{\epsilon_{n}}} L\left(\eta^{\prime}\right) \\
& \leq 2 \epsilon_{n}+\int_{0}^{\infty} \min \left(2 e^{2 n \epsilon_{n}^{2}-n t^{2} / 2}, 1\right) d t \\
& \leq(2+\sqrt{8}) \epsilon_{n}+2 \int_{\sqrt{8} \epsilon_{n}}^{\infty} e^{2 n \epsilon_{n}^{2}-n t^{2} / 2} d t \\
& \leq(2+\sqrt{8}) \epsilon_{n}+\int_{0}^{\infty} e^{-n t^{2} / 4} d t \\
&\left(\text { since } \epsilon_{n}^{2}-t^{2} / 4 \leq-t^{2} / 8 \text { for } t \geq \sqrt{8} \epsilon_{n}\right) \\
&=(2+\sqrt{8}) \epsilon_{n}+\sqrt{\frac{\pi}{n}} .
\end{aligned}
$$

The proof is completed.

Observe that the estimate $\eta_{n}$ is picked from a deterministic class. This, of course, requires quite a bit of preparation and knowledge on behalf of the user. Knowledge of the $\epsilon$-entropy (or at least an upper bound) is absolutely essential. Furthermore, one must be able to construct $\mathcal{F}_{\epsilon}$. This is certainly not computationally simple. Skeleton estimates should therefore be mainly of theoretical importance. They may be used, for example, to establish the existence of estimates with a guaranteed error bound as given in Theorem 28.1. A similar idea in nonparametric density estimation was proposed and worked out in Devroye (1987).

REMARK. In the first step of the proof, we used the inequality

$$
\mathbf{E}\left\{\left|\eta^{\prime}(X)-\eta(X)\right|\right\} \leq \sup _{x}\left|\eta^{\prime}(x)-\eta(x)\right|<\epsilon .
$$

It is clear from this that what we really need is not an $\epsilon$-cover of $\mathcal{F}$ with respect to the supremum norm, but rather, with respect to the $L_{1}(\mu)$ norm. In other words, the skeleton estimate works equally well if the skeleton is an $\epsilon$-cover of $\mathcal{F}$ with respect to the $L_{1}(\mu)$ norm, that is, it is a list of finitely many candidates with the property that for each $\eta^{\prime} \in \mathcal{F}$ there exists an $\eta^{(i)}$ in the list such that $\mathbf{E}\left\{\mid \eta^{\prime}(X)-\right.$ $\left.\eta^{(i)}(X) \mid\right\}<\epsilon$. It follows from the inequality above that the smallest such covering has fewer elements than that of any $\epsilon$-cover of $\mathcal{F}$ with respect to the supremum norm. Therefore, estimates based on such skeletons perform better. In fact, the difference may be essential. As an example, consider the class $\mathcal{F}$ of all regression functions on $[0,1]$ that are monotone increasing. For $\epsilon<1 / 2$, this class does not have a finite $\epsilon$-cover with respect to the supremum norm. However, for any $\mu$ it is possible to find an $\epsilon$-cover of $\mathcal{F}$, with respect to $L_{1}(\mu)$, with not more
than $4^{\lceil 1 / \epsilon\rceil}$ elements (Problem 28.6). Unfortunately, an $\epsilon$-cover with respect to $L_{1}(\mu)$ depends on $\mu$, the distribution of $X$. Since $\mu$ is not known in advance, we cannot construct this better skeleton. However, in some cases, we may have some a priori information about $\mu$. For example, if we know that $\mu$ is a member of a known class of distributions, then we may be able to construct a skeleton that is an $\epsilon$-cover for all measures in the class. In the above example, if we know that $\mu$ has a density bounded by a known number, then there is a finite skeleton with this property (see Problem 28.7). We note here that the basic idea behind the empirical covering method of Buescher and Kumar (1996a), described in Problem 12.14 , is to find a good skeleton based on a fraction of the data. Investigating this question further, we notice that even covering in $L_{1}(\mu)$ is more than what we really need. From the proof of Theorem 28.1, we see that all we need is a skeleton $\mathcal{F}_{\epsilon}$ such that $\inf _{\eta^{\prime} \in \mathcal{F}_{\epsilon}} L\left(\eta^{\prime}\right) \leq L^{*}+\epsilon$ for all $\eta \in \mathcal{F}$. Staying with the example of the class of monotonically increasing $\eta$ 's, we see that we may take in $\mathcal{F}_{\epsilon}$ the functions $\eta^{(i)}(x)=I_{\left\{x \geq q_{i}\right\}}$, where $q_{i}$ is the $i$-th $\epsilon$-quantile of $\mu$, that is, $q_{i}$ is the smallest number $z$ such that $\mathbf{P}\{X \leq z\} \geq i / \epsilon$. This collection of functions forms a skeleton in the required sense with about $(1 / \epsilon)$ elements, instead of the $4^{[1 / \epsilon\rceil}$ obtained by covering in $L_{1}(\mu)$, a significant improvement. Problem 28.8 illustrates another application of this idea. For more work on this we refer to Vapnik (1982), Benedek and Itai (1988), Kulkarni (1991), Dudley, Kulkarni, Richardson, and Zeitouni (1994), and Buescher and Kumar (1996a).

### 28.4 Rate of Convergence

In this section, we take a closer look at the distribution-free upper bound

$$
\mathbb{E}\left\{L\left(\eta_{n}\right)-L^{*}\right\} \leq(2+\sqrt{8}) \epsilon_{n}+\sqrt{\frac{\pi}{n}}
$$

For typical parametric classes (such as the one discussed in a Section 28.2), we have

$$
\mathcal{N}_{\epsilon}=\Theta\left(\frac{1}{\epsilon}\right) .
$$

If we take $\left|\mathcal{F}_{\epsilon}\right|$ close enough to $\mathcal{N}_{\epsilon}$, then $\epsilon_{n}$ is the solution of

$$
\frac{\log \mathcal{N}_{\epsilon}}{\epsilon^{2}} \cong n
$$

or $\epsilon_{n}=\Theta(\log n / \sqrt{n})$, and we achieve a guaranteed rate of $O(\log n / \sqrt{n})$. The same is true for the example of the class of analytic functions discussed earlier.

The situation is different for massive classes such as the Lipschitz functions on $[0,1]^{d}$. Recalling that $\log \mathcal{N}_{\epsilon}=\Theta\left(1 / \epsilon^{d}\right)$ as $\epsilon \downarrow 0$, we note that $\epsilon_{n}=\Theta\left(n^{-1 /(2+d)}\right)$. For this class, we have

$$
\mathbb{E}\left\{L\left(\eta_{n}\right)-L^{*}\right\}=O\left(n^{-\frac{1}{2+d}}\right)
$$

Here, once again, we encounter the phenomenon called the "curse of dimensionality." In order to achieve the performance $\mathbf{E}\left\{L\left(\eta_{n}\right)-L^{*}\right\} \leq \epsilon$, we need a sample of size $n \geq(1 / \epsilon)^{2+d}$, exponentially large in $d$. Note that the class of classifiers defined by this class of functions has infinite vc dimension. The skeleton estimates thus provide a vehicle for dealing with very large classes.

Finally, if we take all $\eta$ on $[0,1]$ for which $s$ derivatives exist, and $\eta^{(s)}$ is Lipschitz with a given constant $c$, similar considerations show that the rate of convergence is $O\left(n^{-(s+1) /(2 s+3)}\right)$, which ranges from $O\left(n^{-1 / 3}\right)$ (at $\left.s=0\right)$ to $O\left(n^{-1 / 2}\right)$ (as $s \rightarrow$ $\infty$ ). As the class becomes smaller, we can guarantee better rates of convergence. Of course, this requires more a priori knowledge about the true $\eta$.

We also note that if $\log \mathcal{N}_{\epsilon} / \epsilon^{2}=2 n$ in Theorem 28.1, then the bound is

$$
\sqrt{\frac{\pi}{n}}+(2+\sqrt{8}) \epsilon_{n}=\sqrt{\frac{\pi}{n}}+(2+\sqrt{8}) \sqrt{\frac{\log \mathcal{N}_{\epsilon_{n}}}{2 n}} .
$$

The error grows only sub-logarithmically in the size of the skeleton set. It grows as the square root of the $\epsilon$-entropy. Roughly speaking (and ignoring the dependence of $\epsilon_{n}$ upon $n$ ), we may say that for the same performance guarantees, doubling the $\epsilon$-entropy implies that we should double the sample size (to keep $\log \mathcal{N}_{\epsilon} / n$ constant). When referring to $\epsilon$-entropy, it is important to keep this sample size interpretation in mind.

## Problems and Exercises

Problem 28.1. Show that the class of functions $e^{-\theta \mid x i}$ on the real line, with parameter $\theta>0$, is not totally bounded.

Problem 28.2. Compute a good upper bound for $\mathcal{N}_{\epsilon}$ as a function of $d$ and $\epsilon$ for the class $\mathcal{F}$ of all functions on $\mathcal{R}^{d}$ given by

$$
\eta^{\prime}(x)= \begin{cases}e^{-\theta\|x\|} & \text { if }\|x\| \leq 1 \\ 0 & \text { otherwise },\end{cases}
$$

where $\theta \in[0,1]$ is a parameter. Repeat this question if $\theta_{1}, \ldots, \theta_{d}$ are in $[0,1]$ and

$$
\eta^{\prime}(x)= \begin{cases}e^{-\sqrt{\sum_{i=1}^{d} \theta_{i} x_{i}^{2}}} & \text { if } \max _{1 \leq i \leq d}\left|x_{i}\right| \leq 1 \\ 0 & \text { otherwise. }\end{cases}
$$

Hint: Both answers are polynomial in $1 / \epsilon$ as $\epsilon \downarrow 0$.
Problem 28.3. Show that $\mathcal{M}_{2 \epsilon} \leq \mathcal{N}_{\epsilon} \leq \mathcal{M}_{\epsilon}$ for any totally bounded set $\mathcal{F}$ (Kolmogorov and Tikhomirov (1961)).

Problem 28.4. Find a class $\mathcal{F}$ of functions $\eta^{\prime}: \mathcal{R}^{d} \rightarrow[0,1]$ such that
(a) for every $\epsilon>0$, it has a finite $\epsilon$-cover;
(b) the vc dimension of $\mathcal{A}=\left\{\left\{x: \eta^{\prime}(x)>1 / 2\right\} ; \eta^{\prime} \in \mathcal{F}\right\}$ is infinite.

Problem 28.5. Show that if $\mathcal{F}=\left\{\eta^{\prime}:[0, \Delta] \rightarrow[0,1], \eta^{\prime}\right.$ is Lipschitz with constant $\left.c\right\}$, then for $\epsilon$ small enough, $\log _{2} \mathcal{N}_{\epsilon} \leq A / \epsilon$, where $A$ is a constant depending upon $\Delta$ and $c$. (This is not as precise as the statement in the text obtained by Kolmogorov and Tikhomirov, but it will give you excellent practice.)

Problem 28.6. Obtain an estimate for the cardinality of the smallest $\epsilon$-cover with respect to the $L_{1}(\mu)$ norm for the class of $\eta$ 's on $[0,1]$ that are increasing. In particular, show that for any $\mu$ it is possible to find an $\epsilon$-cover with $4^{[1 / \epsilon]}$ elements. Can you do something similar for $\eta$ 's on $[0,1]^{d}$ that are increasing in each coordinate?

Problem 28.7. Consider the class of $\eta$ 's on [ 0,1$]$ that are increasing. Show that for every $\epsilon>0$, there is a finite list $\eta^{(1)}, \ldots \eta^{(N)}$ such that for all $\eta$ in the class,

$$
\inf _{i \leq N} \mathbf{E}\left\{\left|\eta(X)-\eta^{(i)}(X)\right|\right\}<\epsilon
$$

whenever $X$ has a density bounded by $B<\infty$. Estimate the smallest such $N$.
Problem 28.8. Assume that $X$ has a bounded density on $[0,1]^{2}$, and that $\eta$ is monotonically increasing in both coordinates. (This is a reasonable assumption in many applications.) Then the set $\left\{x: g^{*}(x)=0\right\}$ is a monotone layer. Consider the following classification rule: take a $k \times k$ grid in $[0,1]^{2}$, and minimize the empirical error over all classifiers $\phi$ such that $\{x: \phi(x)=0\}$ is a monotone layer, and it is a union of cells in the $k \times k$ grid. What is the optimal choice of $k$ ? Obtain an upper bound for $L\left(g_{n}\right)-L^{*}$. Compare your result with that obtained for empirical error minimization over the class of all monotone layers in Section 13.4. Hint: Count the number of different classifiers in the class. Use Hoeffding's inequality and the union-of-events bound for the estimation error. Bound the approximation error using the bounded density assumption.

Problem 28.9. Apply Problem 6.9 to extend the consistency result in Theorem 28.1 as follows. Let $\mathcal{F}$ be a totally bounded class of functions $\eta^{\prime}: \mathcal{R}^{d} \rightarrow[0,1]$ such that $\lambda(\{x:$ $\left.\left.\eta^{\prime}(x)=1 / 2\right\}\right)=0$ for each $\eta^{\prime} \in \mathcal{F}\left(\lambda\right.$ is the Lebesgue measure on $\left.\mathcal{R}^{d}\right)$. Show that there is a sequence $\left\{\epsilon_{n}>0\right\}$ and a sequence of skeletons $\mathcal{F}_{\epsilon_{n}}$ such that if $\eta_{n}$ is the skeleton estimate drawn from $\mathcal{F}_{\epsilon_{n}}$, then $L\left(\eta_{n}\right) \rightarrow \inf _{\eta^{\prime} \in \mathcal{F}} L\left(\eta^{\prime}\right)$ with probability one, whenever $X$ has a density. In particular, $L\left(\eta_{n}\right) \rightarrow L^{*}$ with probability one if the Bayes rule takes the form $g^{*}(x)=I_{\left\{\eta^{\prime}(x)>1 / 2\right\}}$ for some $\eta^{\prime} \in \mathcal{F}$. Note: the true regression function $\eta$ is not required to be in $\mathcal{F}$.

## 29

## Uniform Laws of Large Numbers

### 29.1 Minimizing the Empirical Squared Error

In Chapter 28 the data $D_{n}$ were used to select a function $\eta_{n}$ from a class $\mathcal{F}$ of candidate regression functions $\eta^{\prime}: \mathcal{R}^{d} \rightarrow[0,1]$. The corresponding classification rule $g_{n}$ is $I_{\left\{\eta_{n}>1 / 2\right\}}$. Selecting $\eta_{n}$ was done in two steps: a skeleton-an $\epsilon$-covering-of $\mathcal{F}$ was formed, and the empirical error count was minimized over the skeleton. This method is computationally cumbersome. It is tempting to use some other empirical quantity to select a classifier. Perhaps the most popular among these measures is the empirical squared error:

$$
e_{n}\left(\eta^{\prime}\right)=\frac{1}{n} \sum_{i=1}^{n}\left(\eta^{\prime}\left(X_{i}\right)-Y_{i}\right)^{2} .
$$

Assume now that the function $\eta_{n}$ is selected by minimizing the empirical squared error over $\mathcal{F}$, that is,

$$
e_{n}\left(\eta_{n}\right) \leq e_{n}\left(\eta^{\prime}\right), \quad \eta^{\prime} \in \mathcal{F}
$$

As always, we are interested in the error probability,

$$
L\left(\eta_{n}\right)=\mathbf{P}\left\{I_{\left\{\eta_{n}(x)>1 / 2\right\}} \neq Y \mid D_{n}\right\},
$$

of the resulting classifier. If the true regression function $\eta(x)=\mathbb{P}\{Y=1 \mid X=x\}$ is not in the class $\mathcal{F}$, then it is easy to see that empirical squared error minimization may fail miserably (see Problem 29.1). However, if $\eta \in \mathcal{F}$, then for every $\eta^{\prime} \in \mathcal{F}$
we have

$$
\begin{aligned}
L\left(\eta^{\prime}\right)-L^{*} & \leq 2 \sqrt{\mathbf{E}\left\{\left(\eta^{\prime}(X)-\eta(X)\right)^{2}\right\}} \quad \text { (by Corollary 6.2) } \\
& =2 \sqrt{\mathbf{E}\left\{\left(\eta^{\prime}(X)-Y\right)^{2}\right\}-\mathbf{E}\left\{(\eta(X)-Y)^{2}\right\}} \\
& =2 \sqrt{\mathbf{E}\left\{\left(\eta^{\prime}(X)-Y\right)^{2}\right\}-\inf _{\tilde{\eta} \in \mathcal{F}} \mathbf{E}\left\{(\tilde{\eta}(X)-Y)^{2}\right\}}
\end{aligned}
$$

where the two equalities follow from the fact that $\eta(X)=\mathbf{E}\{Y \mid X\}$. Thus, we have

$$
\begin{align*}
& L\left(\eta_{n}\right)-L^{*} \\
& \quad \leq 2 \sqrt{\mathbf{E}\left\{\left(\eta_{n}(X)-Y\right)^{2} \mid D_{n}\right\}-\inf _{\eta^{\prime} \in \mathcal{F}} \mathbf{E}\left\{\left(\eta^{\prime}(X)-Y\right)^{2}\right\}} \\
& \quad \leq 2 \sqrt{2 \sup _{\eta^{\prime} \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n}\left(\eta^{\prime}\left(X_{i}\right)-Y_{i}\right)^{2}-\mathbf{E}\left\{\left(\eta^{\prime}(X)-Y\right)^{2}\right\}\right|} \tag{29.1}
\end{align*}
$$

by an argument as in the proof of Lemma 8.2. Thus, the method is consistent if the supremum above converges to zero. If we define $Z_{i}=\left(X_{i}, Y_{i}\right)$ and $f\left(Z_{i}\right)=$ $\left(\eta^{\prime}\left(X_{i}\right)-Y_{i}\right)^{2}$, then we see that we need only to bound

$$
\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} f\left(Z_{i}\right)-\mathbf{E}\left\{f\left(Z_{1}\right)\right\}\right|
$$

where $\mathcal{F}$ is a class of bounded functions. In the next four sections we develop upper bounds for such uniform deviations of averages from their expectations. Then we apply these techniques to establish consistency of generalized linear classifiers obtained by minimization of the empirical squared error.

### 29.2 Uniform Deviations of Averages from Expectations

Let $\mathcal{F}$ be a class of real-valued functions defined on $\mathcal{R}^{d}$, and let $Z_{1}, \ldots, Z_{n}$ be i.i.d. $\mathcal{R}^{d}$-valued random variables. We assume that for each $f \in \mathcal{F}, 0 \leq f(x) \leq M$ for all $x \in \mathcal{R}^{d}$ and some $M<\infty$. By Hoeffding's inequality,

$$
\mathbf{P}\left\{\left|\frac{1}{n} \sum_{i=1}^{n} f\left(Z_{i}\right)-\mathbf{E}\left\{f\left(Z_{1}\right)\right\}\right|>\epsilon\right\} \leq 2 e^{-2 n \epsilon^{2} / M^{2}}
$$

for any $f \in \mathcal{F}$. However, it is much less trivial to obtain information about the probabilities

$$
\mathbf{P}\left\{\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} f\left(Z_{i}\right)-\mathbf{E}\left\{f\left(Z_{1}\right)\right\}\right|>\epsilon\right\}
$$

Vapnik and Chervonenkis (1981) were the first to obtain bounds for the probability above. For example, the following simple observation makes Theorems 12.5, 12.8, and 12.10 easy to apply in the new situation:

## Lemma 29.1.

$$
\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} f\left(Z_{i}\right)-\mathbf{E}\{f(Z)\}\right| \leq M \sup _{f \in \mathcal{F}, t>0}\left|\frac{1}{n} \sum_{i=1}^{n} I_{\left\{f\left(Z_{i}\right)>t\right\}}-\mathbf{P}\{f(Z)>t\}\right| .
$$

Proof. Exploiting the identity $\int_{0}^{\infty} \mathbf{P}\{X>t\} d t=\mathbf{E} X$ for nonnegative random variables, we have

$$
\begin{aligned}
& \sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} f\left(Z_{i}\right)-\mathbf{E}\{f(Z)\}\right| \\
& \quad=\sup _{f \in \mathcal{F}}\left|\int_{0}^{\infty}\left(\frac{1}{n} \sum_{i=1}^{n} I_{\left\{f\left(Z_{i}\right)>t\right\}}-\mathbf{P}\{f(Z)>t\}\right) d t\right| \\
& \quad \leq M \sup _{f \in \mathcal{F}, t>0}\left|\frac{1}{n} \sum_{i=1}^{n} I_{\left\{f\left(Z_{i}\right)>t\right\}}-\mathbf{P}\{f(Z)>t\}\right|
\end{aligned}
$$

For example, from Theorem 12.5 and Lemma 29.1 we get
Corollary 29.1. Define the collection of sets

$$
\widehat{\mathcal{F}}=\left\{A_{f, t}: f \in \mathcal{F}, t \in[0, M]\right\}
$$

where for every $f \in \mathcal{F}$ and $t \in[0, M]$ the set $A_{f, t} \in \mathcal{R}^{d}$ is defined as

$$
A_{f, t}=\{z: f(z)>t\}
$$

Then

$$
\mathbf{P}\left\{\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} f\left(Z_{i}\right)-\mathbf{E}\left\{f\left(Z_{1}\right)\right\}\right|>\epsilon\right\} \leq 8 s(\widehat{\mathcal{F}}, n) e^{-n \epsilon^{2} /\left(32 M^{2}\right)}
$$

Example. Consider the empirical squared error minimization problem sketched in the previous section. Let $\mathcal{F}$ be the class of monotone increasing functions $\eta^{\prime}$ : $\mathcal{R} \rightarrow[0,1]$, and let $\eta_{n}$ be the function selected by minimizing the empirical squared error. By (29.1), if $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$ is also monotone increasing, then

$$
\begin{aligned}
& \mathbf{P}\left\{L\left(\eta_{n}\right)-L^{*}>\epsilon\right\} \\
& \quad \leq \mathbf{P}\left\{\sup _{\eta^{\prime} \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n}\left(\eta^{\prime}\left(X_{i}\right)-Y_{i}\right)^{2}-\mathbf{E}\left\{\left(\eta^{\prime}(X)-Y\right)^{2}\right\}\right|>\frac{\epsilon^{2}}{8}\right\} .
\end{aligned}
$$

If $\widehat{\mathcal{F}}$ contains all subsets of $\mathcal{R} \times\{0,1\}$ of the form

$$
A_{\eta^{\prime}, t}=\left\{(x, y):\left(\eta^{\prime}(x)-y\right)^{2}>t\right\}, \quad \eta^{\prime} \in \mathcal{F}, t \in[0,1]
$$

then it is easy to see that its $n$-th shatter coefficient satisfies $s(\widehat{\mathcal{F}}, n) \leq(n / 2+1)^{2}$. Thus, Corollary 29.1 can be applied, and the empirical squared error minimization is consistent.

In many cases, Corollary 29.1 does not provide the best possible bound. To state a similar, but sometimes more useful result, we introduce $l_{1}$-covering numbers. The notion is very similar to that of covering numbers discussed in Chapter 28. The main difference is that here the balls are defined in terms of an $l_{1}$-distance, rather than the supremum norm.

Definition 29.1. Let $A$ be a bounded subset of $\mathcal{R}^{d}$. For every $\epsilon>0$, the $l_{1}$ covering number, denoted by $\mathcal{N}(\epsilon, A)$, is defined as the cardinality of the smallest finite set in $\mathcal{R}^{d}$ such that for every $z \in A$ there is a point $t \in \mathcal{R}^{d}$ in the finite set such that $(1 / d)\|z-t\|_{1}<\epsilon .\left(\|x\|_{1}=\sum_{i=1}^{d}\left|x^{(i)}\right|\right.$ denotes the $l_{1}$-norm of the vector $x=\left(x^{(1)}, \ldots, x^{(d)}\right)$ in $\mathcal{R}^{d}$.) In other words, $\mathcal{N}(\epsilon, A)$ is the smallest number of $l_{1}$-balls of radius $\epsilon d$, whose union contains $A . \log \mathcal{N}(\epsilon, A)$ is often called the metric entropy of $A$.

We will mainly be interested in covering numbers of special sets. Let $z_{1}^{n}=$ $\left(z_{1}, \ldots, z_{n}\right)$ be $n$ fixed points in $\mathcal{R}^{d}$, and define the following set:

$$
\mathcal{F}\left(z_{1}^{n}\right)=\left\{\left(f\left(z_{1}\right), \ldots, f\left(z_{n}\right)\right) ; f \in \mathcal{F}\right\} \subset \mathcal{R}^{n}
$$

The $l_{1}$-covering number of $\mathcal{F}\left(z_{1}^{n}\right)$ is $\mathcal{N}\left(\epsilon, \mathcal{F}\left(z_{1}^{n}\right)\right)$.
If $Z_{1}^{n}=\left(Z_{1}, \ldots, Z_{n}\right)$ is a sequence of i.i.d. random variables, then $\mathcal{N}\left(\epsilon, \mathcal{F}\left(Z_{1}^{n}\right)\right)$ is a random variable, whose expected value plays a central role in our problem:

Theorem 29.1. (Pollard (1984)). For any $n$ and $\epsilon>0$,

$$
\mathbb{P}\left\{\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} f\left(Z_{i}\right)-\mathbb{E}\left\{f\left(Z_{1}\right)\right\}\right|>\epsilon\right\} \leq 8 \mathbf{E}\left\{\mathcal{N}\left(\epsilon / 8, \mathcal{F}\left(Z_{1}^{n}\right)\right)\right\} e^{-n \epsilon^{2} /\left(128 M^{2}\right)}
$$

The proof of the theorem is given in Section 29.4.
REmark. Theorem 29.1 is a generalization of the basic Vapnik-Chervonenkis inequality. To see this, define $l_{\infty}$-covering numbers based on the maximum norm (Vapnik and Chervonenkis (1981)): $\mathcal{N}_{\infty}(\epsilon, A)$ is the cardinality of the smallest finite set in $\mathcal{R}^{d}$ such that for every $z \in A$ there is a point $t \in \mathcal{R}^{d}$ in the set such that $\max _{1 \leq i \leq d}\left|z^{(i)}-t^{(i)}\right|<\epsilon$. If the functions in $\mathcal{F}$ are indicators of sets from a class $\mathcal{A}$ of subsets of $\mathcal{R}^{d}$, then it is easy to see that for every $\epsilon \in(0,1 / 2)$,

$$
\mathcal{N}_{\infty}\left(\epsilon, \mathcal{F}\left(z_{1}^{n}\right)\right)=N_{\mathcal{A}}\left(z_{1}, \ldots, z_{n}\right)
$$

where $N_{\mathcal{A}}\left(z_{1}, \ldots, z_{n}\right)$ is the combinatorial quantity that was used in Definition 12.1 of shatter coefficients. Since

$$
\mathcal{N}\left(\epsilon, \mathcal{F}\left(z_{1}^{n}\right)\right) \leq \mathcal{N}_{\infty}\left(\epsilon, \mathcal{F}\left(z_{1}^{n}\right)\right)
$$

Theorem 29.1 remains true with $l_{\infty}$-covering numbers, therefore, it is a generalization of Theorem 12.5 . To see this, notice that if $\mathcal{F}$ contains indicators of sets of the class $\mathcal{A}$, then

$$
\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} f\left(Z_{i}\right)-\mathbf{E}\left\{f\left(Z_{1}\right)\right\}\right|=\sup _{A \in \mathcal{A}}\left|\frac{1}{n} \sum_{i=1}^{n} I_{\left\{Z_{i} \in A\right\}}-\mathbb{P}\left\{Z_{1} \in A\right\}\right|
$$

For inequalities sharper and more general than Theorem 29.1 we refer to Vapnik (1982), Pollard (1984; 1986), Haussler (1992), and Anthony and Shawe-Taylor (1990).

### 29.3 Empirical Squared Error Minimization

We return to the minimization of the empirical squared error. Let $\mathcal{F}$ be a class of functions $\eta^{\prime}: \mathcal{R}^{d} \rightarrow[0,1]$, containing the true a posteriori function $\eta$. The empirical squared error

$$
\frac{1}{n} \sum_{i=1}^{n}\left(\eta^{\prime}\left(X_{i}\right)-Y_{i}\right)^{2}
$$

is minimized over $\eta^{\prime} \in \mathcal{F}$, to obtain the estimate $\eta_{n}$. The next result shows that empirical squared error minimization is consistent under general conditions. Observe that these are the same conditions that we assumed in Theorem 28.1 to prove consistency of skeleton estimates.

Theorem 29.2. Assume that $\mathcal{F}$ is a totally bounded class of functions. (For the definition see Chapter 28.) If $\eta \in \mathcal{F}$, then the classification rule obtained by minimizing the empirical squared error over $\mathcal{F}$ is strongly consistent, that is,

$$
\lim _{n \rightarrow \infty} L\left(\eta_{n}\right)=L^{*} \quad \text { with probability one }
$$

Proof. Recall that by (29.1),

$$
\begin{aligned}
& \mathbf{P}\left\{L\left(\eta_{n}\right)-L^{*}>\epsilon\right\} \\
& \quad \leq \mathbf{P}\left\{\sup _{\eta^{\prime} \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n}\left(\eta^{\prime}\left(X_{i}\right)-Y_{i}\right)^{2}-\mathbf{E}\left\{\left(\eta^{\prime}(X)-Y\right)^{2}\right\}\right|>\frac{\epsilon^{2}}{8}\right\} .
\end{aligned}
$$

We apply Theorem 29.1 to show that for every $\epsilon>0$, the probability on the righthand side converges to zero exponentially as $n \rightarrow \infty$. To this end, we need to find a suitable upper bound on $\mathbf{E}\left\{\mathcal{N}\left(\epsilon, \mathcal{J}\left(Z_{1}^{n}\right)\right)\right\}$, where $\mathcal{J}$ is the class of functions
$f^{\prime}(x, y)=\left(\eta^{\prime}(x)-y\right)^{2}$ from $\mathcal{R}^{d} \times\{0,1\}$ to $[0,1]$, where $\eta^{\prime} \in \mathcal{F}$, and $Z_{i}=\left(X_{i}, Y_{i}\right)$. Observe that for any $f^{\prime}, \tilde{f} \in \mathcal{J}$,

$$
\begin{aligned}
\frac{1}{n} \sum_{i=1}^{n}\left|f^{\prime}\left(Z_{i}\right)-\tilde{f}\left(Z_{i}\right)\right| & =\frac{1}{n} \sum_{i=1}^{n}\left|\left(\eta^{\prime}\left(X_{i}\right)-Y_{i}\right)^{2}-\left(\tilde{\eta}\left(X_{i}\right)-Y_{i}\right)^{2}\right| \\
& \leq \frac{2}{n} \sum_{i=1}^{n}\left|\eta^{\prime}\left(X_{i}\right)-\tilde{\eta}\left(X_{i}\right)\right| \\
& \leq 2 \sup \left|\eta^{\prime}(x)-\tilde{\eta}(x)\right| .
\end{aligned}
$$

This inequality implies that for every $\epsilon>0$,

$$
\mathbf{E}\left\{\mathcal{N}\left(\epsilon, \mathcal{J}\left(Z_{1}^{n}\right)\right)\right\} \leq \mathcal{N}_{\epsilon / 2},
$$

where $\mathcal{N}_{\epsilon}$ is the covering number of $\mathcal{F}$ defined in Chapter 28. By the assumption of total boundedness, for every $\epsilon>0, \mathcal{N}_{\epsilon / 2}<\infty$. Since $\mathcal{N}_{\epsilon / 2}$ does not depend on $n$, the theorem is proved.

REMARK. The nonasymptotic exponential nature of the inequality in Theorem 29.1 makes it possible to obtain upper bounds for the rate of convergence of $L\left(\eta_{n}\right)$ to $L^{*}$ in terms of the covering numbers $\mathcal{N}_{\epsilon}$ of the class $\mathcal{F}$. However, since we started our analysis by the loose inequality $L\left(\eta^{\prime}\right)-L^{*} \leq 2 \sqrt{\mathbf{E}\left\{\left(\eta^{\prime}(X)-\eta(X)\right)^{2}\right\}}$, the resulting rates are likely to be suboptimal (see Theorem 6.5). Also, the inequality of Theorem 29.1 may be loose in this case. In a somewhat different setup, Barron (1991) developed a proof method based on Bernstein's inequality that is useful for obtaining tighter upper bounds for $L\left(\eta_{n}\right)-L^{*}$ in certain cases.

### 29.4 Proof of Theorem 29.1

The main tricks in the proof resemble those of Theorem 12.5. We can show that for $n \epsilon^{2} \geq 2 M^{2}$,

$$
\mathbf{P}\left\{\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} f\left(Z_{i}\right)-\mathbf{E}\left\{f\left(Z_{1}\right)\right\}\right|>\epsilon\right\} \leq 4 \mathbf{P}\left\{\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} f\left(Z_{i}\right)\right|>\frac{\epsilon}{4}\right\},
$$

where $\sigma_{1}, \ldots, \sigma_{n}$ are i.i.d. $\{-1,1\}$-valued random variables, independent of the $Z_{i}$ 's, with $\mathbf{P}\left\{\sigma_{i}=1\right\}=\mathbf{P}\left\{\sigma_{i}=-1\right\}=1 / 2$. The only minor difference with Theorem 12.5 appears when Chebyshev's inequality is applied. We use the fact that by boundedness, $\operatorname{Var}\left(f\left(Z_{1}\right)\right) \leq M^{2} / 4$ for every $f \in \mathcal{F}$.

Now, take a minimal $\epsilon / 8$-covering of $\mathcal{F}\left(Z_{1}^{n}\right)$, that is, $M=\mathcal{N}\left(\epsilon / 8, \mathcal{F}\left(Z_{1}^{n}\right)\right)$ functions $g_{1}, \ldots, g_{M}$ such that for every $f \in \mathcal{F}$ there is a $g^{*} \in\left\{g_{1}, \ldots, g_{M}\right\}$ with

$$
\frac{1}{n} \sum_{i=1}^{n}\left|f\left(Z_{i}\right)-g^{*}\left(Z_{i}\right)\right| \leq \frac{\epsilon}{8} .
$$

For any function $f$, we have

$$
\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} f\left(Z_{i}\right)\right| \leq \frac{1}{n} \sum_{i=1}^{n}\left|\sigma_{i} f\left(Z_{i}\right)\right|=\frac{1}{n} \sum_{i=1}^{n}\left|f\left(Z_{i}\right)\right|
$$

and thus

$$
\begin{aligned}
\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} f\left(Z_{i}\right)\right| & \leq\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} g^{*}\left(Z_{i}\right)\right|+\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i}\left(f\left(Z_{i}\right)-g^{*}\left(Z_{i}\right)\right)\right| \\
& \leq\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} g^{*}\left(Z_{i}\right)\right|+\frac{1}{n} \sum_{i=1}^{n}\left|f\left(Z_{i}\right)-g^{*}\left(Z_{i}\right)\right|
\end{aligned}
$$

As $(1 / n) \sum_{i=1}^{n}\left|g^{*}\left(Z_{i}\right)-f\left(Z_{i}\right)\right| \leq \epsilon / 8$,

$$
\begin{aligned}
& \mathbf{P}\left\{\left.\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} f\left(Z_{i}\right)\right|>\frac{\epsilon}{4} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} \\
& \quad \leq \mathbf{P}\left\{\left.\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} g^{*}\left(Z_{i}\right)\right|+\frac{1}{n} \sum_{i=1}^{n}\left|f\left(Z_{i}\right)-g^{*}\left(Z_{i}\right)\right|>\frac{\epsilon}{4} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} \\
& \quad \leq \mathbf{P}\left\{\left.\max _{g_{j}}\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} g_{j}\left(Z_{i}\right)\right|>\frac{\epsilon}{8} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} .
\end{aligned}
$$

Now that we have been able to convert the "sup" into a "max," we can use the union bound:

$$
\begin{aligned}
& \mathbf{P}\left\{\left.\max _{g_{j}}\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} g_{j}\left(Z_{i}\right)\right|>\frac{\epsilon}{8} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} \\
& \quad \leq \mathcal{N}\left(\frac{\epsilon}{8}, \mathcal{F}\left(Z_{1}^{n}\right)\right) \max _{g_{j}} \mathbf{P}\left\{\left.\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} g_{j}\left(Z_{i}\right)\right|>\frac{\epsilon}{8} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} .
\end{aligned}
$$

We need only find a uniform bound for the probability following the "max." This, however, is easy, since after conditioning on $Z_{1}, \ldots, Z_{n}, \sum_{i=1}^{n} \sigma_{i} g_{j}\left(Z_{i}\right)$ is the sum of independent bounded random variables whose expected value is zero. Therefore, Hoeffding's inequality gives

$$
\mathbf{P}\left\{\left.\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} g_{j}\left(Z_{i}\right)\right|>\frac{\epsilon}{8} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\} \leq 2 e^{-n \epsilon^{2} /\left(128 M^{2}\right)}
$$

In summary,

$$
\begin{aligned}
& \mathbf{P}\left\{\sup _{f \in \mathcal{F}}\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} f\left(Z_{i}\right)\right|>\frac{\epsilon}{4}\right\} \\
& \quad \leq \mathbf{E}\left\{\mathbf{P}\left\{\left.\max _{g_{j}}\left|\frac{1}{n} \sum_{i=1}^{n} \sigma_{i} g_{j}\left(Z_{i}\right)\right|>\frac{\epsilon}{8} \right\rvert\, Z_{1}, \ldots, Z_{n}\right\}\right\}
\end{aligned}
$$

$$
\leq 2 \mathbf{E}\left\{\mathcal{N}\left(\frac{\epsilon}{8}, \mathcal{F}\left(Z_{1}^{n}\right)\right)\right\} e^{-n \epsilon^{2} /\left(128 M^{2}\right)}
$$

The theorem is proved.

Properties of $\mathcal{N}\left(\epsilon, \mathcal{F}\left(Z_{1}^{n}\right)\right)$ will be studied in the next section.

### 29.5 Covering Numbers and Shatter Coefficients

In this section we study covering numbers, and relate them to shatter coefficients of certain classes of sets. As in Chapter 28, we introduce $l_{1}$-packing numbers. Let $\mathcal{F}$ be a class of functions on $\mathcal{R}^{d}$, taking their values in $[0, M]$. Let $\mu$ be an arbitrary probability measure on $\mathcal{R}^{d}$. Let $g_{1}, \ldots, g_{m}$ be a finite collection of functions from $\mathcal{F}$ with the property that for any two of them

$$
\int_{\mathcal{R}^{d}}\left|g_{i}(x)-g_{j}(x)\right| \mu(d x) \geq \epsilon .
$$

The largest $m$ for which such a collection exists is called the packing number of $\mathcal{F}$ (relative to $\mu$ ), and is denoted by $\mathcal{M}(\epsilon, \mathcal{F})$. If $\mu$ places probability $1 / n$ on each of $z_{1}, \ldots, z_{n}$, then by definition $\mathcal{M}(\epsilon, \mathcal{F})=\mathcal{M}\left(\epsilon, \mathcal{F}\left(z_{1}^{n}\right)\right)$, and it is easy to see (Problem 28.3) that

$$
\mathcal{M}\left(2 \epsilon, \mathcal{F}\left(z_{1}^{n}\right)\right) \leq \mathcal{N}\left(\epsilon, \mathcal{F}\left(z_{1}^{n}\right)\right) \leq \mathcal{M}\left(\epsilon, \mathcal{F}\left(z_{1}^{n}\right)\right)
$$

An important feature of a class of functions $\mathcal{F}$ is the vc dimension $V_{\mathcal{F}+}$ of

$$
\mathcal{F}^{+}=\{\{(x, t): t \leq f(x)\} ; f \in \mathcal{F}\} .
$$

This is clarified by the following theorem, which is a slight refinement of a result by Pollard (1984), which is based on Dudley's (1978) work. It connects the packing number of $\mathcal{F}$ with the shatter coefficients of $\mathcal{F}^{+}$. See also Haussler (1992) for a somewhat different argument.

Theorem 29.3. Let $\mathcal{F}$ be a class of $[0, M]$-valued functions on $\mathcal{R}^{d}$. For every $\epsilon>0$ and probability measure $\mu$,

$$
\mathcal{M}(\epsilon, \mathcal{F}) \leq s\left(\mathcal{F}^{+}, k\right)
$$

where $k=\left\lceil\frac{M}{\epsilon} \log \frac{e \in \mathcal{M}^{2}(\epsilon, \mathcal{F})}{2 M}\right\rceil$.
Proof. Let $\left\{g_{1}, g_{2}, \ldots, g_{m}\right\}$ be an arbitrary $\epsilon$-packing of $\mathcal{F}$ of size $m \leq \mathcal{M}(\epsilon, \mathcal{F})$. The proof is in the spirit of the probabilistic method of combinatorics (see, e.g., Spencer (1987)). To prove the inequality, we create $k$ random points on $\mathcal{R}^{d} \times[0, M]$ in the following way, where $k$ is a positive integer specified later. We generate $k$ independent random variables $S_{1}, \ldots, S_{k}$ on $\mathcal{R}^{d}$ with common distribution
$\mu$, and independently of this, we generate another $k$ independent random variables $T_{1}, \ldots, T_{k}$, uniformly distributed on $[0, M]$. This yields $k$ random pairs $\left(S_{1}, T_{1}\right), \ldots,\left(S_{k}, T_{k}\right)$. For any two functions $g_{i}$ and $g_{j}$ in an $\epsilon$-packing, the probability that the sets $G_{i}=\left\{(x, t): t \leq g_{i}(x)\right\}$ and $G_{j}=\left\{(x, t): t \leq g_{j}(x)\right\}$ pick the same points from our random set of $k$ points is bounded as follows:

$$
\begin{aligned}
\mathbf{P} & \left\{G_{i} \text { and } G_{j} \text { pick the same points }\right\} \\
& =\prod_{l=1}^{k}\left(1-\mathbf{P}\left\{\left(S_{l}, T_{l}\right) \in G_{i} \Delta G_{j}\right\}\right) \\
& =\left(1-\mathbf{E}\left\{\mathbf{P}\left\{\left(S_{1}, T_{1}\right) \in G_{i} \Delta G_{j}\right\} \mid S_{1}\right\}\right)^{k} \\
& =\left(1-\frac{1}{M} \mathbf{E}\left\{\left|g_{i}\left(S_{1}\right)-g_{j}\left(S_{1}\right)\right|\right\}\right)^{k} \\
& \leq(1-\epsilon / M)^{k} \\
& \leq e^{-k \epsilon / M}
\end{aligned}
$$

where we used the definition of the functions $g_{1}, \ldots, g_{m}$. Observe that the expected number of pairs $\left(g_{i}, g_{j}\right)$ of these functions, such that the corresponding sets $G_{i}=$ $\left\{(x, t): t \leq g_{i}(x)\right\}$ and $G_{j}=\left\{(x, t): t \leq g_{j}(x)\right\}$ pick the same points, is bounded by

$$
\begin{aligned}
& \mathbb{E}\left\{\left\{\left\{\left(g_{i}, g_{j}\right) ; G_{i} \text { and } G_{j} \text { pick the same points }\right\} \mid\right\}\right. \\
& \\
& =\binom{m}{2} \mathbf{P}\left\{G_{i} \text { and } G_{j} \text { pick the same points }\right\} \leq\binom{ m}{2} e^{-k \epsilon / M}
\end{aligned}
$$

Since for $k$ randomly chosen points the average number of pairs that pick the same points is bounded by $\binom{m}{2} e^{-k \epsilon / M}$, there exist $k$ points in $\mathcal{R}^{d} \times[0, M]$, such that the number of pairs $\left(g_{i}, g_{j}\right)$ that pick the same points is actually bounded by $\binom{m}{2} e^{-k \epsilon / M}$. For each such pair we can add one more point in $\mathcal{R}^{d} \times[0, M]$ such that the point is contained in $G_{i} \Delta G_{j}$. Thus, we have obtained a set of no more than $k+\binom{m}{2} e^{-k \epsilon / M}$ points such that the sets $G_{1}, \ldots, G_{m}$ pick different subsets of it. Since $k$ was arbitrary, we can choose it to minimize this expression. This yields $\left\lfloor\frac{M}{\epsilon} \log \left(e \epsilon\binom{m}{2} / M\right)\right\rfloor$ points, so the shatter coefficient of $\mathcal{F}^{+}$corresponding to this number must be greater than $m$, which proves the statement.

The meaning of Theorem 29.3 is best seen from the following simple corollary:
Corollary 29.2. Let $\mathcal{F}$ be a class of $[0, M]$-valued functions on $\mathcal{R}^{d}$. For every $\epsilon>0$ and probability measure $\mu$,

$$
\mathcal{M}(\epsilon, \mathcal{F}) \leq\left(\frac{4 e M}{\epsilon} \log \frac{2 e M}{\epsilon}\right)^{V_{\mathcal{F +}}}
$$

Proof. Recall that Theorem 13.2 implies

$$
s\left(\mathcal{F}^{+}, k\right) \leq\left(\frac{e k}{V_{\mathcal{F}+}}\right)^{V_{\mathcal{F}^{+}}}
$$

The inequality follows from Theorem 29.3 by straightforward calculation. The details are left as an exercise (Problem 29.2).

Recently Haussler (1991) was able to get rid of the "log" factor in the above upper bound. He proved that if $\epsilon=k / n$ for an integer $k$, then

$$
\mathcal{M}(\epsilon, \mathcal{F}) \leq e(d+1)\left(\frac{2 e}{\epsilon}\right)^{V_{\mathcal{F}^{+}}}
$$

The quantity $V_{\mathcal{F}^{+}}$is sometimes called the pseudo dimension of $\mathcal{F}$ (see Problem 29.3). It follows immediately from Theorem 13.9 that if $\mathcal{F}$ is a linear space of functions of dimension $r$, then its pseudo dimension is at most $r+1$. A few more properties are worth mentioning:

Theorem 29.4. (Wenocur and Dudley, (1981)). Let $g: \mathcal{R}^{d} \rightarrow \mathcal{R}$ be an arbitrary function, and consider the class of functions $\mathcal{G}=\{g+f ; f \in \mathcal{F}\}$. Then

$$
V_{\mathcal{G}^{+}}=V_{\mathcal{F}^{+}} .
$$

PROOF. If the points $\left(s_{1}, t_{1}\right), \ldots,\left(s_{k}, t_{k}\right) \in \mathcal{R}^{d} \times \mathcal{R}$ are shattered by $\mathcal{F}^{+}$, then the points $\left(s_{1}, t_{1}+g\left(s_{1}\right)\right), \ldots,\left(s_{k}, t_{k}+g\left(s_{k}\right)\right)$ are shattered by $\mathcal{G}^{+}$. This proves

$$
V_{\mathcal{G}^{+}} \geq V_{\mathcal{F}^{+}}
$$

The proof of the other inequality is similar.

Theorem 29.5. (Nolan and Pollard (1987); Dudley, (1987)). Let g: [0, M] $\rightarrow$ $\mathcal{R}$ be a fixed nondecreasing function, and define the class $\mathcal{G}=\{g \circ f ; f \in \mathcal{F}\}$. Then

$$
V_{\mathcal{G}^{+}} \leq V_{\mathcal{F}^{+}}
$$

Proof. Assume that $n \leq V_{\mathcal{G}^{+}}$, and let the functions $f_{1}, \ldots, f_{2^{n}} \in \mathcal{F}$ be such that the binary vector

$$
\left(I_{\left\{g\left(f_{j}\left(s_{1}\right)\right) \geq t_{1}\right\}}, \ldots, I_{\left\{g\left(f_{j}\left(s_{n}\right)\right) \geq t_{n}\right\}}\right)
$$

takes all $2^{n}$ values if $j=1, \ldots, 2^{n}$. For all $1 \leq i \leq n$ define the numbers

$$
u_{i}=\min _{1 \leq j \leq 2^{n}}\left\{f_{j}\left(s_{i}\right): g\left(f_{j}\left(s_{i}\right)\right) \geq t_{i}\right\}
$$

and

$$
l_{i}=\max _{1 \leq j \leq 2^{n}}\left\{f_{j}\left(s_{i}\right): g\left(f_{j}\left(s_{i}\right)\right)<t_{i}\right\}
$$

By the monotonicity of $g, u_{i}>l_{i}$. Then the binary vector

$$
\left(I_{\left\{f_{j}\left(s_{1}\right) \geq \frac{u_{1}+l_{l}}{2}\right\}}, \ldots, I_{\left\{f_{j}\left(s_{n}\right) \geq \frac{u_{n}+l_{n}}{2}\right\}}\right)
$$

takes the same value as

$$
\left(I_{\left\{g\left(f_{j}\left(s_{1}\right)\right) \geq t_{1}\right\}}, \ldots, I_{\left\{g\left(f_{j}\left(s_{n}\right)\right) \geq t_{n}\right\}}\right)
$$

for every $j \leq 2^{n}$. Therefore, the pairs

$$
\left(s_{1}, \frac{u_{1}+l_{1}}{2}\right), \ldots,\left(s_{n}, \frac{u_{n}+l_{n}}{2}\right)
$$

are shattered by $\mathcal{F}^{+}$, which proves the theorem.

Next we present a few results about covering numbers of classes of functions whose members are sums or products of functions from other classes. Similar results can be found in Nobel (1992), Nolan and Pollard (1987), and Pollard (1990).

Theorem 29.6. Let $\mathcal{F}_{1}, \ldots, \mathcal{F}_{k}$ be classes of real functions on $\mathcal{R}^{d}$. Forn arbitrary, fixed points $z_{1}^{n}=\left(z_{1}, \ldots, z_{n}\right)$ in $\mathcal{R}^{d}$, define the sets $\mathcal{F}_{1}\left(z_{1}^{n}\right), \ldots, \mathcal{F}_{k}\left(z_{1}^{n}\right)$ in $\mathcal{R}^{n}$ by

$$
\mathcal{F}_{j}\left(z_{1}^{n}\right)=\left\{\left(f_{j}\left(z_{1}\right), \ldots, f_{j}\left(z_{n}\right)\right) ; f_{j} \in \mathcal{F}_{j}\right\}
$$

$j=1, \ldots, k$. Also, introduce

$$
\mathcal{F}\left(z_{1}^{n}\right)=\left\{\left(f\left(z_{1}\right), \ldots, f\left(z_{n}\right)\right) ; f \in \mathcal{F}\right\}
$$

for the class of functions

$$
\mathcal{F}=\left\{f_{1}+\ldots+f_{k} ; f_{j} \in \mathcal{F}_{j}, \quad j=1, \ldots, k\right\}
$$

Then for every $\epsilon>0$ and $z_{1}^{n}$

$$
\mathcal{N}\left(\epsilon, \mathcal{F}\left(z_{1}^{n}\right)\right) \leq \prod_{j=1}^{k} \mathcal{N}\left(\epsilon / k, \mathcal{F}_{j}\left(z_{1}^{n}\right)\right)
$$

Proof. Let $S_{1}, \ldots, S_{k} \subset \mathcal{R}^{n}$ be minimal $\epsilon / k$-coverings of $\mathcal{F}_{1}\left(z_{1}^{n}\right), \ldots, \mathcal{F}_{k}\left(z_{1}^{n}\right)$, respectively. This implies that for any $f_{j} \in \mathcal{F}_{j}$ there is a vector $s_{j}=\left(s_{j}^{(1)}, \ldots, s_{j}^{(n)}\right) \in$ $S_{j}$ such that

$$
\frac{1}{n} \sum_{i=1}^{n}\left|f_{j}\left(z_{i}\right)-s_{j}^{(i)}\right|<\epsilon / k,
$$

for every $j=1, \ldots, k$. Moreover, $\left|S_{j}\right|=\mathcal{N}\left(\epsilon / k, \mathcal{F}_{j}\left(z_{1}^{n}\right)\right)$. We show that

$$
S=\left\{s_{1}+\ldots+s_{k} ; s_{j} \in S_{j}, j=1, \ldots, k\right\}
$$

is an $\epsilon$-covering of $\mathcal{F}\left(z_{1}^{n}\right)$. This follows immediately from the triangle inequality, since for any $f_{1}, \ldots, f_{k}$ there is $s_{1}, \ldots, s_{k}$ such that

$$
\begin{aligned}
& \frac{1}{n} \sum_{i=1}^{n}\left|f_{1}\left(z_{i}\right)+\ldots+f_{k}\left(z_{i}\right)-\left(s_{1}^{(i)}+\ldots+s_{k}^{(i)}\right)\right| \\
& \quad \leq \frac{1}{n} \sum_{i=1}^{n}\left|f_{1}\left(z_{i}\right)-s_{1}^{(i)}\right|+\ldots+\frac{1}{n} \sum_{i=1}^{n}\left|f_{k}\left(z_{i}\right)-s_{k}^{(i)}\right|<k \frac{\epsilon}{k}
\end{aligned}
$$

Theorem 29.7. (Pollard (1990)). Let $\mathcal{F}$ and $\mathcal{G}$ be classes of real functions on $\mathcal{R}^{d}$, bounded by $M_{1}$ and $M_{2}$, respectively. (That is, e.g., $|f(x)| \leq M_{1}$ for every $x \in \mathcal{R}^{d}$ and $f \in \mathcal{F}$.) For arbitrary fixed points $z_{1}^{n}=\left(z_{1}, \ldots, z_{n}\right)$ in $\mathcal{R}^{d}$ define the sets $\mathcal{F}\left(z_{1}^{n}\right)$ and $\mathcal{G}\left(z_{1}^{n}\right)$ in $\mathcal{R}^{n}$ as in Theorem 29.6. Introduce

$$
\mathcal{J}\left(z_{1}^{n}\right)=\left\{\left(h\left(z_{1}\right), \ldots, h\left(z_{n}\right)\right) ; h \in \mathcal{J}\right\}
$$

for the class of functions

$$
\mathcal{J}=\{f g ; f \in \mathcal{F}, g \in \mathcal{G}\}
$$

Then for every $\epsilon>0$ and $z_{1}^{n}$

$$
\mathcal{N}\left(\epsilon, \mathcal{J}\left(z_{1}^{n}\right)\right) \leq \mathcal{N}\left(\frac{\epsilon}{2 M_{2}}, \mathcal{F}\left(z_{1}^{n}\right)\right) \cdot \mathcal{N}\left(\frac{\epsilon}{2 M_{1}}, \mathcal{G}\left(z_{1}^{n}\right)\right)
$$

Proof. Let $S \subset\left[-M_{1}, M_{1}\right]^{n}$ be an $\epsilon /\left(2 M_{2}\right)$-covering of $\mathcal{F}\left(z_{1}^{n}\right)$, that is, for any $f \in \mathcal{F}$ there is a vector $s=\left(s^{(1)}, \ldots, s^{(n)}\right) \in S$ such that

$$
\frac{1}{n} \sum_{i=1}^{n}\left|f\left(z_{i}\right)-s^{(i)}\right|<\frac{\epsilon}{2 M_{2}}
$$

It is easy to see that $S$ can be chosen such that $|S|=\mathcal{N}\left(\epsilon /\left(2 M_{2}\right), \mathcal{F}\left(z_{1}^{n}\right)\right)$. Similarly, let $T \subset\left[-M_{2}, M_{2}\right]$ be an $\epsilon /\left(2 M_{1}\right)$-covering of $\mathcal{G}\left(z_{1}^{n}\right)$ with $|T|=\mathcal{N}\left(\epsilon /\left(2 M_{1}\right)\right.$, $\left.\mathcal{G}\left(z_{1}^{n}\right)\right)$ such that for any $g \in \mathcal{G}$ there is a $t=\left(t^{(1)}, \ldots, t^{(n)}\right) \in T$ with

$$
\frac{1}{n} \sum_{i=1}^{n}\left|g\left(z_{i}\right)-t^{(i)}\right|<\frac{\epsilon}{2 M_{1}}
$$

We show that the set

$$
U=\{s t ; s \in S, t \in T\}
$$

is an $\epsilon$-covering of $\mathcal{J}\left(z_{1}^{n}\right)$. Let $f \in \mathcal{F}$ and $g \in \mathcal{G}$ be arbitrary and $s \in S$ and $t \in T$ the corresponding vectors such that

$$
\frac{1}{n} \sum_{i=1}^{n}\left|f\left(z_{i}\right)-s^{(i)}\right|<\frac{\epsilon}{2 M_{2}} \quad \text { and } \quad \frac{1}{n} \sum_{i=1}^{n}\left|g\left(z_{i}\right)-t^{(i)}\right|<\frac{\epsilon}{2 M_{1}}
$$

Then

$$
\begin{aligned}
& \frac{1}{n} \sum_{i=1}^{n}\left|f\left(z_{i}\right) g\left(z_{i}\right)-s^{(i)} t^{(i)}\right| \\
& \quad=\frac{1}{n} \sum_{i=1}^{n}\left|f\left(z_{i}\right)\left(t^{(i)}+g\left(z_{i}\right)-t^{(i)}\right)-s^{(i)} t^{(i)}\right| \\
& \quad=\frac{1}{n} \sum_{i=1}^{n}\left|t^{(i)}\left(f\left(z_{i}\right)-s^{(i)}\right)+f\left(z_{i}\right)\left(g\left(z_{i}\right)-t^{(i)}\right)\right| \\
& \quad \leq \frac{M_{2}}{n} \sum_{i=1}^{n}\left|f\left(z_{i}\right)-s^{(i)}\right|+\frac{M_{1}}{n} \sum_{i=1}^{n}\left|g\left(z_{i}\right)-t^{(i)}\right|<\epsilon
\end{aligned}
$$

### 29.6 Generalized Linear Classification

In this section we use the uniform laws of large numbers discussed in this chapter to prove that squared error minimization over an appropriately chosen class of generalized linear classifiers yields a universally consistent rule. Consider the class $\mathcal{C}^{\left(k_{n}\right)}$ of generalized linear classifiers, whose members are functions $\phi: \mathcal{R}^{d} \rightarrow$ $\{0,1\}$ of the form

$$
\phi(x)= \begin{cases}0 & \text { if } \sum_{j=1}^{k_{n}} a_{j} \psi_{j}(x) \leq 0 \\ 1 & \text { otherwise }\end{cases}
$$

where $\psi_{1}, \ldots, \psi_{k_{n}}$ are fixed basis functions, and the coefficients $a_{1}, \ldots, a_{k_{n}}$ are arbitrary real numbers. The training sequence $D_{n}$ is used to determine the coefficients $a_{i}$. In Chapter 17 we studied the behavior of the classifier whose coefficients are picked to minimize the empirical error probability

$$
\widehat{L}_{n}=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi\left(X_{i}\right)=Y_{i}\right\}} .
$$

Instead of minimizing the empirical error probability $\widehat{L}_{n}(\phi)$, several authors suggested minimizing the empirical squared error

$$
\frac{1}{n} \sum_{i=1}^{n}\left(\sum_{j=1}^{k_{n}} a_{j} \psi_{j}\left(X_{i}\right)-\left(2 Y_{i}-1\right)\right)^{2}
$$

(see, e.g., Duda and Hart (1973), Vapnik (1982)). This is rather dangerous. For example, for $k=1$ and $d=1$ it is easy to find a distribution such that the error probability of the linear classifier that minimizes the empirical squared error converges to $1-\epsilon$, while the error probability of the best linear classifier is $\epsilon$, where $\epsilon$ is an arbitrarily small positive number (Theorem 4.7). Clearly, similar examples can
be found for any fixed $k$. This demonstrates powerfully the danger of minimizing squared error instead of error count. Minimizing the latter yields a classifier whose average error probability is always within $O(\sqrt{\log n / n})$ of the optimum in the class, for fixed $k$. We note here that in some special cases minimization of the two types of error are equivalent (see Problem 29.5). Interestingly though, if $k_{n} \rightarrow \infty$ as $n$ increases, we can obtain universal consistency by minimizing the empirical squared error.

Theorem 29.8. Let $\psi_{1}, \psi_{2}, \ldots$ be a sequence of bounded functions with $\left|\psi_{j}(x)\right| \leq$ 1 such that the set of all finite linear combinations of the $\psi_{j}$ 's

$$
\bigcup_{k=1}^{\infty}\left\{\sum_{j=1}^{k} a_{j} \psi_{j}(x) ; a_{1}, a_{2}, \ldots \in \mathcal{R}\right\}
$$

is dense in $L_{2}(\mu)$ on all balls of the form $\{x:\|x\| \leq M\}$ for any probability measure $\mu$. Let the coefficients $a_{1}^{*}, \ldots, a_{k_{n}}^{*}$ minimize the empirical squared error

$$
\frac{1}{n} \sum_{i=1}^{n}\left(\sum_{j=1}^{k_{n}} a_{j} \psi_{j}\left(X_{i}\right)-\left(2 Y_{i}-1\right)\right)^{2}
$$

under the constraint $\sum_{j=1}^{k_{n}}\left|a_{j}\right| \leq b_{n}, b_{n} \geq 1$. Define the generalized linear classifier $g_{n}$ by

$$
g_{n}(x)= \begin{cases}0 & \text { if } \sum_{j=1}^{k_{n}} a_{j}^{*} \psi_{j}(x) \leq 0 \\ 1 & \text { otherwise }\end{cases}
$$

If $k_{n}$ and $b_{n}$ satisfy

$$
k_{n} \rightarrow \infty, \quad b_{n} \rightarrow \infty \text { and } \frac{k_{n} b_{n}^{4} \log \left(b_{n}\right)}{n} \rightarrow 0
$$

then $\mathbf{E}\left\{L\left(g_{n}\right)\right\} \rightarrow L^{*}$ for all distributions of $(X, Y)$, that is, the rule $g_{n}$ is universally consistent. If we assume additionally that $b_{n}^{4} \log n=o(n)$, then $g_{n}$ is strongly universally consistent.

Proof. Let $\delta>0$ be arbitrary. Then there exists a constant $M$ such that $\mathbf{P}\{\|X\|>$ $M\}<\delta$. Thus,

$$
L\left(g_{n}\right)-L^{*} \leq \delta+\mathbf{P}\left\{g_{n}(X) \neq Y,\|X\| \leq M \mid D_{n}\right\}-\mathbf{P}\left\{g^{*}(X) \neq Y,\|X\| \leq M\right\}
$$

It suffices to show that $\mathbf{P}\left\{g_{n}(X) \neq Y,\|X\| \leq M \mid D_{n}\right\}-\mathbf{P}\left\{g^{*}(X) \neq Y,\|X\| \leq\right.$ $M\} \rightarrow 0$ in the required sense for every $M>0$. Introduce the notation $f_{n}^{*}(x)=$ $\sum_{j=1}^{k_{n}} a_{j}^{*} \psi_{j}(x)$. By Corollary 6.2 , we see that

$$
\begin{aligned}
& \mathbf{P}\left\{g_{n}(X) \neq Y,\|X\| \leq M \mid D_{n}\right\}-\mathbf{P}\left\{g^{*}(X) \neq Y,\|X\| \leq M\right\} \\
& \leq \sqrt{\int_{\|x\| \leq M}\left(f_{n}^{*}(x)-(2 \eta(x)-1)\right)^{2} \mu(d x)}
\end{aligned}
$$

We prove that the right-hand side converges to zero in probability. Observe that since $\mathbf{E}\{2 Y-1 \mid X=x\}=2 \eta(x)-1$, for any function $h(x)$,

$$
\begin{aligned}
& (h(x)-(2 \eta(x)-1))^{2} \\
& \quad=\mathbf{E}\left\{(h(X)-(2 Y-1))^{2} \mid X=x\right\}-\mathbf{E}\left\{((2 Y-1)-(2 \eta(X)-1))^{2} \mid X=x\right\}
\end{aligned}
$$

(see Chapter 2), therefore, denoting the class of functions over which we minimize by

$$
\mathcal{F}_{n}=\left\{\sum_{j=1}^{k_{n}} a_{j} \psi_{j} ; \sum_{j=1}^{k_{n}}\left|a_{j}\right| \leq b_{n}\right\},
$$

we have

$$
\begin{aligned}
\int_{\|x\| \leq M}( & \left.f_{n}^{*}(x)-(2 \eta(x)-1)\right)^{2} \mu(d x) \\
= & \mathbf{E}\left\{\left(f_{n}^{*}(X)-(2 Y-1)\right)^{2} I_{\{\|X\| \leq M\}} \mid D_{n}\right\} \\
& -\mathbf{E}\left\{((2 Y-1)-(2 \eta(X)-1))^{2} I_{\{\|X\| \leq M\}}\right\} \\
= & \left(\mathbf{E}\left\{\left(f_{n}^{*}(X)-(2 Y-1)\right)^{2} I_{\{\|X\| \leq M\}} \mid D_{n}\right\}\right. \\
& \left.\quad \inf _{f \in \mathcal{F}_{n}} \mathbf{E}\left\{(f(X)-(2 Y-1))^{2} I_{\{\|X\| \leq M\}}\right\}\right) \\
& +\inf _{f \in \mathcal{F}_{n}} \mathbf{E}\left\{(f(X)-(2 Y-1))^{2} I_{\{\|X\| \leq M\}}\right\} \\
& -\mathbf{E}\left\{((2 Y-1)-(2 \eta(X)-1))^{2} I_{\{\|X\| \leq M\}}\right\}
\end{aligned}
$$

The last two terms may be combined to yield

$$
\inf _{f \in \mathcal{F}_{n}} \int_{\|x\| \leq M}(f(x)-(2 \eta(x)-1))^{2} \mu(d x)
$$

which converges to zero by the denseness assumption. To prove that the first term converges to zero in probability, observe that we may assume without loss of generality that $\mathbf{P}\{\|X\|>M\}=0$. As in the proof of Lemma 8.2 , it is easy to show that

$$
\begin{aligned}
\mathbf{E}\{ & \left.\left(f_{n}^{*}(X)-(2 Y-1)\right)^{2} I_{\{\|X\| \leq M\}} \mid D_{n}\right\} \\
& \quad-\inf _{f \in \mathcal{F}_{n}} \mathbf{E}\left\{(f(X)-(2 Y-1))^{2} I_{\{\|X\| \leq M\}}\right\} \\
= & \mathbf{E}\left\{\left(f_{n}^{*}(X)-(2 Y-1)\right)^{2} \mid D_{n}\right\}-\inf _{f \in \mathcal{F}_{n}} \mathbf{E}\left\{(f(X)-(2 Y-1))^{2}\right\} \\
\leq & 2 \sup _{f \in \mathcal{F}_{n}}\left|\frac{1}{n} \sum_{i=1}^{n}\left(f\left(X_{i}\right)-\left(2 Y_{i}-1\right)\right)^{2}-\mathbf{E}\left\{(f(X)-(2 Y-1))^{2}\right\}\right| \\
= & 2 \sup _{h \in \mathcal{J}}\left|\frac{1}{n} \sum_{i=1}^{n} h\left(X_{i}, Y_{i}\right)-\mathbf{E}\{h(X, Y)\}\right|,
\end{aligned}
$$

where the class of functions $\mathcal{J}$ is defined by

$$
\mathcal{J}=\left\{h(x, y)=(f(x)-(2 y-1))^{2} ; f \in \mathcal{F}_{n}\right\}
$$

Observe that since $|2 y-1|=1$ and $\left|\psi_{j}(x)\right| \leq 1$, we have

$$
0 \leq h(x, y) \leq\left(\sum_{j=1}^{k_{n}}\left|a_{j}\right|+1\right)^{2} \leq 2\left(b_{n}^{2}+1\right) \leq 4 b_{n}^{2}
$$

Therefore, Theorem 29.1 asserts that

$$
\begin{aligned}
& \mathbf{P}\left\{\mathbf{E}\left\{\left(f_{n}^{*}(X)-(2 Y-1)\right)^{2} \mid D_{n}\right\}\right. \\
&\left.-\inf _{f \in \mathcal{F}_{n}} \mathbf{E}\left\{(f(X)-(2 Y-1))^{2}\right\}>\epsilon\right\} \\
& \leq \mathbf{P}\left\{\sup _{h \in \mathcal{J}}\left|\frac{1}{n} \sum_{i=1}^{n} h\left(X_{i}, Y_{i}\right)-\mathbf{E}\{h(X, Y)\}\right|>\epsilon / 2\right\} \\
& \leq 8 \mathbf{E}\left\{\mathcal{N}\left(\frac{\epsilon}{16}, \mathcal{J}\left(Z_{1}^{n}\right)\right)\right\} e^{-n \epsilon^{2} /\left(512\left(4 b_{n}^{2}\right)^{2}\right)},
\end{aligned}
$$

where $Z_{1}^{n}=\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$. Next, for fixed $z_{1}^{n}$, we estimate the covering number $\mathcal{N}\left(\epsilon / 16, \mathcal{J}\left(z_{1}^{n}\right)\right)$. For arbitrary $f_{1}, f_{2} \in \mathcal{F}_{n}$, consider the functions $h_{1}(x, y)=\left(f_{1}(x)-(2 y-1)\right)^{2}$ and $h_{2}(x, y)=\left(f_{2}(x)-(2 y-1)\right)^{2}$. Then for any probability measure $v$ on $\mathcal{R}^{d} \times\{0,1\}$,

$$
\begin{aligned}
\int & \left|h_{1}(x, y)-h_{2}(x, y)\right| v(d(x, y)) \\
& =\int\left|\left(f_{1}(x)-(2 y-1)\right)^{2}-\left(f_{2}(x)-(2 y-1)\right)^{2}\right| \nu(d(x, y)) \\
& \leq \int 2\left|f_{1}(x)-f_{2}(x)\right|\left(b_{n}+1\right) v(d(x, y)) \\
& \leq 4 b_{n} \int\left|f_{1}(x)-f_{2}(x)\right| \mu(d x)
\end{aligned}
$$

where $\mu$ is the marginal measure for $v$ on $\mathcal{R}^{d}$. Thus, for any $z_{1}^{n}=\left(x_{1}, y_{1}\right), \ldots$, ( $x_{n}, y_{n}$ ) and $\epsilon$,

$$
\mathcal{N}\left(\epsilon, \mathcal{J}\left(z_{1}^{n}\right)\right) \leq \mathcal{N}\left(\frac{\epsilon}{4 b_{n}}, \mathcal{F}_{n}\left(x_{1}^{n}\right)\right)
$$

Therefore, it suffices to estimate the covering number corresponding to $\mathcal{F}_{n}$. Since $\mathcal{F}_{n}$ is a subset of a linear space of functions, we have $V_{\mathcal{F}_{n}^{+}} \leq k_{n}+1$ (Theorem 13.9). By Corollary 29.2,

$$
\mathcal{N}\left(\frac{\epsilon}{4 b_{n}}, \mathcal{F}_{n}\left(x_{1}^{n}\right)\right) \leq\left(\frac{8 e b_{n}}{\epsilon /\left(4 b_{n}\right)} \log \frac{4 e b_{n}}{\epsilon /\left(4 b_{n}\right)}\right)^{k_{n}+1} \leq\left(\frac{32 e b_{n}^{2}}{\epsilon}\right)^{2\left(k_{n}+1\right)}
$$

Summarizing, we have

$$
\begin{aligned}
\mathbf{P}\{\mathbf{E}\{ & \left(f_{n}^{*}(X)-(2 Y-1)\right)^{2}\left\{D_{n}\right\} \\
& \left.-\inf _{f \in \mathcal{F}_{n}} \mathbf{E}\left\{(f(X)-(2 Y-1))^{2}\right\}>\epsilon\right\} \\
\leq & 8\left(\frac{2^{10} e b_{n}^{2}}{\epsilon}\right)^{2\left(k_{n}+1\right)} e^{-n \epsilon^{2} /\left(2^{13} b_{n}^{4}\right)},
\end{aligned}
$$

which goes to zero if $k_{n} b_{n}^{4} \log \left(b_{n}\right) / n \rightarrow 0$. The proof of the theorem is completed.
It is easy to see that if we assume additionally that $b_{n}^{4} \log n / n \rightarrow 0$, then strong universal consistency follows by applying the Borel-Cantelli lemma to the last probability.

REMARK. Minimization of the squared error is attractive because there are efficient algorithms to find the minimizing coefficients, while minimizing the number of errors committed on the training sequence is computationally more difficult. If the dimension $k$ of the generalized linear classifier is fixed, then stochastic approximation asymptotically provides the minimizing coefficients. For more information about this we refer to Robbins and Monro (1951), Kiefer and Wolfowitz (1952), Dvoretzky (1956), Fabian (1971), Tsypkin (1971), Nevelson and Khasminskii (1973), Kushner (1984), Ruppert (1991), and Ljung, Pflug, and Walk (1992). For example, Györfi (1984) proved that if $\left(U_{1}, V_{1}\right),\left(U_{2}, V_{2}\right), \ldots$ form a stationary and ergodic sequence, in which each pair is distributed as the bounded random variable pair $(U, V) \in \mathcal{R}^{k} \times \mathcal{R}$, and the vector of coefficients $a=\left(a_{1}, \ldots, a_{k}\right)$ minimizes

$$
R(a)=\mathbf{E}\left\{\left(a^{T} U-V\right)^{2}\right\}
$$

and $a^{(0)} \in \mathcal{R}^{k}$ is arbitrary, then the sequence of coefficient vectors defined by

$$
a^{(n+1)}=a^{(n)}-\frac{1}{n+1}\left(a^{(n)^{T}} U_{n+1}-V_{n+1}\right) U_{n+1}
$$

satisfies

$$
\lim _{n \rightarrow \infty} R\left(a^{(n)}\right)=R(a) \text { a.s. }
$$

## Problems and Exercises

Problem 29.1. Find a class $\mathcal{F}$ containing two functions $\eta_{1}, \eta_{2}: \mathcal{R} \rightarrow[0,1]$ and a distribution of $(X, Y)$ such that $\min \left(L\left(\eta_{1}\right), L\left(\eta_{2}\right)\right)=L^{*}$, but as $n \rightarrow \infty$, the probability

$$
\mathbf{P}\left\{L\left(\eta_{n}\right)=\max \left(L\left(\eta_{1}\right), L\left(\eta_{2}\right)\right)\right\}
$$

converges to one, where $\eta_{n}$ is selected from $\mathcal{F}$ by minimizing the empirical squared error.

Problem 29.2. Prove Corollary 29.2.
Problem 29.3. Let $\mathcal{F}$ be a class of functions on $\mathcal{R}^{d}$, taking their values in [ $\left.0, M\right]$. Haussler (1992) defines the pseudo dimension of $\mathcal{F}$ as the largest integer $n$ for which there exist $n$ points in $\mathcal{R}^{d}, z_{1}, \ldots, z_{n}$, and a vector $v=\left(v^{(1)}, \ldots, v^{(n)}\right) \in \mathcal{R}^{n}$ such that the binary $n$-vector

$$
\left(I_{\left\{v^{(1)}+f\left(z_{1}\right)>0\right\}}, \ldots, I_{\left\{v^{(n)}+f\left(z_{n}\right)>0\right\}}\right)
$$

takes all $2^{n}$ possible values as $f$ ranges through $\mathcal{F}$. Prove that the pseudo dimension of $\mathcal{F}$ equals the quantity $V_{\mathcal{F}^{+}}$defined in the text.

Problem 29.4. Consistency of clustering. Let $X, X_{1}, \ldots, X_{n}$ be i.i.d. random variables in $\mathcal{R}^{d}$, and assume that there is a number $0<M<\infty$ such that $\mathbf{P}\left\{X \in[-M, M]^{d}\right\}=1$. Take the empirically optimal clustering of $X_{1}, \ldots, X_{n}$, that is, find the points $a_{1}, \ldots, a_{k}$ that minimize the empirical squared error:

$$
e_{n}\left(a_{1}, \ldots, a_{k}\right)=\frac{1}{n} \sum_{i=1}^{n} \min _{1 \leq j \leq k}\left\|a_{j}-X_{i}\right\|^{2}
$$

The error of the clustering is defined by the mean squared error

$$
e\left(a_{1}, \ldots, a_{k}\right)=\mathbf{E}\left\{\min _{1 \leq j \leq k}\left\|a_{j}-X\right\|^{2} \mid X_{1}, \ldots, X_{n}\right\}
$$

Prove that if $a_{1}, \ldots, a_{k}$ denote the empirically optimal cluster centers, then

$$
e\left(a_{1}, \ldots, a_{k}\right)-\inf _{b_{1}, \ldots, b_{k} \in \mathcal{R}^{d}} e\left(b_{1}, \ldots, b_{k}\right) \leq 2 \sup _{b_{1}, \ldots, b_{k} \in \mathcal{R}^{d}}\left|e_{n}\left(b_{1}, \ldots, b_{k}\right)-e\left(b_{1}, \ldots, b_{k}\right)\right|
$$

and that for every $\epsilon>0$

$$
\mathbf{P}\left\{\sup _{b_{1}, \ldots, b_{k} \in \mathcal{R}^{d}}\left|e_{n}\left(b_{1}, \ldots, b_{k}\right)-e\left(b_{1}, \ldots, b_{k}\right)\right|>\epsilon\right\} \leq 4 e^{8} n^{2 k(d+1)} e^{-n \epsilon^{2} /\left(32 M^{4}\right)} .
$$

Conclude that the error of the empirically optimal clustering converges to that of the truly optimal one as $n \rightarrow \infty$. (Pollard (1981; 1982), Linder, Lugosi, and Zeger (1994)). Hint: For the first part proceed as in the proof of Lemma 8.2. For the second part use the technique shown in Corollary 29.1. To compute the vc dimension, exploit Corollary 13.2.

Problem 29.5. Let $\psi_{1}, \ldots, \psi_{k}$ be indicator functions of cells of a $k$-way partition of $\mathcal{R}^{d}$. Consider generalized linear classifiers based on these functions. Show that the classifier obtained by minimizing the number of errors made on the training sequence is the same as for the classifier obtained by minimizing the empirical squared error. Point out that this is just the histogram classifier based on the partition defined by the $\psi_{i}$ 's (Csibi (1975)).

## 30

## Neural Networks

### 30.1 Multilayer Perceptrons

The linear discriminant or perceptron (see Chapter 4) makes a decision

$$
\phi(x)= \begin{cases}0 & \text { if } \psi(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

based upon a linear combination $\psi(x)$ of the inputs,

$$
\begin{equation*}
\psi(x)=c_{0}+\sum_{i=1}^{d} c_{i} x^{(i)}=c_{0}+c^{T} x \tag{30.1}
\end{equation*}
$$

where the $c_{i}$ 's are weights, $x=\left(x^{(1)}, \ldots, x^{(d)}\right)^{T}$, and $c=\left(c_{1}, \ldots, c_{d}\right)^{T}$. This is called a neural network without hidden layers (see Figure 4.1).

In a (feed-forward) neural network with one hidden layer, one takes

$$
\begin{equation*}
\psi(x)=c_{0}+\sum_{i=1}^{k} c_{i} \sigma\left(\psi_{i}(x)\right) \tag{30.2}
\end{equation*}
$$

where the $c_{i}$ 's are as before, and each $\psi_{i}$ is of the form given in (30.1): $\psi_{i}(x)=$ $b_{i}+\sum_{j=1}^{d} a_{i j} x^{(j)}$ for some constants $b_{i}$ and $a_{i j}$. The function $\sigma$ is called a sigmoid. We define sigmoids to be nondecreasing functions with $\sigma(x) \rightarrow-1$ as $x \downarrow-\infty$ and $\sigma(x) \rightarrow 1$ as $x \uparrow \infty$. Examples include:
(1) the threshold sigmoid

$$
\sigma(x)= \begin{cases}-1 & \text { if } x \leq 0 \\ 1 & \text { if } x>0\end{cases}
$$

(2) the standard, or logistic, sigmoid

$$
\sigma(x)=\frac{1-e^{-x}}{1+e^{-x}}
$$

(3) the arctan sigmoid

$$
\sigma(x)=\frac{2}{\pi} \arctan (x)
$$

(4) the gaussian sigmoid

$$
\sigma(x)=2 \int_{-\infty}^{x} \frac{1}{\sqrt{2 \pi}} e^{-u^{2} / 2} d u-1
$$



Figure 30.1. A neural network with one hidden layer. The hidden neurons are those within the frame.


Figure 30.2. The threshold, standard, arctan, and gaussian sigmoids.

For early discussion of multilayer perceptrons, see Rosenblatt (1962), Barron (1975), Nilsson (1965), and Minsky and Papert (1969). Surveys may be found in Barron and Barron (1988), Ripley (1993; 1994), Hertz, Krogh, and Palmer (1991), and Weiss and Kulikowski (1991).

In the perceptron with one hidden layer, we say that there are $k$ hidden neuronsthe output of the $i$-th hidden neuron is $u_{i}=\sigma\left(\psi_{i}(x)\right)$. Thus, (30.2) may be rewritten as

$$
\psi(x)=c_{0}+\sum_{i=1}^{k} c_{i} u_{i}
$$

which is similar in form to (30.1). We may continue this process and create multilayer feed-forward neural networks. For example, a two-hidden-layer perceptron uses

$$
\psi(x)=c_{0}+\sum_{i=1}^{l} c_{i} z_{i}
$$

where

$$
z_{i}=\sigma\left(d_{i 0}+\sum_{j=1}^{k} d_{i j} u_{j}\right), \quad 1 \leq i \leq l
$$

and

$$
u_{j}=\sigma\left(b_{j}+\sum_{i=1}^{d} a_{j i} x^{(i)}\right), \quad 1 \leq j \leq k
$$

and the $d_{i j}$ 's, $b_{j}$ 's, and $a_{j i}$ 's are constants. The first hidden layer has $k$ hidden neurons, while the second hidden layer has $l$ hidden neurons.


Figure 30.3. A feed-forward neural network with two hidden layers.

The step from perceptron to a one-hidden-layer neural network is nontrivial. We know that linear discriminants cannot possibly lead to universally consistent rules. Fortunately, one-hidden-layer neural networks yield universally consistent discriminants provided that we allow $k$, the number of hidden neurons, to grow unboundedly with $n$. The interest in neural networks is undoubtedly due to the possibility of implementing them directly via processors and circuits. As the hardware is fixed beforehand, one does not have the luxury to let $k$ become a function of $n$, and thus, the claimed universal consistency is a moot point. We will deal with both fixed architectures and variable-sized neural networks. Because of the universal consistency of one-hidden-layer neural networks, there is little theoretical gain in considering neural networks with more than one hidden layer. There may, however, be an information-theoretic gain as the number of hidden neurons needed to achieve the same performance may be substantially reduced. In fact, we will make a case for two hidden layers, and show that after two hidden layers, little is gained for classification.

For theoretical analysis, the neural networks are rooted in a classical theorem by Kolmogorov (1957) and Lorentz (1976) which states that every continuous function $f$ on $[0,1]^{d}$ can be written as

$$
f(x)=\sum_{i=1}^{2 d+1} F_{i}\left(\sum_{j=1}^{d} G_{i j}\left(x^{(j)}\right)\right)
$$

where the $G_{i j}$ 's and the $F_{i}$ 's are continuous functions whose form depends on $f$. We will see that neural networks approximate any measurable function with arbitrary precision, despite the fact that the form of the sigmoids is fixed beforehand.

As an example, consider $d=2$. The function $x^{(1)} x^{(2)}$ is rewritten as

$$
\frac{1}{4}\left(\left(x^{(1)}+x^{(2)}\right)^{2}-\left(x^{(1)}-x^{(2)}\right)^{2}\right)
$$

which is in the desired form. However, it is much less obvious how one would rewrite more general continuous functions. In fact, in neural networks, we approximate the $G_{i j}$ 's and $F_{i}$ 's by functions of the form $\sigma\left(b+a^{T} x\right)$ and allow the number of tunable coefficients to be high enough such that any continuous function may be represented -though no longer rewritten exactly in the form of Kolmogorov and Lorentz. We discuss other examples of approximations based upon such representations in a later section.


Figure 30.4. The general Kolmogorov-Lorentz representation of a continuous function.

### 30.2 Arrangements

A finite set $\mathcal{A}$ of hyperplanes in $\mathcal{R}^{d}$ partitions the space into connected convex polyhedral pieces of various dimensions. Such a partition $\mathcal{P}=\mathcal{P}(\mathcal{A})$ is called an arrangement. An arrangement is called simple if any $d$ hyperplanes of $\mathcal{A}$ have a unique point in common and if $d+1$ hyperplanes have no point in common.


Figure 30.5. An arrangement of five lines in the plane.


Figure 30.6. An arrangement classifier.

A simple arrangement creates polyhedral cells. Interestingly, the number of these cells is independent of the actual configuration of the hyperplanes. In particular,
the number of cells is exactly $2^{k}$ if $d \geq k$, and

$$
\sum_{i=0}^{d}\binom{k}{i} \quad \text { if } d<k
$$

where $|\mathcal{A}|=k$. For a proof of this, see Problem 22.1, or Lemma 1.2 of Edelsbrunner (1987). For general arrangements, this is merely an upper bound.

We may of course use arrangements for designing classifiers. We let $g_{\mathcal{A}}$ be the natural classifier obtained by taking majority votes over all $Y_{i}$ 's for which $X_{i}$ is in the same cell of the arrangement $\mathcal{P}=\mathcal{P}(\mathcal{A})$ as $x$.

All classifiers discussed in this section possess the property that they are invariant under linear transformations and universally consistent (in some cases, we assume that $X$ has a density, but that is only done to avoid messy technicalities).

If we fix $k$ and find that $\mathcal{A}$ with $|\mathcal{A}|=k$ for which the empirical error

$$
L_{n}^{(R)}=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{g_{\mathcal{A}}\left(X_{i}\right) \neq Y_{i}\right\}}
$$

is minimal, we obtain-perhaps at great computational expense-the empirical risk optimized classifier. There is a general theorem for such classifiers-see, for example, Corollary 23.2-the conditions of which are as follows:
(1) It must be possible to select a given sequence of $\mathcal{A}$ 's for which $L_{n}\left(g_{\mathcal{A}}\right)$ (the conditional probability of error with $g_{\mathcal{A}}$ ) tends to $L^{*}$ in probability. But if $k \rightarrow \infty$, we may align the hyperplanes with the axes, and create a cubic histogram, for which, by Theorem 6.2 , we have consistency if the grid expands to $\infty$ and the cell sizes in the grid shrink to 0 . Thus, as $k \rightarrow \infty$, this condition holds trivially.
(2) The collection $\mathcal{G}=\left\{g_{\mathcal{A}}\right\}$ is not too rich, in the sense that $n / \log \mathcal{S}(\mathcal{G}, n)$ $\rightarrow \infty$, where $\mathcal{S}(\mathcal{G}, n)$ denotes the shatter coefficient of $\mathcal{G}$, that is, the maximal number of ways $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)$ can be split by sets of the form

$$
\left(\bigcup_{A \in \mathcal{P}(\mathcal{A})} A \times\{0\}\right) \cup\left(\bigcup_{A \in \mathcal{P}(\mathcal{A})} A \times\{1\}\right)
$$

If $|\mathcal{A}|=1$, we know that $\mathcal{S}(\mathcal{G}, n) \leq 2\left(n^{d}+1\right)$ (see Chapter 13). For $|\mathcal{A}|=k$, a trivial upper bound is

$$
\left(2\left(n^{d}+1\right)\right)^{k}
$$

The consistency condition is fulfilled if $k=o(n / \log n)$.
We have
Theorem 30.1. The empirical-risk-optimized arrangement classifier based upon arrangements with $|\mathcal{A}| \leq k$ has $\mathbb{E}\left\{L_{n}\right\} \rightarrow L^{*}$ for all distributions if $k \rightarrow \infty$ and $k=o(n / \log n)$.

Arrangements can also be made from the data at hand in a simpler way. Fix $k$ points $X_{1}, \ldots, X_{k}$ in general position and look at all possible $\binom{k}{d}$ hyperplanes you can form with these points. These form your collection $\mathcal{A}$, which defines your arrangement. No optimization of any kind is performed. We take the natural classifier obtained by a majority vote within the cells of the partition over $\left(X_{k+1}, Y_{k+1}\right), \ldots,\left(X_{n}, Y_{n}\right)$.


Figure 30.7. Arrangement determined by $k=4$ data points on the plane.

Here we cannot apply the powerful consistency theorem mentioned above. Also, the arrangement is no longer simple. Nevertheless, the partition of space depends on the $X_{i}$ 's only, and thus Theorem 6.1 (together with Lemma 20.1) is useful. The rule thus obtained is consistent if $\operatorname{diam}(A(X)) \rightarrow 0$ in probability and the number of cells is $o(n)$, where $A(X)$ is the cell to which $X$ belongs in the arrangement. As the number of cells is certainly not more than

$$
\sum_{i=0}^{d}\binom{k^{\prime}}{i}
$$

where $k^{\prime}=\binom{k}{d}$, we see that the number of cells divided by $n$ tends to zero if

$$
k^{d^{2}} / n \rightarrow 0
$$

This puts a severe restriction on the growth of $k$. However, it is easy to prove the following:

Lemma 30.1. If $k \rightarrow \infty$, then $\operatorname{diam}(A(X)) \rightarrow 0$ in probability whenever $X$ has a density.

Proof. As noted in Chapter 20 (see Problem 20.6), the set of all $x$ for which for all $\epsilon>0$, we have $\mu\left(x+\epsilon Q_{i}\right)>0$ for all quadrants $Q_{1}, \ldots, Q_{2^{d}}$ having one vertex at $(0,0, \ldots, 0)$ and sides of length one, has $\mu$-measure one. For such $x$, if at least one of the $X_{i}$ 's $(i \leq k)$ falls in each of the $2^{d}$ quadrants $x+\epsilon Q_{i}$, then $\operatorname{diam}(A(x)) \leq 2 d \epsilon($ see Figure 30.8).


Figure 30.8. The diameter of the cell containing $x$ is less than $4 \epsilon$ if there is a data point in each of the four quadrants of size $\epsilon$ around $x$.

Therefore, for arbitrary $\epsilon>0$,

$$
\mathbf{P}\{\operatorname{diam}(A(x))>2 d \epsilon\} \leq 2^{d}\left(1-\min _{1 \leq i \leq 2^{d}} \mu\left(x+\epsilon Q_{i}\right)\right)^{k} \rightarrow 0 .
$$

Thus, by the Lebesgue dominated convergence theorem,

$$
\mathbf{P}\{\operatorname{diam}(A(X))>2 d \epsilon\} \rightarrow 0 .
$$

Theorem 30.2. The arrangement classifier defined above is consistent whenever $X$ has a density and

$$
k \rightarrow \infty, \quad k^{d^{2}}=o(n)
$$

The theorem points out that empirical error minimization over a finite set of arrangements can also be consistent. Such a set may be formed as the collection of arrangements consisting of hyperplanes through $d$ points of $X_{1}, \ldots, X_{k}$. As nothing new is added here to the discussion, we refer the reader to Problem 30.1.

So how do we deal with arrangements in a computer? Clearly, to reach a cell, we find for each hyperplane $A \in \mathcal{A}$ the side to which $x$ belongs. If $f(x)=a^{T} x+a_{0}$, then $f(x)>0$ in one halfplane, $f(x)=0$ on the hyperplane, and $f(x)<0$ in the other halfplane. If $\mathcal{A}=\left\{A_{1}, \ldots, A_{k}\right\}$, the vector $\left(I_{\left\{H_{1}(x)>0\right\}}, \ldots, I_{\left\{H_{k}(x)>0\right\}}\right)$ describes the cell to which $x$ belongs, where $H_{i}(x)$ is a linear function that is positive if $x$ is on one side of the hyperplane $A_{i}$, negative if $x$ is on the other side of $A_{i}$, and 0 if $x \in A_{i}$. A decision is thus reached in time $O(k d)$. More importantly, the whole process is easily parallelizable and can be pictured as a battery of perceptrons. It is easy to see that the classifier depicted in Figure 30.9 is identical to the arrangement classifier. In neural network terminology, the first hidden layer of neurons corresponds to just $k$ perceptrons (and has $k(d+1)$ weights or parameters, if you wish). The first layer outputs a $k$-vector of bits that pinpoints the precise location of $x$ in the cells of the arrangement. The second layer only assigns a class (decision) to each cell of the arrangement by firing up one neuron. It has $2^{k}$ neurons (for class assignments), but of course, in natural classifiers, these neurons do not require training or learning-the majority vote takes care of that.


Figure 30.9. Arrangement classifier realized by a two-hiddenlayer neural network. Each of the $2^{k}$ cells in the second hidden layer performs an "and" operation: the output of node "101" is 1 if its three inputs are 1,0, and 1, respectively. Otherwise its output is 0 . Thus, one and only one of the $2^{k}$ outputs is 1 .

If a more classical second layer is needed-without boolean operations-let $b=\left(b_{1}, \ldots, b_{k}\right)$ be the $k$-vector of bits seen at the output of the first layer. Assign a perceptron in the second layer to each region of the arrangement and define the output $z \in\{-1,1\}$ to be $\sigma\left(\sum_{i=1}^{k} c_{j} b_{j}-k+1 / 2\right)$, where $c_{j} \in\{-1,1\}$ are weights. For each region of the arrangement, we have a description in terms of $c=\left(c_{1}, \ldots, c_{k}\right)$. The argument of the sigmoid function is $1 / 2$ if $2 b_{j}-1=c_{j}$ for all $j$ and is negative otherwise. Hence $z=1$ if and only if $2 b-1=c$. Assume we now take a decision based upon the sign of

$$
\sum_{l} w_{l} z_{l}+w_{0}
$$

where the $w_{l}$ 's are weights and the $z_{l}$ 's are the outputs of the second hidden layer. Assume that we wish to assign class 1 to $s$ regions in the arrangement and class 0 to $t$ other regions. For a class 1 region $l$, set $w_{l}=1$, and for a class 0 region, set $w_{l}=-1$. Define $w_{0}=1+s-t$. Then, if $z_{j}=1, z_{i}=-1, i \neq j$,

$$
\sum_{l} w_{l} z_{l}+w_{0}=w_{j}+w_{0}-\sum_{i \neq j} w_{i}= \begin{cases}1 & \text { if } w_{j}=1 \\ -1 & \text { if } w_{j}=-1\end{cases}
$$



Figure 30.10. The second hidden layer of a two-hidden-layer neural network with threshold sigmoids in the first layer. For each $k$-vector of bits $b=\left(b_{1}, \ldots, b_{k}\right)$ at the output of the first layer, we may find a decision $g(b) \in\{0,1\}$. Now return to the two-hidden-layer network of Figure 30.9 and assign the values $g(b)$ to the neurons in the second hidden layer to obtain an equivalent network.

Thus, every arrangement classifier corresponds to a neural network with two hidden layers, and threshold units. The correspondence is also reciprocal. Assume someone shows a two-hidden-layer neural network with the first hidden layer as above-thus, outputs consist of a vector of $k$ bits-and with the second hidden layer consisting once again a battery of perceptrons (see Figure 30.10). Whatever happens in the second hidden layer, the decision is just a function of the configuration of $k$ input bits. The output of the first hidden layer is constant over each region of the arrangement defined by the hyperplanes given by the input weights of the units of the first layer. Thus, the neural network classifier with threshold units in the first hidden layer is equivalent to an arrangement classifier with the same number of hyperplanes as units in the first hidden layer. The equivalence with tree classifiers is described in Problem 30.2.

Of course, equivalence is only valid up to a certain point. If the number of neurons in the second layer is small, then neural networks are more restricted. This could be an advantage in training. However, the majority vote in an arrangement classifier avoids training of the second layer's neurons altogether, and offers at the same time an easier interpretation of the classifier. Conditions on consistency of general two-hidden-layer neural networks will be given in Section 30.4.

### 30.3 Approximation by Neural Networks

Consider first the class $\mathcal{C}^{(k)}$ of classifiers (30.2) that contains all neural network classifiers with the threshold sigmoid and $k$ hidden nodes in two hidden layers. The training data $D_{n}$ are used to select a classifier from $\mathcal{C}^{(k)}$. For good performance of the selected rule, it is necessary that the best rule in $\mathcal{C}^{(k)}$ has probability of error close to $L^{*}$, that is, that

$$
\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)-L^{*}
$$

is small. We call this quantity the approximation error. Naturally, for fixed $k$, the approximation error is positive for most distributions. However, for large $k$, it is expected to be small. The question is whether the last statement is true for all distributions of $(X, Y)$. We showed in the section on arrangements that the class of two-hidden-layer neural network classifiers with $m$ nodes in the first layer and $2^{m}$ nodes in the second layer contains all arrangement classifiers with $m$ hyperplanes. Therefore, for $k=m+2^{m}$, the class of all arrangement classifiers with $m$ hyperplanes is a subclass of $\mathcal{C}^{(k)}$. From this, we easily obtain the following approximation result:

Theorem 30.3. If ${ }^{(k)}$ is the class of all neural network classifiers with the threshold sigmoid and $k$ neurons in two hidden layers, then

$$
\lim _{k \rightarrow \infty} \inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)-L^{*}=0
$$

for all distributions of $(X, Y)$.
It is more surprising that the same property holds if $\mathcal{C}^{(k)}$ is the class of one-hidden-layer neural networks with $k$ hidden neurons, and an arbitrary sigmoid. More precisely, $\mathcal{C}^{(k)}$ is the class of classifiers

$$
\phi(x)= \begin{cases}0 & \text { if } \psi(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

where $\psi$ is as in (30.2).
By Theorem 2.2, we have

$$
L(\phi)-L^{*} \leq 2 \mathbf{E}\{|\psi(X)-\eta(X)|\}
$$

where $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$. Thus, $\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)-L^{*} \rightarrow 0$ as $k \rightarrow \infty$ if

$$
\mathbf{E}\left\{\left|\psi_{k}(x)-\eta(x)\right|\right\} \rightarrow 0
$$

for some sequence $\left\{\psi_{k}\right\}$ with $\phi_{k} \in \mathcal{C}^{(k)}$ for $\phi_{k}(x)=I_{\left\{\psi_{k}(x)>1 / 2\right\}}$. For universal consistency, we need only assure that the family of $\psi$ 's can approximate any $\eta$ in the $L_{1}(\mu)$ sense. In other words, the approximation error $\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)-L^{*}$ converges to zero if the class of functions $\psi$ is dense in $L_{1}(\mu)$ for every $\mu$. Another sufficient condition for this-but of course much too severe-is that the class
$\mathcal{F}$ of functions $\psi$ becomes dense in the $L_{\infty}$ (supremum-) norm in the space of continuous functions $C[a, b]^{d}$ on $[a, b]^{d}$, where $[a, b]^{d}$ denotes the hyperrectangle of $\mathcal{R}^{d}$ defined by its opposite vertices $a$ and $b$, for any $a$ and $b$.

Lemma 30.2. Assume that a sequence of classes of functions $\mathcal{F}_{k}$ becomes dense in the $L_{\infty}$ norm in the space of continuous functions $C[a, b]^{d}$ (where $[a, b]^{d}$ is the hyperrectangle of $\mathcal{R}^{d}$ defined by $\left.a, b\right)$. In other words, assume that for every $a, b \in \mathcal{R}^{d}$, and every bounded function $g$,

$$
\lim _{k \rightarrow \infty} \inf _{f \in \mathcal{F}_{k}} \sup _{x \in[a, b]^{d}}|f(x)-g(x)|=0
$$

Then for any distribution of $(X, Y)$,

$$
\lim _{k \rightarrow \infty} \inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)-L^{*}=0
$$

where $\mathcal{C}^{(k)}$ is the class of classifiers $\phi(x)=I_{\{\psi(x)>1 / 2\}}$ for $\psi \in \mathcal{F}_{k}$.
Proof. For fixed $\epsilon>0$, find $a, b$ such that $\mu\left([a, b]^{d}\right) \geq 1-\epsilon / 3$, where $\mu$ is the probability measure of $X$. Choose a continuous function $\widehat{\eta}$ vanishing off $[a, b]^{d}$ such that

$$
\mathbf{E}\{|\eta(X)-\widehat{\eta}(X)|\} \leq \frac{\epsilon}{6}
$$

Find $k$ and $f \in \mathcal{F}_{k}$ such that

$$
\sup _{x \in[a, b]^{d}}|f(x)-\widehat{\eta}(x)| \leq \frac{\epsilon}{6}
$$

For $\phi(x)=I_{\{f(x)>1 / 2\}}$, we have, by Theorem 2.2,

$$
\begin{aligned}
L(\phi)-L^{*} & \leq 2 \mathbf{E}\left\{|f(X)-\eta(X)| I_{\left\{X \in[a, b]^{d}\right\}}\right\}+\frac{\epsilon}{3} \\
& \leq 2 \mathbf{E}\left\{|f(X)-\widehat{\eta}(X)| I_{\left\{X \in[a, b]^{d}\right\}}\right\}+2 \mathbf{E}\{|\widehat{\eta}(X)-\eta(X)|\}+\frac{\epsilon}{3} \\
& \leq 2 \sup _{x \in[a, b]^{d}}|f(x)-\widehat{\eta}(x)|++2 \mathbf{E}\{|\widehat{\eta}(X)-\eta(X)|\}+\frac{\epsilon}{3} \\
& \leq \epsilon
\end{aligned}
$$

This text is basically about all such good families, such as families that are obtainable by summing kernel functions, and histogram families. The first results for approximation with neural networks with one hidden layer appeared in 1989, when Cybenko (1989), Hornik, Stinchcombe, and White (1989), and Funahashi (1989) proved independently that feedforward neural networks with one hidden layer are dense with respect to the supremum norm on bounded sets in the set of continuous functions. In other words, by taking $k$ large enough, every continuous function
on $\mathcal{R}^{d}$ can be approximated arbitrarily closely, uniformly over any bounded set by functions realized by neural networks with one hidden layer. For a survey of various denseness results we refer to Barron (1989) and Hornik (1993). The proof given here uses ideas of Chen, Chen, and Liu (1990). It uses the denseness of the class of trigonometric polynomials in the $L_{\infty}$ sense for $C[0,1]^{d}$ (this is a special case of the Stone-Weierstrass theorem; see Theorem A. 9 in the Appendix), that is, by functions of the form

$$
\sum_{i}\left(\alpha_{i} \cos \left(2 \pi a_{i}^{T} x\right)+\beta_{i} \sin \left(2 \pi b_{i}^{T} x\right)\right),
$$

where $a_{i}, b_{i}$ are integer-valued vectors of $\mathcal{R}^{d}$.
Theorem 30.4. For every continuous function $f:[a, b]^{d} \rightarrow \mathcal{R}$ and for every $\epsilon>0$, there exists a neural network with one hidden layer and function $\psi(x)$ as in (30.2) such that

$$
\sup _{x \in[a, b]^{d}}|f(x)-\psi(x)|<\epsilon
$$

Proof. We prove the theorem for the threshold sigmoid

$$
\sigma(x)= \begin{cases}-1 & \text { if } x \leq 0 \\ 1 & \text { if } x>0\end{cases}
$$

The extension to general nondecreasing sigmoids is left as an exercise (Problem 30.3). Fix $\epsilon>0$. We take the Fourier series approximation of $f(x)$. By the Stone-Weierstrass theorem (Theorem A.9), there exists a large positive integer $M$, nonzero real coefficients $\alpha_{1}, \ldots, \alpha_{M}, \beta_{1}, \ldots, \beta_{M}$, and integers $m_{i, j}$ for $i=1, \ldots, M, j=1, \ldots, d$, such that

$$
\sup _{x \in[a, b]^{d}}\left|\sum_{i=1}^{M}\left(\alpha_{i} \cos \left(\frac{\pi}{a} m_{i}^{T} x\right)+\beta_{i} \sin \left(\frac{\pi}{a} m_{i}^{T} x\right)\right)-f(x)\right|<\frac{\epsilon}{2},
$$

where $m_{i}=\left(m_{i, 1}, \ldots, m_{i, d}\right), i=1, \ldots, M$. It is clear that every continuous function on the real line that is zero outside some bounded interval can be arbitrarily closely approximated uniformly on the interval by one-dimensional neural networks, that is, by functions of the form

$$
\sum_{i=1}^{k} c_{i} \sigma\left(a_{i} x+b_{i}\right)+c_{0}
$$

Just observe that the indicator function of an interval $[b, c]$ may be written as $\sigma(x-b)+\sigma(-x+c)$. This implies that bounded functions such as sin and cos can be approximated arbitrarily closely by neural networks. In particular, there exist neural networks $u_{i}(x), v_{i}(x)$ with $i=1, \ldots, M$, (i.e., mappings from $\mathcal{R}^{d}$ to $\mathcal{R}$ ) such that

$$
\sup _{x \in[a, b]^{d}}\left|u_{i}(x)-\cos \left(\frac{\pi}{a} m_{i}^{T} x\right)\right|<\frac{\epsilon}{4 M\left|\alpha_{i}\right|}
$$

and

$$
\sup _{x \in[a, b]^{d}}\left|v_{i}(x)-\sin \left(\frac{\pi}{a} m_{i}^{T} x\right)\right|<\frac{\epsilon}{4 M\left|\beta_{i}\right|}
$$

Therefore, applying the triangle inequality we get
$\sup _{x \in[a, b]^{d}}\left|\sum_{i=1}^{M}\left(\alpha_{i} \cos \left(\frac{\pi}{a} m_{i}^{T} x\right)+\beta_{i} \sin \left(\frac{\pi}{a} m_{i}^{T} x\right)\right)-\sum_{i=1}^{M}\left(\alpha_{i} u_{i}(x)+\beta_{i} v_{i}(x)\right)\right|<\frac{\epsilon}{2}$.
Since the $u_{i}$ 's and $v_{i}$ 's are neural networks, their linear combination

$$
\psi(x)=\sum_{i=1}^{M}\left(\alpha_{i} u_{i}(x)+\beta_{i} v_{i}(x)\right)
$$

is a neural network too and, in fact,

$$
\begin{aligned}
& \sup _{x \in[a, b]^{d}}|f(x)-\psi(x)| \\
& \quad \leq \sup _{x \in[a, b]^{d}}\left|f(x)-\sum_{i=1}^{M}\left(\alpha_{i} \cos \left(\frac{\pi}{a} m_{i}^{T} x\right)+\beta_{i} \sin \left(\frac{\pi}{a} m_{i}^{T} x\right)\right)\right| \\
& \quad+\sup _{x \in[a, b]^{d}}\left|\sum_{i=1}^{M}\left(\alpha_{i} \cos \left(\frac{\pi}{a} m_{i}^{T} x\right)+\beta_{i} \sin \left(\frac{\pi}{a} m_{i}^{T} x\right)\right)-\psi(x)\right| \\
& \quad<\frac{2 \epsilon}{2}=\epsilon
\end{aligned}
$$

The convergence may be arbitrarily slow for some $f$. By restricting the class of functions, it is possible to obtain upper bounds for the rate of convergence. For an example, see Barron (1993). The following corollary of Theorem 30.4 is obtained via Lemma 30.2:

Corollary 30.1. Let $\mathcal{C}^{(k)}$ contain all neural network classifiers defined by networks of one hidden layer with $k$ hidden nodes, and an arbitrary sigmoid $\sigma$. Then for any distribution of $(X, Y)$,

$$
\lim _{k \rightarrow \infty} \inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)-L^{*}=0
$$

The above convergence also holds if the range of the parameters $a_{i j}, b_{i}, c_{i}$ is restricted to an interval $\left[-\beta_{k}, \beta_{k}\right]$, where $\lim _{k \rightarrow \infty} \beta_{k}=\infty$.

REMARK. It is also true that the class of one-hidden-layer neural networks with $k$ hidden neurons becomes dense in $L_{1}(\mu)$ for every probability measure $\mu$ on $\mathcal{R}^{d}$ as $k \rightarrow \infty$ (see Problem 30.4). Then Theorem 2.2 may be used directly to prove Corollary 30.1.

In practice, the network architecture (i.e., $k$ in our case) is given to the designer, who can only adjust the parameters $a_{i j}, b_{i}$, and $c_{i}$, depending on the data $D_{n}$. In this respect, the above results are only of theoretical interest. It is more interesting to find out how far the error probability of the chosen rule is from $\inf _{\phi \in C_{k}} L(\phi)$. We discuss this problem in the next few sections.

### 30.4 VC Dimension

Assume now that the data $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$ are used to tune the parameters of the network. To choose a classifier from $\mathcal{C}^{(k)}$, we focus on the difference between the probability of error of the selected rule and that of the best classifier in $\mathcal{C}^{(k)}$. Recall from Chapters 12 and 14 that the vC dimension $V_{\mathcal{C}^{(k)}}$ of the class $\mathcal{C}^{(k)}$ determines the performance of some learning algorithms. Theorem 14.5 tells us that no method of picking a classifier from $\mathcal{C}^{(k)}$ can guarantee better than $\Omega\left(\sqrt{V_{\mathcal{C}^{(k)}} / n}\right)$ performance uniformly for all distributions. Thus, for meaningful distribution-free performance guarantees, the sample size $n$ has to be significantly larger than the vC dimension. On the other hand, by Corollary 12.1, there exists a way of choosing the parameters of the network-namely, by minimization of the empirical error probability-such that the obtained classifier $\phi_{n}^{*}$ satisfies

$$
\mathbf{E}\left\{L\left(\phi_{n}^{*}\right)\right\}-\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi) \leq 16 \sqrt{\frac{V_{\mathcal{C}^{(k)}} \log n+4}{2 n}}
$$

for all distributions. On the other hand, if $V_{\mathcal{C}^{(k)}}=\infty$, then for any $n$ and any rule, some bad distributions exist that induce very large error probabilities (see Theorem 14.3).

We start with a universal lower bound on the VC dimension of networks with one hidden layer.

Theorem 30.5. (BAUM (1988)). Let $\sigma$ be an arbitrary sigmoid and consider the class $\mathcal{C}^{(k)}$ of neural net classifiers with $k$ nodes in one hidden layer. Then

$$
V_{\mathcal{C}^{(k)}} \geq 2\left\lfloor\frac{k}{2}\right\rfloor d
$$

Proof. We prove the statement for the threshold sigmoid, and leave the extension as an exercise (Problem 30.7). We need to show that there is a set of $n=2\lfloor k / 2\rfloor d$ points in $\mathcal{R}^{d}$ that can be shattered by sets of the form $\{x: \psi(x)>1 / 2\}$, where $\psi$ is a one-layer neural network of $k$ hidden nodes. Clearly, it suffices to prove this for even $k$. In fact, we prove more: if $k$ is even, any set of $n=k d$ points in general position can be shattered (points are in general position if no $d+1$ points fall on the same $d$-1-dimensional hyperplane). Let $\left\{x_{1}, \ldots, x_{n}\right\}$ be a set of $n=k d$ such points. For each subset of this set, we construct a neural network $\psi$ with $k$ hidden nodes such that $\psi\left(x_{i}\right)>1 / 2$ if and only if $x_{i}$ is a member of this subset. We may
assume without loss of generality that the cardinality of the subset to be picked out is at most $n / 2$, since otherwise we can use $1 / 2-\psi(x)$, where $\psi$ picks out the complement of the subset. Partition the subset into at most $n /(2 d)=k / 2$ groups, each containing at most $d$ points. For each such group, there exists a hyperplane $a^{T} x+b=0$ that contains these points, but no other point from $\left\{x_{1}, \ldots, x_{n}\right\}$. Moreover, there exists a small positive number $h$ such that $a^{T} x_{i}+b \in[-h, h]$ if and only if $x_{i}$ is among this group of at most $d$ points. Therefore, the simple network

$$
\sigma\left(a^{T} x+b+h\right)+\sigma\left(-a^{T} x-b+h\right)
$$

is larger than 0 on $x_{i}$ for exactly these $x_{i}$ 's. Denote the vectors $a$, and parameters $b, h$ obtained for the $k / 2$ groups by $a_{1}, \ldots, a_{k / 2}, b_{1}, \ldots, b_{k / 2}$, and $h_{1}, \ldots, h_{k / 2}$. Let $h=\min _{j \leq k / 2} h_{j}$. It is easy to see that

$$
\psi(x)=\sum_{j=1}^{k / 2}\left(\sigma\left(a_{j}^{T} x+b_{j}+h\right)+\sigma\left(-a_{j}^{T} x-b_{j}+h\right)\right)+\frac{1}{2}
$$

is larger than $1 / 2$ for exactly the desired $x_{i}$ 's. This network has $k$ hidden nodes.

Theorem 30.5 implies that there is no hope for good performance guarantees unless the sample size is much larger than $k d$. Recall Chapter 14, where we showed that $n \gg V_{\mathcal{C}^{(k)}}$ is necessary for a guaranteed small error probability, regardless of the method of tuning the parameters. Bartlett (1993) improved Theorem 30.5 in several ways. For example, he proved that

$$
V_{\mathcal{C}^{(k)}} \geq d \min \left(k, \frac{2^{d}}{d^{2} / 2+d+1}\right)+1
$$

Bartlett also obtained similar lower bounds for not fully connected networks-see Problem 30.9—and for two-hidden-layer networks.

Next we show that for the threshold sigmoid, the bound of Theorem 30.5 is tight up to a logarithmic factor, that is, the vC dimension is at most of the order of $k d \log k$.

Theorem 30.6. (BaUM AND HaUssler (1989)). Let $\sigma$ be the threshold sigmoid and let $\mathcal{C}^{(k)}$ be the class of neural net classifiers with $k$ nodes in the hidden layer. Then the shatter coefficients satisfy

$$
\begin{aligned}
\mathcal{S}\left(\mathcal{C}^{(k)}, n\right) & \leq\left(\sum_{i=0}^{d+1}\binom{n}{i}\right)^{k}\left(\sum_{i=0}^{k+1}\binom{n}{i}\right) \\
& \leq\left(\frac{n e}{d+1}\right)^{k(d+1)}\left(\frac{n e}{k+1}\right)^{k+1} \leq(n e)^{k d+2 k+1}
\end{aligned}
$$

which implies that for all $k, d \geq 1$,

$$
V_{\mathcal{C}^{(k)}} \leq(2 k d+4 k+2) \log _{2}(e(k d+2 k+1))
$$

Proof. Fix $n$ points $x_{1}, \ldots, x_{n} \in \mathcal{R}^{d}$. We bound the number of different values of the vector $\left(\phi\left(x_{1}\right), \ldots, \phi\left(x_{n}\right)\right)$ as $\phi$ ranges through $\mathcal{C}^{(k)}$. A node $j$ in the hidden layer realizes a dichotomy of the $n$ points by a hyperplane split. By Theorems 13.9 and 13.3, this can be done at most $\sum_{i=0}^{d+1}\binom{n}{i} \leq(n e /(d+1))^{d+1}$ different ways. The different splittings obtained at the $k$ nodes determine the $k$-dimensional input of the output node. Different choices of the parameters $c_{0}, c_{1}, \ldots, c_{k}$ of the output node determine different $k$-dimensional linear splits of the $n$ input vectors. This cannot be done in more than $\sum_{i=0}^{k+1}\binom{n}{i} \leq(n e /(k+1))^{k+1}$ different ways for a fixed setting of the $a_{i j}$ and $b_{i}$ parameters. This altogether yields at most

$$
\left(\sum_{i=0}^{d+1}\binom{n}{i}\right)^{k}\left(\sum_{i=0}^{k+1}\binom{n}{i}\right)
$$

different dichotomies of the $n$ points $x_{1}, \ldots, x_{n}$, as desired. The bound on the VC dimension follows from the fact that $V_{\mathcal{C}^{(k)}} \leq n$ if $\mathcal{S}\left(\mathcal{C}^{(k)}, n\right) \geq 2^{n}$. $\square$

For threshold sigmoids, the gap between the lower and upper bounds above is logarithmic in $k d$. Notice that the VC dimension is about the number of weights (or tunable parameters) $w=k d+2 k+1$ of the network. Surprisingly, Maass (1994) proved that for networks with at least two hidden layers, the upper bound has the right order of magnitude, that is, the vc dimension is $\Omega(w \log w)$. A simple application of Theorems 30.4 and 30.6 provides the next consistency result that was pointed out in Faragó and Lugosi (1993):

Theorem 30.7. Let $\sigma$ be the threshold sigmoid. Let $g_{n}$ be a classifier from $\mathcal{C}^{(k)}$ that minimizes the empirical error

$$
\widehat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi\left(X_{i}\right) \neq Y_{i}\right\}}
$$

over $\phi \in \mathcal{C}^{(k)}$. If $k \rightarrow \infty$ such that $k \log n / n \rightarrow 0$ as $n \rightarrow \infty$, then $g_{n}$ is strongly universally consistent, that is,

$$
\lim _{n \rightarrow \infty} L\left(g_{n}\right)=L^{*}
$$

with probability one for all distributions of $(X, Y)$.

Proof. By the usual decomposition into approximation and estimation errors,

$$
L\left(g_{n}\right)-L^{*}=\left(L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)\right)+\left(\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)-L^{*}\right)
$$

The second term on the right-hand side tends to zero by Corollary 30.1. For the estimation error, by Theorems 12.6 and 30.6,

$$
\mathbf{P}\left\{L\left(g_{n}\right)-\inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)>\epsilon\right\} \leq 8 \mathcal{S}\left(\mathcal{C}^{(k)}, n\right) e^{-n \epsilon^{2} / 128}
$$

$$
\leq 8(n e)^{(k d+2 k+1)} e^{-n \epsilon^{2} / 128}
$$

which is summable if $k=o(n / \log n)$.

The theorem assures us that if $\sigma$ is the threshold sigmoid, then a sequence of properly sized networks may be trained to asymptotically achieve the optimum probability of error, regardless of what the distribution of $(X, Y)$ is. For example, $k \approx \sqrt{n}$ will do the job. However, this is clearly not the optimal choice in the majority of cases. Since Theorem 30.6 provides suitable upper bounds on the vC dimension of each class $\mathcal{C}^{(k)}$, one may use complexity regularization as described in Chapter 18 to find a near-optimum size network.

Unfortunately, the situation is much less clear for more general, continuous sigmoids. The VC dimension then depends on the specific sigmoid. It is not hard to see that the vc dimension of $\mathcal{C}^{(k)}$ with an arbitrary nondecreasing sigmoid is always larger than or equal to that with the threshold sigmoid (Problem 30.8). Typically, the vc dimension of a class of such networks is significantly larger than that for the threshold sigmoid. In fact, it can even be infinite! Macintyre and Sontag (1993) demonstrated the existence of continuous, infinitely many times differentiable monotone increasing sigmoids such that the vc dimension of $\mathcal{C}^{(k)}$ is infinite if $k \geq 2$. Their sigmoids have little squiggles, creating the large variability. It is even more surprising that infinite VC dimension may occur for even smoother sigmoids, whose second derivative is negative for $x>0$ and positive for $x<0$. In Chapter 25 (see Problem 25.11) we basically proved the following result. The details are left to the reader (Problem 30.13).

Theorem 30.8. There exists a sigmoid $\sigma$ that is monotone increasing, continuous, concave on $(0, \infty)$, and convex on $(-\infty, 0)$, such that $V_{\mathcal{C}^{(k)}}=\infty$ for each $k \geq 8$.

We recall once again that infinite vc dimension implies that there is no hope of obtaining nontrivial distribution-free upper bounds on

$$
L\left(g_{n}\right)-\inf _{\phi \in C^{(k)}} L(\phi)
$$

no matter how the training sequence $D_{n}$ is used to select the parameters of the neural network. However, as we will see later, it is still possible to obtain universal consistency. Finiteness of the VC dimension has been proved for many types of sigmoids. Maass (1993) and Goldberg and Jerrum (1993) obtain upper bounds for piecewise polynomial sigmoids. The results of Goldberg and Jerrum (1993) apply for general classes parametrized by real numbers, e.g., for classes of neural networks with the sigmoid

$$
\sigma(x)= \begin{cases}1-1 /(2 x+2) & \text { if } x \geq 0 \\ 1 /(2-2 x) & \text { if } x<0\end{cases}
$$

Macintyre and Sontag (1993) prove $V_{\mathcal{C}^{(k)}}<\infty$ for a large class of sigmoids, which includes the standard, arctan, and gaussian sigmoids. While finiteness is useful, the
lack of an explicit tight upper bound on $V_{\mathcal{C}^{(k)}}$ prevents us from getting meaningful upper bounds on the performance of $g_{n}$, and also from applying the structural risk minimization of Chapter 18. For the standard sigmoid, and for networks with $k$ hidden nodes and $w$ tunable weights Karpinski and Macintyre (1994) recently reported the upper bound

$$
V \leq \frac{k w(k w-1)}{2}+w(1+2 k)+w(1+3 k) \log (3 w+6 k w+3)
$$

See also Shawe-Taylor (1994).
Unfortunately, the consistency result of Theorem 30.7 is only of theoretical interest, as there is no efficient algorithm to find a classifier that minimizes the empirical error probability. Relatively little effort has been made to solve this important problem. Faragó and Lugosi (1993) exhibit an algorithm that finds the empirically optimal network. However, their method takes time exponential in $k d$, which is intractable even for the smallest toy problems. Much more effort has been invested in the tuning of networks by minimizing the empirical squared error, or the empirical $L_{1}$ error. These problems are also computationally demanding, but numerous suboptimal hill-climbing algorithms have been used with some success. Most famous among these is the back propagation algorithm of Rumelhart, Hinton, and Williams (1986). Nearly all known algorithms that run in reasonable time may get stuck at local optima, which results in classifiers whose probability of error is hard to predict. In the next section we study the error probability of neural net classifiers obtained by minimizing empirical $L_{p}$ errors.

We end this section with a very simple kind of one-layer network. The committee machine (see, e.g., Nilsson (1965) and Schmidt (1994)) is a special case of a one-hidden-layer neural network of the form (30.2) with $c_{0}=0, c_{1}=\cdots=c_{k}=1$, and the threshold sigmoid.


Figure 30.11. The committee machine has fixed weights at the output of the hidden layer.

Committee machines thus use a majority vote over the outcomes of the hidden neurons. It is interesting that the lower bound of Theorem 30.5 remains valid when $\mathcal{C}^{(k)}$ is the class of all committee machines with $k$ neurons in the hidden
layer. It is less obvious, however, that the class of committee machines is large enough for the asymptotic property

$$
\lim _{k \rightarrow \infty} \inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)=L^{*}
$$

for all distributions of $(X, Y)$ (see Problem 30.6).


Figure 30.12. A partition of the plane determined by a committee machine with 5 hidden neurons. The total vote is shown in each region. The region in which we decide " 0 " is shaded.

## $30.5 \quad L_{1}$ Error Minimization

In the previous section, we obtained consistency for the standard threshold sigmoid networks by empirical risk minimization. We could not apply the same methodology for general sigmoids simply because the vc dimension for general sigmoidal networks is not bounded. It is bounded for certain classes of sigmoids, and for those, empirical risk minimization yields universally consistent classifiers. Even if the VC dimension is infinite, we may get consistency, but this must then be proved by other methods, such as methods based upon metric entropy and covering numbers (see Chapters 28 and 29, as well as the survey by Haussler (1992)). One could also train the classifier by minimizing another empirical criterion, which is exactly what we will do in this section. We will be rewarded with a general consistency theorem for all sigmoids.

For $1 \leq p<\infty$, the empirical $L_{p}$ error of a neural network $\psi$ is defined by

$$
J_{n}^{(p)}(\psi)=\left(\frac{1}{n} \sum_{i=1}^{n}\left|\psi\left(X_{i}\right)-Y_{i}\right|^{p}\right)^{1 / p}
$$

The most interesting cases are $p=1$ and $p=2$. For $p=2$ this is just the empirical squared error, while $p=1$ yields the empirical absolute error. Often it makes sense to attempt to choose the parameters of the network $\psi$ such that $J_{n}^{(p)}(\psi)$ is
minimized. In situations where one is not only interested in the number of errors, but also how robust the decision is, such error measures may be meaningful. In other words, these error measures penalize even good decisions if $\psi$ is close to the threshold value 0 . Minimizing $J_{n}^{(p)}$ is like finding a good regression function estimate. Our concern is primarily with the error probability. In Chapter 4 we already highlighted the dangers of squared error minimization and $L_{p}$ errors in general. Here we will concentrate on the consistency properties. We minimize the empirical error over a class of functions, which should not be too large to avoid overfitting. However, the class should be large enough to contain a good approximation of the target function. Thus, we let the class of candidate functions grow with the sample size $n$, as in Grenander's "method of sieves" (Grenander (1981)). Its consistency and rates of convergence have been widely studied primarily for least squares regression function estimation and nonparametric maximum likelihood density estimation-see Geman and Hwang (1982), Gallant (1987), and Wong and Shen (1992).

Remark. Regression function estimation. In the regression function estimation setup, White (1990) proved consistency of neural network estimates based on squared error minimization. Barron (1991; 1994) used a complexity-regularized modification of these error measures to obtain the fastest possible rate of convergence for nonparametric neural network estimates. Haussler (1992) provides a general framework for empirical error minimization, and provides useful tools for handling neural networks. Various consistency properties of nonparametric neural network estimates have been proved by White (1991), Mielniczuk and Tyrcha (1993), and Lugosi and Zeger (1995).

We only consider the $p=1$ case, as the generalization to other values of $p$ is straightforward. Define the $L_{1}$ error of a function $\psi: \mathcal{R}^{d} \rightarrow \mathcal{R}$ by

$$
J(\psi)=\mathbf{E}\{|\psi(X)-Y|\} .
$$

We pointed out in Problem 2.12 that one of the functions minimizing $J(\psi)$ is the Bayes rule $g^{*}$ whose error is denoted by

$$
J^{*}=\inf _{\psi} J(\psi)=J\left(g^{*}\right) .
$$

Then clearly, $J^{*}=L^{*}$. We have also seen that if we define a decision by

$$
g(x)= \begin{cases}0 & \text { if } \psi(x) \leq 1 / 2 \\ 1 & \text { otherwise, }\end{cases}
$$

then its error probability $L(g)=\mathbf{P}\{g(X) \neq Y\}$ satisfies the inequality

$$
L(g)-L^{*} \leq J(\psi)-J^{*} .
$$

Our approach is to select a neural network from a suitably chosen class of networks by minimizing the empirical error

$$
J_{n}(\psi)=\frac{1}{n} \sum_{i=1}^{n}\left|\psi\left(X_{i}\right)-Y_{i}\right| .
$$

Denoting this function by $\psi_{n}$, according to the inequality above, the classifier

$$
g_{n}(x)= \begin{cases}0 & \text { if } \psi_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

is consistent if the $L_{1}$ error

$$
J\left(\psi_{n}\right)=\mathbf{E}\left\{\left|\psi_{n}(X)-Y\right| \mid D_{n}\right\}
$$

converges to $J^{*}$ in probability. Convergence with probability one provides strong consistency. For universal convergence, the class over which the minimization is performed has to be defined carefully. The following theorem shows that this may be achieved by neural networks with $k$ nodes, in which the range of the output weights $c_{0}, c_{1}, \ldots, c_{k}$ is restricted.

Theorem 30.9. (LUGOSI AND ZEGER (1995)). Let $\sigma$ be an arbitrary sigmoid. Define the class $\mathcal{F}_{n}$ of neural networks by

$$
\mathcal{F}_{n}=\left\{\sum_{i=1}^{k_{n}} c_{i} \sigma\left(a_{i}^{T} x+b_{i}\right)+c_{0}: \quad a_{i} \in \mathcal{R}^{d}, b_{i} \in \mathcal{R}, \sum_{i=0}^{k_{n}}\left|c_{i}\right| \leq \beta_{n}\right\}
$$

and let $\psi_{n}$ be a function that minimizes the empirical $L_{1}$ error

$$
J_{n}(\psi)=\frac{1}{n} \sum_{i=1}^{n}\left|\psi\left(X_{i}\right)-Y_{i}\right|
$$

over $\psi \in \mathcal{F}_{n}$. If $k_{n}$ and $\beta_{n}$ satisfy

$$
\lim _{n \rightarrow \infty} k_{n}=\infty, \quad \lim _{n \rightarrow \infty} \beta_{n}=\infty, \quad \text { and } \lim _{n \rightarrow \infty} \frac{k_{n} \beta_{n}^{2} \log \left(k_{n} \beta_{n}\right)}{n}=0
$$

then the classification rule

$$
g_{n}(x)= \begin{cases}0 & \text { if } \psi_{n}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

is universally consistent.

Remark. Strong universal consistency may also be shown by imposing slightly more restrictive conditions on $k_{n}$ and $\beta_{n}$ (see Problem 30.16).

Proof. By the argument preceding the theorem, it suffices to prove that $J\left(\psi_{n}\right)-$ $J^{*} \rightarrow 0$ in probability. Write

$$
J\left(\psi_{n}\right)-J^{*}=\left(J\left(\psi_{n}\right)-\inf _{\psi \in \mathcal{F}_{n}} J(\psi)\right)+\left(\inf _{\psi \in \mathcal{F}_{n}} J(\psi)-J^{*}\right)
$$

To handle the approximation error-the second term on the right-hand side-let $\psi^{\prime} \in \mathcal{F}_{n}$ be a function such that

$$
\mathbf{E}\left\{\left|\psi^{\prime}(X)-g^{*}(X)\right|\right\} \leq \mathbb{E}\left\{\left|\psi(X)-g^{*}(X)\right|\right\}
$$

for each $\psi \in \mathcal{F}_{n}$. The existence of such a function may be seen by noting that $\mathbb{E}\left\{\left|\psi(X)-g^{*}(X)\right|\right\}$ is a continuous function of the parameters $a_{i}, b_{i}, c_{i}$ of the neural network $\psi$. Clearly,

$$
\begin{aligned}
\inf _{\psi \in \mathcal{F}_{n}} J(\psi)-J^{*} & \leq J\left(\psi^{\prime}\right)-J^{*} \\
& =\mathbf{E}\left\{\left|\psi^{\prime}(X)-Y\right|\right\}-\mathbf{E}\left\{\left|g^{*}(X)-Y\right|\right\} \\
& \leq \mathbf{E}\left\{\left|\psi^{\prime}(X)-g^{*}(X)\right|\right\}
\end{aligned}
$$

which converges to zero as $n \rightarrow \infty$, by Problem 30.4. We start the analysis of the estimation error by noting that as in Lemma 8.2, we have

$$
\begin{aligned}
J\left(\psi_{n}\right)-\inf _{\psi \in \mathcal{F}_{n}} J(\psi) & \leq 2 \sup _{\psi \in \mathcal{F}_{n}}\left|J(\psi)-J_{n}(\psi)\right| \\
& =2 \sup _{\psi \in \mathcal{F}_{n}}\left|\mathbf{E}\{|\psi(X)-Y|\}-\frac{1}{n} \sum_{i=1}^{n}\right| \psi\left(X_{i}\right)-Y_{i}| |
\end{aligned}
$$

Define the class $\mathcal{M}_{n}$ of functions on $\mathcal{R}^{d} \times\{0,1\}$ by

$$
\begin{aligned}
& \mathcal{M}_{n} \\
& =\left\{m(x, y)=\left|\sum_{i=1}^{k_{n}} c_{i} \sigma\left(a_{i}^{T} x+b_{i}\right)+c_{0}-y\right|: a_{i} \in \mathcal{R}^{d}, b_{i} \in \mathcal{R}, \sum_{i=0}^{k_{n}}\left|c_{i}\right| \leq \beta_{n}\right\} .
\end{aligned}
$$

Then the previous bound becomes

$$
2 \sup _{m \in \mathcal{M}_{n}}\left|\mathbf{E}\{m(X, Y)\}-\frac{1}{n} \sum_{i=1}^{n} m\left(X_{i}, Y_{i}\right)\right| .
$$

Such quantities may be handled by the uniform law of large numbers of Theorem 29.1, which applies to classes of uniformly bounded functions. Indeed, for each $m \in \mathcal{M}_{n}$

$$
\begin{aligned}
m(x, y) & =\left|\sum_{i=1}^{k_{n}} c_{i} \sigma\left(a_{i}^{T} x+b_{i}\right)+c_{0}-y\right| \\
& \leq 2 \max \left(\sum_{i=1}^{k_{n}} c_{i} \sigma\left(a_{i}^{T} x+b_{i}\right)+c_{0}, 1\right) \\
& \leq 2 \max \left(\sum_{i=1}^{k_{n}}\left|c_{i}\right|, 1\right) \\
& \leq 2 \beta_{n},
\end{aligned}
$$

if $n$ is so large that $\beta_{n} \geq 1$. For such $n$ 's, Theorem 29.1 states that

$$
\begin{aligned}
\mathbf{P}\left\{\sup _{m \in \mathcal{M}_{n}} \mid \mathbf{E}\{m(X, Y)\}\right. & \left.\left.-\frac{1}{n} \sum_{i=1}^{n} m\left(X_{i}, Y_{i}\right) \right\rvert\,>\epsilon\right\} \\
\leq & 8 \mathbf{E}\left\{\mathcal{N}\left(\epsilon / 8, \mathcal{M}_{n}\left(D_{n}\right)\right)\right\} e^{-n \epsilon^{2} /\left(512 \beta_{n}^{2}\right)}
\end{aligned}
$$

where $\mathcal{N}\left(\epsilon, \mathcal{M}_{n}\left(D_{n}\right)\right)$ denotes the $l_{1}$-covering number of the random set

$$
\mathcal{M}_{n}\left(D_{n}\right)=\left\{\left(m\left(X_{1}, Y_{1}\right), \ldots, m\left(X_{n}, Y_{n}\right)\right): m \in \mathcal{M}_{n}\right\} \subset \mathcal{R}^{n}
$$

defined in Chapter 29. All we need now is to estimate these covering numbers. Observe that for $m_{1}, m_{2} \in \mathcal{M}_{n}$ with $m_{1}(x, y)=\left|\psi_{1}(x)-y\right|$ and $m_{2}(x, y)=$ $\left|\psi_{2}(x)-y\right|$, for any probability measure $v$ on $\mathcal{R}^{d} \times\{0,1\}$,

$$
\int\left|m_{1}(x, y)-m_{2}(x, y)\right| \nu(d(x, y)) \leq \int\left|\psi_{1}(x)-\psi_{2}(x)\right| \mu(d x)
$$

where $\mu$ is the marginal of $v$ on $\mathcal{R}^{d}$. Therefore, it follows that $\mathcal{N}\left(\epsilon, \mathcal{M}_{n}\left(D_{n}\right)\right) \leq$ $\mathcal{N}\left(\epsilon, \mathcal{F}_{n}\left(X_{1}^{n}\right)\right)$, where $X_{1}^{n}=\left(X_{1}, \ldots, X_{n}\right)$. It means that an upper bound on the covering number of the class of neural networks $\mathcal{F}_{n}$ is also an upper bound on the quantity that interests us. This bounding may be done by applying lemmas from Chapters 13 and 30. Define the following three collections of functions:

$$
\begin{aligned}
\mathcal{G}_{1} & =\left\{a^{T} x+b ; a \in \mathcal{R}^{d}, b \in \mathcal{R}\right\} \\
\mathcal{G}_{2} & =\left\{\sigma\left(a^{T} x+b\right) ; a \in \mathcal{R}^{d}, b \in \mathcal{R}\right\} \\
\mathcal{G}_{3} & =\left\{c \sigma\left(a^{T} x+b\right) ; a \in \mathcal{R}^{d}, b \in \mathcal{R}, c \in\left[-\beta_{n}, \beta_{n}\right]\right\}
\end{aligned}
$$

By Theorem 13.9, the VC dimension of the class of sets

$$
\mathcal{G}_{1}^{+}=\left\{(x, t): t \leq \psi(x), \psi \in \mathcal{G}_{1}\right\}
$$

is $V_{\mathcal{G}_{1}^{+}} \leq d+2$. This implies by Lemma 29.5 that $V_{\mathcal{G}_{2}^{+}} \leq d+2$, so by Corollary 29.2, for any $x_{1}^{n}=\left(x_{1}, \ldots, x_{n}\right)$,

$$
\mathcal{N}\left(\epsilon, \mathcal{G}_{2}\left(x_{1}^{n}\right)\right) \leq 2\left(\frac{4 e}{\epsilon}\right)^{2(d+2)}
$$

where $\mathcal{G}_{2}\left(x_{1}^{n}\right)=\left\{z \in \mathcal{R}^{n}: z=\left(g\left(x_{1}\right), \ldots, g\left(x_{n}\right)\right), g \in \mathcal{G}_{2}\right\}$. Now, using similar notations, Theorem 29.7 allows us to estimate covering numbers of $\mathcal{G}_{3}\left(x_{1}^{n}\right)$ :

$$
\mathcal{N}\left(\epsilon, \mathcal{G}_{3}\left(x_{1}^{n}\right)\right) \leq \frac{4}{\epsilon} \mathcal{N}\left(\epsilon /\left(2 \beta_{n}\right), \mathcal{G}_{2}\left(x_{1}^{n}\right)\right) \leq\left(\frac{8 e \beta_{n}}{\epsilon}\right)^{2 d+5}
$$

if $\beta_{n}>2 / e$. Finally, we can apply Lemma 29.6 to obtain

$$
\begin{aligned}
\mathcal{N}\left(\epsilon, \mathcal{F}_{n}\left(x_{1}^{n}\right)\right) & \leq \frac{2 \beta_{n}\left(k_{n}+1\right)}{\epsilon} \mathcal{N}\left(\epsilon /\left(k_{n}+1\right), \mathcal{G}_{3}\left(x_{1}^{n}\right)\right)^{k_{n}} \\
& \leq\left(\frac{8 e\left(k_{n}+1\right) \beta_{n}}{\epsilon}\right)^{k_{n}(2 d+5)+1}
\end{aligned}
$$

Thus, substituting this bound into the probability inequality above, we get for $n$ large enough,

$$
\begin{aligned}
\mathbf{P}\left\{\sup _{\psi \in \mathcal{F}_{n}}|\mathbf{E}| \psi(X)-Y\left|-\frac{1}{n} \sum_{j=1}^{n}\right| \psi\left(X_{j}\right)-Y_{j}| |>\epsilon\right\} \\
\leq 8\left(\frac{64 e\left(k_{n}+1\right) \beta_{n}}{\epsilon}\right)^{k_{n}(2 d+5)+1} e^{-n \epsilon^{2} /\left(512 \beta_{n}^{2}\right)},
\end{aligned}
$$

which tends to zero if

$$
\frac{k_{n} \beta_{n}^{2} \log \left(k_{n} \beta_{n}\right)}{n} \rightarrow 0
$$

concluding the proof.

There are yet other ways to obtain consistency for general sigmoidal networks. We may restrict the network by discretizing the values of the coefficients in some way-thus creating a sieve with a number of members that is easy to enumerateand applying complexity regularization (Chapter 18). This is the method followed by Barron (1988; 1991).

### 30.6 The Adaline and Padaline

Widrow (1959) and Widrow and Hoff (1960) introduced the Adaline, and Specht ( 1967 ; 1990) studied polynomial discriminant functions such as the Padaline. Looked at formally, the discriminant function $\psi$ used in the decision $g(x)=$ $I_{\{\psi(x)>0\}}$ is of a polynomial nature, with $\psi$ consisting of sums of monomials like

$$
a\left(x^{(1)}\right)^{i_{1}} \cdots\left(x^{(d)}\right)^{i_{d}}
$$

where $i_{1}, \ldots, i_{d} \geq 0$ are integers, and $a$ is a coefficient. The order of a monomial is $i_{1}+\cdots+i_{d}$. Usually, all terms up to, and including those of order $r$ are included. Widrow's Adaline (1959) has $r=1$. The total number of monomials of order $r$ or less does not exceed $(r+1)^{d}$. The motivation for developing these discriminants is that only up to $(r+1)^{d}$ coefficients need to be trained and stored. In applications in which data continuously arrive, the coefficients may be updated on-line and the data can be discarded. This property is, of course, shared with standard neural networks. In most cases, order $r$ polynomial discriminants are not translation invariant. Minimizing a given criterion on-line is a phenomenal task, so Specht noted that training is not necessary if the $a$ 's are chosen so as to give decisions that are close to those of the kernel method with normal kernels.

For example, if $K(u)=e^{-u^{2} / 2}$, the kernel method picks a smoothing factor $h$
(such that $h \rightarrow 0$ and $n h^{d} \rightarrow \infty$; see Chapter 10) and uses

$$
\begin{aligned}
\psi(x) & =\frac{1}{n} \sum_{i=1}^{n}\left(2 Y_{i}-1\right) K\left(\frac{\left\|x-X_{i}\right\|}{h}\right) \\
& =\frac{1}{n} \sum_{i=1}^{n}\left(2 Y_{i}-1\right) e^{-x^{T} x /\left(2 h^{2}\right)} e^{x^{T} X_{i} / h^{2}} e^{-X_{i}^{T} X_{i} /\left(2 h^{2}\right)}
\end{aligned}
$$

The same decision is obtained if we use

$$
\psi(x)=\frac{1}{n} \sum_{i=1}^{n}\left(2 Y_{i}-1\right) e^{x^{T} X_{i} / h^{2}-X_{i}^{T} X_{i} /\left(2 h^{2}\right)}
$$

Now, approximate this by using Taylor's series expansion and truncating to the order $r$ terms. For example, the coefficient of $\left(x^{(1)}\right)^{i_{1}} \cdots\left(x^{(d)}\right)^{i_{d}}$ in the expansion of $\psi(x)$ would be, if $\sum_{j=1}^{d} i_{j}=i$,

$$
\frac{1}{i_{1}!\cdots i_{d}!} \cdot \frac{1}{h^{2 i}} \cdot \sum_{j=1}^{n}\left(2 Y_{j}-1\right)\left(X_{j}^{(1)}\right)^{i_{1}} \cdots\left(X_{j}^{(d)}\right)^{i_{d}} e^{-X_{j}^{T} X_{j} /\left(2 h^{2}\right)}
$$

These sums are easy to update on-line, and decisions are based on the sign of the order $r$ truncation $\psi_{r}$ of $\psi$. The classifier is called the Padaline. Specht notes that overfitting in $\psi_{r}$ does not occur due to the fact that overfitting does not occur for the kernel method based on $\psi$. His method interpolates between the latter method, generalized linear discrimination, and generalizations of the perceptron.

For fixed $r$, the Padaline defined above is not universally consistent (for the same reason linear discriminants are not universally consistent), but if $r$ is allowed to grow with $n$, the decision based on the sign of $\psi_{r}$ becomes universally consistent (Problem 30.14). Recall, however, that Padaline was not designed with a variable $r$ in mind.

### 30.7 Polynomial Networks

Besides Adaline and Padaline, there are several ways of constructing polynomial many-layered networks in which basic units are of the form $\left(x^{(1)}\right)^{i_{1}} \ldots$ $\left(x^{(k)}\right)^{i_{k}}$ for inputs $x^{(1)}, \ldots, x^{(k)}$ to that level. Pioneers in this respect are Gabor (1961), Ivakhnenko $(1968$; 1971) who invented the GMDH method-the group method of data handling-and Barron (1975). See also Ivakhnenko, Konovalenko, Tulupchuk, and Tymchenko (1968) and Ivakhnenko, Petrache, and Krasyts'kyy (1968). These networks can be visualized in the following way:


Figure 30.13. Simple polynomial network: Each $g_{i}$ represents a simple polynomial function of its input. In Barron's work (Barron (1975), Barron and Barron (1988)), the $g_{i}$ 's are sometimes 2-input elements of the form $g_{i}(x, y)=a_{i}+b_{i} x+c_{i} y+d_{i} x y$.

If the $g_{i}$ 's are Barron's quadratic elements, then the network shown in Figure 30.13 represents a particular polynomial of order 8 in which the largest degree of any $x^{(i)}$ is at most 4 . The number of unknown coefficients is 36 in the example shown in the figure, while in a full-fledged order-8 polynomial network, it would be much larger.

For training these networks, many strategies have been proposed by Barron and his associates. Ivakhnenko (1971), for example, trains one layer at a time and lets only the best neurons in each layer survive for use as input in the next layer.

It is easy to see that polynomial networks, even with only two inputs per node, and with degree in each cell restricted to two, but with an unrestricted number of layers, can implement any polynomial in $d$ variables. As the polynomials are dense in the $L_{\infty}$ sense on $C[a, b]^{d}$ for all $a, b \in \mathcal{R}^{d}$, we note by Lemma 30.2 that such networks include a sequence of classifiers approaching the Bayes error for any distribution. Consider several classes of polynomials:

$$
\mathcal{G}_{1}=\left\{\psi=a_{1} \psi_{1}+\cdots+a_{k} \psi_{k}\right\},
$$

where $\psi_{1}, \ldots, \psi_{k}$ are fixed monomials, but the $a_{i}$ 's are free coefficients.

$$
\mathcal{G}_{2}=\left\{\psi \in \mathcal{G}_{1}: \psi_{1}, \ldots, \psi_{k} \text { are monomials of order } \leq r\right\} .
$$

(The order of a monomial $\left(x^{(1)}\right)^{i_{1}} \cdots\left(x^{(d)}\right)^{i_{d}}$ is $i_{1}+\cdots+i_{d}$.)
$\mathcal{G}_{3}=\left\{\psi \in \mathcal{G}_{1}: \psi_{1}, \ldots, \psi_{k}\right.$ are monomials of order $\leq r$, but $k$ is not fixed $\}$.
As $k \rightarrow \infty, \mathcal{G}_{1}$ does not generally become dense in $C[a, b]^{d}$ in the $L_{\infty}$ sense unless we add $\psi_{k+1}, \psi_{k+2}, \ldots$ in a special way. However, $\mathcal{G}_{3}$ becomes dense as $r \rightarrow \infty$ by Theorem A. 9 and $\mathcal{G}_{2}$ becomes dense as both $k \rightarrow \infty$ and $r \rightarrow \infty$ (as $\mathcal{G}_{2}$ contains a subclass of $\mathcal{G}_{3}$ for a smaller $r$ depending upon $k$ obtained by including in $\psi_{1}, \ldots, \psi_{k}$ all monomials in increasing order).

The vc dimension of the class of classifiers based on $\mathcal{G}_{1}$ is not more than $k$ (see Theorem 13.9). The vc dimension of classifiers based upon $\mathcal{G}_{2}$ does not exceed
those of $\mathcal{G}_{3}$, which in turn is nothing but the number of possible monomials of order $\leq r$. A simple counting argument shows that this is bounded by $(r+1)^{d}$. See also Anthony and Holden (1993).

These simple bounds may be used to study the consistency of polynomial networks.

Let us take a fixed structure network in which all nodes are fixed-they have at most $k$ inputs with $k \geq 2$ fixed and represent polynomials of order $\leq r$ with $r \geq 2$ fixed. For example, with $r=2$, each cell with input $z_{1}, \ldots, z_{k}$ computes $\sum_{i} a_{i} \psi_{i}\left(z_{1}, \ldots, z_{k}\right)$, where the $a_{i}$ 's are coefficients and the $\psi_{i}$ 's are fixed monomials of order $r$ or less, and all such monomials are included. Assume that the layout is fixed and is such that it can realize all polynomials of order $\leq s$ on the input $x^{(1)}, \ldots, x^{(d)}$. One way of doing this is to realize all polynomials of order $\leq r$ by taking all $\binom{d}{k}$ possible input combinations, and to repeat at the second level with $\binom{\binom{d}{k}}{k}$ cells of neurons, and so forth for a total of $s / r$ layers of cells. This construction is obviously redundant but it will do for now. Then note that the VC dimension is not more than $(s+1)^{d}$, as noted above. If we choose the best coefficients in the cells by empirical risk minimization, then the method is consistent:

Theorem 30.10. In the fixed-structure polynomial network described above, if $L_{n}$ is the probability of error of the empirical risk minimizer, then $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ if $s \rightarrow \infty$ and $s=o\left(n^{1 / d}\right)$.

Proof. Apply Lemma 30.2 and Theorem 12.6.

Assume a fixed-structure network as above such that all polynomials of order $\leq s$ are realized plus some other ones, while the number of layers of cells is not more than $l$. Then the vC dimension is not more than $(r l+1)^{d}$ because the maximal order is not more than $l r$. Hence, we have consistency under the same conditions as above, that is, $s \rightarrow \infty$, and $l=o\left(n^{1 / d}\right)$. Similar considerations can now be used in a variety of situations.

### 30.8 Kolmogorov-Lorentz Networks and Additive Models

Answering one of Hilbert's famous questions, Kolmogorov (1957) and Lorentz (1976) (see also Sprecher (1965) and Hecht-Nielsen (1987)) obtained the following interesting representation of any continuous function on $[0,1]^{d}$.

Theorem 30.11. (Kolmogorov (1957); LORENTZ (1976)). Let $f$ be continuous on $[0,1]^{d}$. Then $f$ can be rewritten as follows: let $\delta>0$ be an arbitrary constant, and choose $0<\epsilon \leq \delta$ rational. Then

$$
f=\sum_{k=1}^{2 d+1} g\left(z_{k}\right)
$$

where $g: \mathcal{R} \rightarrow \mathcal{R}$ is a continuous function (depending upon $f$ and $\epsilon$ ), and each $z_{k}$ is rewritten as

$$
z_{k}=\sum_{j=1}^{d} \lambda^{k} \psi\left(x^{(j)}+\epsilon k\right)+k
$$

Here $\lambda$ is real and $\psi$ is monotonic and increasing in its argument. Also, both $\lambda$ and $\psi$ are universal (independent of $f$ ), and $\psi$ is Lipschitz: $|\psi(x)-\psi(y)| \leq c|x-y|$ for some $c>0$.

The Kolmogorov-Lorentz theorem states that $f$ may be represented by a very simple network that we will call the Kolmogorov-Lorentz network. What is amazing is that the first layer is fixed and known beforehand. Only the mapping $g$ depends on $f$. This representation immediately opens up new revenues of pursuitwe need not mix the input variables. Simple additive functions of the input variables suffice to represent all continuous functions.


Figure 30.14. The Kolmogorov-Lorentz network of Theorem 30.11.

To explain what is happening here, we look at the interleaving of bits to make one-dimensional numbers out of $d$-dimensional vectors. For the sake of simplicity, let $f:[0,1]^{2} \rightarrow \mathcal{R}$. Let

$$
x=0 . x_{1} x_{2} x_{3} \ldots \quad \text { and } \quad y=0 . y_{1} y_{2} y_{3} \ldots
$$

be the binary expansions of $x$ and $y$, and consider a representation for the function $f(x, y)$. The bit-interleaved number $z \in[0,1]$ has binary expansion

$$
z=0 . x_{1} y_{1} x_{2} y_{2} x_{3} y_{3} \ldots
$$

and may thus be written as $z=\phi(x)+(1 / 2) \phi(y)$, where

$$
\phi(x)=0 . x_{1} 0 x_{2} 0 x_{3} 0 \ldots
$$

and thus

$$
\frac{1}{2} \phi(y)=0.0 y_{1} 0 y_{2} 0 y_{3} \ldots
$$

We can also retrieve $x$ and $y$ from $z$ by noting that

$$
x=\psi_{1}(z), \quad \psi_{1}(z)=0 . z_{1} z_{3} z_{5} z_{7} \ldots,
$$

and

$$
y=\psi_{2}(z), \quad \psi_{2}(z)=0 . z_{2} z_{4} z_{6} z_{8} \ldots
$$

Therefore,

$$
\begin{aligned}
f(x, y) & =f\left(\psi_{1}(z), \psi_{2}(z)\right) \\
& \stackrel{\text { def }}{=} g(z) \quad \text { (a one-dimensional function of } z) \\
& =g\left(\phi(x)+\frac{1}{2} \phi(y)\right) .
\end{aligned}
$$

The function $\phi$ is strictly monotone increasing. Unfortunately, $\phi, \psi_{1}$, and $\psi_{2}$ are not continuous. Kolmogorov's theorem for this special case is as follows:

Theorem 30.12. (Kolmogorov (1957)). There exist five monotone functions $\phi_{1}, \ldots, \phi_{5}:[0,1] \rightarrow \mathcal{R}$ satisfying $\left|\phi_{i}\left(x_{1}\right)-\phi_{i}\left(x_{2}\right)\right| \leq\left|x_{1}-x_{2}\right|$, with the following property: for every $f \in C[0,1]^{2}$ (the continuous functions on $[0,1]^{2}$ ), there exists a continuous function $g$ such that for all $\left(x_{1}, x_{2}\right) \in[0,1]^{2}$,

$$
f\left(x_{1}, x_{2}\right)=\sum_{i=1}^{5} g\left(\phi_{i}\left(x_{1}\right)+\frac{1}{2} \phi_{i}\left(x_{2}\right)\right) .
$$

The difference with pure bit-interleaving is that now the $\phi_{i}$ 's are continuous and $g$ is continuous whenever $f$ is. Also, just as in our simle example, Kolmogorov gives an explicit construction for $\phi_{1}, \ldots, \phi_{5}$.

Kolmogorov's theorem may be used to show the denseness of certain classes of functions that may be described by networks. There is one pitfall however: any such result must involve at least one neuron or cell that has a general function in it, and we are back at square one, because a general function, even on only one input, may be arbitrarily complicated and wild.

Additive models include, for example, models such as

$$
\alpha+\sum_{i=1}^{d} \psi_{i}\left(x^{(i)}\right)
$$

where the $\psi_{i}$ 's are unspecified univariate functions (Friedman and Silverman (1989), Hastie and Tibshirani (1990)). These are not powerful enough to approximate all functions. A generalized additive model is

$$
\sigma\left(\alpha+\sum_{i=1}^{d} \psi_{i}\left(x^{(i)}\right)\right)
$$

(Hastie and Tibshirani (1990)), where $\sigma$ is now a given or unspecified function. From Kolmogorov's theorem, we know that the model

$$
\sum_{k=1}^{2 d+1} \psi\left(\alpha_{k}+\sum_{i=1}^{d} \psi_{i, k}\left(x^{(i)}\right)\right)
$$

with $\psi_{i, k}, \psi$ unspecified functions, includes all continuous functions on compact sets and is thus ideally suited for constructing networks. In fact, we may take $\alpha_{k}=k$ and take all $\psi_{i, k}$ 's as specified in Kolmogorov's theorem. This leaves only $\psi$ as the unknown.

Now, consider the following: any continuous univariate function $f$ can be approximated on bounded sets to within $\epsilon$ by simple combinations of threshold sigmoids $\sigma$ of the form

$$
\sum_{i=1}^{k} a_{i} \sigma\left(x-c_{i}\right)
$$

where $a_{i}, c_{i}, k$ are variable. This leads to a two-hidden-layer neural network representation related to that of Kůrková (1992), where only the last layer has unknown coefficients for a total of $2 k$.

Theorem 30.13. Consider a network classifier of the form described above in which $\sigma(\cdot)$ is the threshold sigmoid, and the $a_{i}$ 's and $c_{i}$ 's are found by empirical error minimization. Then $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ for all distributions of $(X, Y)$, if $k \rightarrow \infty$ and $k \log n / n \rightarrow 0$.

Proof. We will only outline the proof. First observe that we may approximate all functions on $C[a, b]^{d}$ by selecting $k$ large enough. By Lemma 30.2 and Theorem 12.6, it suffices to show that the VC dimension of our class of classifiers is $o(n / \log n)$. Considering $\alpha_{k}+\sum_{i=1}^{d} \psi_{i, k}\left(x^{(i)}\right)$ as new input elements, called $y_{k}$, $1 \leq k \leq 2 d+1$, we note that the VC dimension is not more than that of the classifiers based on

$$
\sum_{k=1}^{2 d+1} \sum_{i=1}^{k} a_{i} \sigma\left(y_{j}-c_{i}\right)
$$

which in turn is not more than that of the classifiers given by

$$
\sum_{i=1}^{k(2 d+1)} b_{l} \sigma\left(z_{l}-d_{l}\right)
$$

where $\left\{b_{l}\right\},\left\{d_{l}\right\}$ are parameters, and $\left\{z_{l}\right\}$ is an input sequence. By Theorem 13.9, the vc dimension is not more than $k(2 d+1)$. This concludes the proof.

For more results along these lines, we refer to Kůrková (1992).

### 30.9 Projection Pursuit

In projection pursuit (Friedman and Tukey (1974), Friedman and Stuetzle (1981), Friedman, Stuetzle, and Schroeder (1984), Huber (1985), Hall (1989), Flick, Jones, Priest, and Herman (1990)), one considers functions of the form

$$
\begin{equation*}
\psi(x)=\sum_{j=1}^{k} \psi_{j}\left(b_{j}+a_{j}^{T} x\right) \tag{30.3}
\end{equation*}
$$

where $b_{j} \in \mathcal{R}, a_{j} \in \mathcal{R}^{d}$ are constants and $\psi_{1}, \ldots, \psi_{k}$ are fixed functions. This is related to, but not a special case of, one-hidden-layer neural networks. Based upon the Kolmogorov-Lorentz representation theorem, we may also consider

$$
\begin{equation*}
\psi(x)=\sum_{j=1}^{d} \psi_{j}\left(x^{(j)}\right) \tag{30.4}
\end{equation*}
$$

for fixed functions $\psi_{j}$ (Friedman and Silverman (1989), Hastie and Tibshirani (1990)). In (30.4) and (30.3), the $\psi_{j}$ 's may be approximated in turn by spline functions or other nonparametric constructs. This approach is covered in the literature on generalized additive models (Stone (1985), Hastie and Tibshirani (1990)).

The class of functions $e^{a^{T} x}, a \in \mathcal{R}^{d}$, satisfies the conditions of the StoneWeierstrass theorem (Theorem A.9) and is therefore dense in the $L_{\infty}$ norm on $C[a, b]^{d}$ for any $a, b \in \mathcal{R}^{d}$ (see, e.g., Diaconis and Shahshahani (1984)).

As a corollary, we note that the same denseness result applies to the family

$$
\begin{equation*}
\sum_{i=1}^{k} \psi_{i}\left(a_{i}^{T} x\right) \tag{30.5}
\end{equation*}
$$

where $k \geq 1$ is arbitrary and $\psi_{1}, \psi_{2}, \ldots$ are general functions. The latter result is at the basis of projection pursuit methods for approximating functions, where one tries to find vectors $a_{i}$ and functions $\psi_{i}$ that approximate a given function very well.

REMARK. In some cases, approximations by functions as in (30.5) may be exact. For example,

$$
x^{(1)} x^{(2)}=\frac{1}{4}\left(x^{(1)}+x^{(2)}\right)^{2}-\frac{1}{4}\left(x^{(1)}-x^{(2)}\right)^{2}
$$

and

$$
\max \left(x^{(1)}, x^{(2)}\right)=\frac{1}{2}\left|x^{(1)}+x^{(2)}\right|+\frac{1}{2}\left|x^{(1)}-x^{(2)}\right|
$$

Theorem 30.14. (Diaconis and Shahshahani (1984)). Let $m$ be a positive integer. There are $\binom{m+d-1}{m}$ distinct vectors $a_{j} \in \mathcal{R}^{d}$ such that any homogeneous polynomial $f$ of order $m$ can be written as

$$
f(x)=\sum_{j=1}^{\substack{m+d-1 \\ m}} c_{j}\left(a_{j}^{T} x\right)^{m}
$$

for some real numbers $c_{j}$.
Every polynomial of order $m$ over $\mathcal{R}^{d}$ is a homogeneous polynomial of order $m$ over $\mathcal{R}^{d+1}$ by replacing the constant 1 by a component $x^{(d+1)}$ raised to an appropriate power. Thus, any polynomial of order $m$ over $\mathcal{R}^{d}$ may be decomposed exactly by

$$
f(x)=\sum_{j=1}^{\substack{m+d \\ m}} c_{j}\left(a_{j}^{T} x+b_{j}\right)^{m}
$$

for some real numbers $b_{j}, c_{j}$ and vectors $a_{j} \in \mathcal{R}^{d}$.
Polynomials may thus be represented exactly in the form $\sum_{i=1}^{k} \phi_{i}\left(a_{i}^{T} x\right)$ with $k=\binom{m+d}{m}$. As the polynomials are dense in $C[a, b]^{d}$, we have yet another proof that $\left\{\sum_{i=1}^{k} \phi_{i}\left(a_{i}^{T} x\right)\right\}$ is dense in $C[a, b]^{d}$. See Logan and Shepp (1975) or Logan (1975) for other proofs.

The previous discussion suggests at least two families from which to select a discriminant function. As usual, we let $g(x)=I_{\{\psi(x)>0\}}$ for a discriminant function $\psi$. Here $\psi$ could be picked from

$$
\mathcal{F}_{m}=\left\{\sum_{j=1}^{m} c_{j}\left(a_{j}^{T} x+b_{j}\right)^{m} ; b_{j}, c_{j} \in \mathcal{R}, a_{j} \in \mathcal{R}^{d}\right\}
$$

or

$$
\mathcal{F}_{m}^{\prime}=\left\{\sum_{j=1}^{m} c_{j} e^{a_{j}^{T} x} ; c_{j} \in \mathcal{R}, a_{j} \in \mathcal{R}^{d}\right\},
$$

where $m$ is sufficiently large. If we draw $\psi$ by minimizing the empirical error (admittedly at a tremendous computational cost), then convergence may result if $m$ is not too large. We need to know the vc dimension of the classes of classifiers corresponding to $\mathcal{F}_{m}$ and $\mathcal{F}_{m}^{\prime}$. Note that $\mathcal{F}_{m}$ coincides with all polynomials of order $m$ and each such polynomial is the sum of at most $(m+1)^{d}$ monomials. If we invoke Lemma 30.2 and Theorem 12.6, then we get

Theorem 30.15. Empirical risk minimization to determine $\left\{a_{j}, b_{j}, c_{j}\right\}$ in $\mathcal{F}_{m}$ leads to a universally consistent classifier provided that $m \rightarrow \infty$ and $m=o\left(n^{1 / d} / \log n\right)$.

Projection pursuit is very powerful and not at all confined to our limited discussion above. In particular, there are many other ways of constructing good consistent classifiers that do not require extensive computations such as empirical error minimization.

### 30.10 Radial Basis Function Networks

We may perform discrimination based upon networks with functions of the form

$$
\psi(x)=\sum_{i=1}^{k} a_{i} K\left(\frac{x-x_{i}}{h_{i}}\right)
$$

(and decision $g(x)=I_{\{\psi(x)>0\}}$ ), where $k$ is an integer, $a_{1}, \ldots, a_{k} h_{1}, \ldots, h_{k}$ are constants, $x_{1}, \ldots, x_{k} \in \mathcal{R}^{d}$, and $K$ is a kernel function (such as $K(u)=e^{-\|u\|^{2}}$ or $\left.K(u)=1 /\left(1+\|u\|^{2}\right)\right)$. In this form, $\psi$ covers several well-known methodologies:
(1) The kernel rule (Chapter 10): Take $k=n, a_{i}=2 Y_{i}-1, h_{i}=h, x_{i}=X_{i}$. With this choice, for a large class of kernels, we are guaranteed convergence if $h \rightarrow 0$ and $n h^{d} \rightarrow \infty$. This approach is attractive as no difficult optimization problem needs to be solved.
(2) The potential function method. In Bashkirov, Braverman, and Muchnik (1964), the parameters are $k=n, h_{i}=h, x_{i}=X_{i}$. The weights $a_{i}$ are picked to minimize the empirical error on the data, and $h$ is held fixed. The original kernel suggested there is $K(u)=1 /\left(1+\|u\|^{2}\right)$.
(3) Linear discrimination. For $k=2, K(u)=e^{-\|u\|^{2}}, h_{i} \equiv h, a_{1}=1, a_{2}=-1$, the set $\{x: \psi(x)>0\}$ is a linear halfspace. This, of course, is not universally consistent. Observe that the separating hyperplane is the collection of all points $x$ at equal distance from $x_{1}$ and $x_{2}$. By varying $x_{1}$ and $x_{2}$, all hyperplanes may be obtained.
(4) Radial basis function (RBF) neural networks (e.g., Powell (1987), Broomhead and Lowe (1988), Moody and Darken (1989), Poggio and Girosi (1990), Xu, Krzyżak, and Oja (1993), Xu, Krzyżak, and Yuille (1994), and Krzyżak, Linder, and Lugosi (1993)). An even more general function $\psi$ is usually employed here:

$$
\psi(x)=\sum_{i=1}^{k} c_{i} K\left(\left(x-x_{i}\right)^{T} A_{i}\left(x-x_{i}\right)\right)+c_{0}
$$

where the $A_{i}$ 's are tunable $d \times d$ matrices.
(5) Sieve methods. Grenander (1981) and Geman and Hwang (1982) advocate the use of maximum likelihood methods to find suitable values for the tunable parameters in $\psi$ (for $k, K$ fixed beforehand) subject to certain compactness constraints on these parameters to control the abundance of choices one may have. If we were to use empirical error minimization, we would find, if $k \geq n$, that all data points can be correctly classified (take the $h_{i}$ 's small enough, set $a_{i}=2 Y_{i}-1, x_{i}=X_{i}, k=n$ ), causing overfitting. Hence, $k$ must be smaller than $n$ if parameters are picked in this manner. Practical ways of choosing the parameters are discussed by Kraaijveld and Duin (1991),
and Chou and Chen (1992). In both (4) and (5), the $x_{i}$ 's may be thought of as representative prototypes, the $a_{i}$ 's as weights, and the $h_{i}$ 's as the radii of influence. As a rule, $k$ is much smaller than $n$ as $x_{1}, \ldots, x_{k}$ summarizes the information present at the data.

To design a consistent RBF neural network classifier, we may proceed as in (1). We may also take $k \rightarrow \infty$ but $k=o(n)$. Just let $\left(x_{1}, \ldots, x_{k}\right) \equiv\left(X_{1}, \ldots, X_{k}\right)$, $\left(a_{1}, \ldots, a_{k} \equiv 2 Y_{1}-1, \ldots, 2 Y_{k}-1\right)$, and choose $A_{i}$ or $h_{i}$ to minimize a given error criterion based upon $X_{k+1}, \ldots, X_{n}$, such as

$$
\hat{L}_{n}(g)=\frac{1}{n-k} \sum_{i=k+1}^{n} I_{\left\{g\left(X_{i}\right) \neq Y_{i}\right\}},
$$

where $g(x)=I_{\{\psi(x)>0\}}$, and $\psi$ is as in (1). This is nothing but data splitting (Chapter 22). Convergence conditions are described in Theorem 22.1.

A more ambitious person might try empirical risk minimization to find the best $x_{1}, \ldots, x_{k}, a_{1}, \ldots, a_{k}, A_{1}, \ldots, A_{k}(d \times d$ matrices as in (4)) based upon

$$
\frac{1}{n} \sum_{i=1}^{n} I_{\left\{g\left(X_{i}\right) \neq Y_{i}\right\}} .
$$

If $k \rightarrow \infty$, the class of rules contains a consistent subsequence, and therefore, it suffices only to show that the vc dimension is $o(n / \log n)$. This is a difficult task and some kernels yield infinite vc dimension, even if $d=1$ and $k$ is very small (see Chapter 25). However, there is a simple argument if $K=I_{R}$ for a simple set $R$. Let

$$
\mathcal{A}=\left\{\{x: x=a+A y, y \in R\}: a \in \mathcal{R}^{d}, A \text { a } d \times d \text { matrix }\right\} .
$$

If $R$ is a sphere, then $\mathcal{A}$ is the class of all ellipsoids. The number of ways of shattering a set $\left\{x_{1}, \ldots, x_{n}\right\}$ by intersecting with members from $\mathcal{A}$ is not more than

$$
n^{d(d+1) / 2+1}
$$

(see Theorem 13.9, Problem 13.10). The number of ways of shattering a set $\left\{x_{1}, \ldots, x_{n}\right\}$ by intersecting with sets of the form

$$
\left\{a_{1} I_{R_{1}}+\cdots+a_{k} I_{R_{k}}>0: R_{1}, \ldots, R_{k} \in \mathcal{A}, a_{1}, \ldots, a_{k} \in \mathcal{R}\right\}
$$

is not more than the product of all ways of shattering by intersections with $R_{1}$, with $R_{2}$, and so forth, that is,

$$
n^{k(d(d+1) / 2+1)} .
$$

The logarithm of the shatter coefficient is $o(n)$ if $k=o(n / \log n)$. Thus, by Corollary 12.1, we have

Theorem 30.16. (KRZYŻAK, Linder, And Lugosi (1993)). If we take $k \rightarrow \infty$, $k=o(n / \log n)$ in the RBF classifier (4) in which $K=I_{R}, R$ being the unit ball of
$\mathcal{R}^{d}$, and in which all the parameters are chosen by empirical risk minimization, then $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ for all distributions of $(X, Y)$. Furthermore, if $\mathcal{G}_{k}$ is the class of all RBF classifiers with $k$ prototypes,

$$
\mathbf{E}\left\{L_{n}\right\}-\inf _{g \in \mathcal{G}_{k}} L(g) \leq 16 \sqrt{\frac{k(d(d+1) / 2+1) \log n+\log (8 e)}{2 n}}
$$

The theorem remains valid (with modified constants in the error estimate) when $R$ is a hyperrectangle or polytope with a bounded number of faces. However, for more general $K$, the vc dimension is more difficult to evaluate.

For general kernels, consistent RBF classifiers can be obtained by empirical $L_{1}$ or $L_{2}$ error minimization (Problem 30.32). However, no efficient practical algorithms are known to compute the minima.

Finally, as suggested by Chou and Chen (1992) and Kraaijveld and Duin (1991), it is a good idea to place $x_{1}, \ldots, x_{k}$ by $k$-means clustering or another clustering method and to build an RBF classifier with those values or by optimization started at the given cluster centers.

## Problems and Exercises

Problem 30.1. Let $k, l$ be integers, with $d \leq k<n$, and $l \leq\binom{ k}{d}$. Assume $X$ has a density. Let $\mathcal{A}$ be a collection of hyperplanes drawn from the $\binom{k}{d}$ possible hyperplanes through $d$ points of $\left\{X_{1}, \ldots, X_{k}\right\}$, and let $g_{\mathcal{A}}$ be the corresponding natural classifier based upon the arrangement $\mathcal{P}(\mathcal{A})$. Take $l$ such collections $\mathcal{A}$ at random and with replacement, and pick the best $\mathcal{A}$ by minimizing $\widehat{L}_{n}\left(g_{\mathcal{A}}\right)$, where

$$
\widehat{L}_{n}\left(g_{\mathcal{A}}\right)=\frac{1}{n-k} \sum_{i=k+1}^{n} I_{\left\{Y_{i} \neq \xi_{\mathcal{A}}\left(X_{i}\right)\right\}} .
$$

Show that the selected classifier is consistent if $l \rightarrow \infty, l^{d}=o(n), n /(l \log k) \rightarrow \infty$. (Note: this is applicable with $k=\lfloor n / 2\rfloor$, that is, half the sample is used to define $\mathcal{A}$, and the other half is used to pick a classifier empirically.)

Problem 30.2. We are given a tree classifier with $k$ internal linear splits and $k+1$ leaf regions (a BSP tree). Show how to combine the neurons in a two-hidden-layer perceptron with $k$ and $k+1$ hidden neurons in the two hidden layers so as to obtain a decision that is identical to the tree-based classifier (Brent (1991), Sethi (1990; 1991)). For more on the equivalence of decision trees and neural networks, see Meisel (1990), Koutsougeras and Papachristou (1989), or Golea and Marchand (1990). Hint: Mimic the argument for arrangements in text.

Problem 30.3. Extend the proof of Theorem 30.4 so that it includes any nondecreasing sigmoid with $\lim _{x \rightarrow-\infty} \sigma(x)=-1$ and $\lim _{x \rightarrow \infty} \sigma(x)=1$. Hint: If $t$ is large, $\sigma(t x)$ approximates the threshold sigmoid.

Problem 30.4. This exercise states denseness in $L_{1}(\mu)$ for any probability measure $\mu$. Show that for every probability measure $\mu$ on $\mathcal{R}^{d}$, every measurable function $f: \mathcal{R}^{d} \rightarrow \mathcal{R}$
with $\int|f(x)| \mu(d x)<\infty$, and every $\epsilon>0$, there exists a neural network with one hidden layer and function $\psi(x)$ as in (30.2) such that

$$
\int|f(x)-\psi(x)| \mu(d x)<\epsilon
$$

(Hornik (1991)). Hint: Proceed as in Theorem 30.4, considering the following:
(1) Approximate $f$ in $L_{1}(\mu)$ by a continuous function $g(x)$ that is zero outside some bounded set $B \subset \mathcal{R}^{d}$. Since $g(x)$ is bounded, its maximum $\beta=\max _{x \in \mathcal{R}^{d}}|g(x)|$ is finite.
(2) Now, choose $a$ to be a positive number large enough so that both $B \subset[-a, a]^{d}$ and $\mu\left([-a, a]^{d}\right)$ is large.
(3) Extend the restriction of $g(x)$ to $[-a, a]^{d}$ periodically by tiling over all of $\mathcal{R}^{d}$. The obtained function, $\widehat{g}(x)$ is a good approximation of $g(x)$ in $L_{1}(\mu)$.
(4) Take the Fourier series approximation of $\widehat{g}(x)$, and use the Stone-Weierstrass theorem as in Theorem 30.4.
(5) Observe that every continuous function on the real line that is zero outside some bounded interval can be arbitrarily closely approximated uniformly over the whole real line by one-dimensional neural networks. Thus, bounded functions such as the sine and cosine functions can be approximated arbitrarily closely by neural networks in $L_{1}(\nu)$ for any probability measure $v$ on $\mathcal{R}$.
(6) Apply the triangle inequality to finish the proof.

Problem 30.5. Generalize the previous exercise for denseness in $L_{p}(\mu)$. More precisely, let $1 \leq p<\infty$. Show that for every probability measure $\mu$ on $\mathcal{R}^{d}$, every measurable function $f: \mathcal{R}^{d} \rightarrow \mathcal{R}$ with $\int|f(x)|^{p} \mu(d x)<\infty$, and every $\epsilon>0$, there exists a neural network with one hidden layer $h(x)$ such that

$$
\left(\int|f(x)-h(x)|^{p} \mu(d x)\right)^{1 / p}<\epsilon
$$

(Hornik (1991)).
Problem 30.6. Committee machines. Let $\mathcal{C}^{(k)}$ be the class of all committee machines. Prove that for all distributions of $(X, Y)$,

$$
\lim _{k \rightarrow \infty} \inf _{\phi \in \mathcal{C}^{(k)}} L(\phi)=L^{*}
$$

Hint: For a one-hidden-layer neural network with coefficients $c_{i}$ in (30.2), approximate the $c_{i}$ 's by discretization (truncation to a grid of values), and note that $c_{i} \psi_{i}(x)$ may thus be approximated in a committee machine by a sufficient number of identical copies of $\psi_{i}(x)$. This only forces the number of neurons to be a bit larger.

Problem 30.7. Prove Theorem 30.5 for arbitrary sigmoids. Hint: Approximate the threshold sigmoid by $\sigma(t x)$ for a sufficiently large $t$.

Problem 30.8. Let $\sigma$ be a nondecreasing sigmoid with $\sigma(x) \rightarrow-1$ if $x \rightarrow-\infty$ and $\sigma(x) \rightarrow 1$ if $x \rightarrow \infty$. Denote by $\mathcal{C}_{\sigma}^{(k)}$ the class of corresponding neural network classifiers with $k$ hidden layers. Show that $V_{\mathcal{C}_{\alpha}^{(k)}} \geq V_{\mathcal{C}^{(k)}}$, where $\mathcal{C}^{(k)}$ is the class corresponding to the threshold sigmoid.

Problem 30.9. This exercise generalizes Theorem 30.5 for not fully connected neural networks with one hidden layer. Consider the class $\mathcal{C}^{(k)}$ of one-hidden-layer neural networks with the threshold sigmoid such that each of the $k$ nodes in the hidden layer are connected to $d_{1}, d_{2}, \ldots, d_{k}$ inputs, where $1 \leq d_{i} \leq d$. More precisely, $\mathcal{C}^{(k)}$ contains all classifiers based on functions of the form

$$
\psi(x)=c_{0}+\sum_{i=1}^{k} c_{i} \sigma\left(\psi_{i}(x)\right), \quad \text { where } \quad \psi_{i}(x)=\sum_{j=1}^{d_{i}} a_{j} x^{\left(m_{i, j}\right)},
$$

where for each $i,\left(m_{i, 1}, \ldots, m_{i, d_{i}}\right)$ is a vector of distinct positive integers not exceeding $d$. Show that

$$
V_{\mathcal{C}^{(k)}} \geq \sum_{i=1}^{k} d_{i}+1
$$

if the $a_{j}$ 's, $c_{i}$ 's and $m_{i, j}$ 's are the tunable parameters (Bartlett (1993)).
Problem 30.10. Let $\sigma$ be a sigmoid that takes $m$ different values. Find upper bounds on the vC dimension of the class $\mathcal{C}_{\sigma}^{(k)}$.

Problem 30.11. Consider a one-hidden-layer neural network $\psi$. If $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ are fixed and all $X_{i}$ 's are different, show that with $n$ hidden neurons, we are always able to tune the weights such that $Y_{i}=\psi\left(X_{i}\right)$ for all $i$. (This remains true if $Y \in \mathcal{R}$ instead of $Y \in\{0,1\}$.) The property above describes a situation of overfitting that occurs when the neural network becomes too "rich"-recall also that the vc dimension, which is at least $d$ times the number of hidden neurons, must remain smaller than $n$ for any meaningful training.

Problem 30.12. The Bernstein perceptron. Consider the following perceptron for onedimensional data:

$$
\phi(x)= \begin{cases}1 & \text { if } \sum_{i=1}^{k} a_{i} x^{i}(1-x)^{k-i}>1 / 2 \\ 0 & \text { otherwise } .\end{cases}
$$

Let us call this the Bernstein perceptron since it involves Bernstein polynomials. If $n$ data points are collected, how would you choose $k$ (as a function of $n$ ) and how would you adjust the weights (the $a_{i}$ 's) to make sure that the Bernstein perceptron is consistent for all distributions of $(X, Y)$ with $\mathbf{P}\{X \in[0,1]\}=1$ ? Can you make the Bernstein perceptron consistent for all distributions of $(X, Y)$ on $\mathcal{R} \times\{0,1\}$ ?


Figure 30.15. The Bernstein perceptron.

Problem 30.13. Use the ideas of Section 25.5 and Problem 25.11 to prove Theorem 30.8.

Problem 30.14. Consider Specht's Padaline with $r=r_{n} \uparrow \infty$. Let $h=h_{n} \downarrow 0$, and $n h^{d} \rightarrow \infty$. Show that for any distribution of $(X, Y), \mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$.

Problem 30.15. Bounded first layers. Consider a feed-forward neural network with any number of layers, with only one restriction, that is, the first layer has at most $k<d$ outputs $z_{1}, \ldots, z_{k}$, where

$$
z_{j}=\sigma\left(b_{j}+\sum_{i=1}^{d} a_{j i} x^{(i)}\right)
$$

$x=\left(x^{(1)}, \ldots, x^{(d)}\right)$ is the input, and $\sigma$ is an arbitrary function $\mathcal{R} \rightarrow \mathcal{R}$ (not just a sigmoid). The integer $k$ remains fixed. Let $A$ denote the $k \times(d+1)$ matrix of weights $a_{j i}, b_{j}$.
(1) If $L^{*}(A)$ is the Bayes error for a recognition problem based upon ( $Z, Y$ ), with $Z=\left(Z_{1}, \ldots, Z_{k}\right)$ and

$$
Z_{j}=\sigma\left(b_{j}+\sum_{i=1}^{d} a_{j i} X^{(i)}\right), \quad X=\left(X^{(1)}, \ldots, X^{(d)}\right)
$$

then show that for some distribution of $(X, Y), \inf _{A} L^{*}(A)>L^{*}$, where $L^{*}$ is the Bayes probability of error for ( $X, Y$ ).
(2) If $k \geq d$ however, show that for any strictly monotonically increasing sigmoid $\sigma$, $\inf _{A} L^{*}(A)=L^{*}$.
(3) Use (1) to conclude that any neural network based upon a first layer with $k<d$ outputs is not consistent for some distribution, regardless of how many layers it has (note however, that the inputs of each layer are restricted to be the outputs of the previous layer).


Figure 30.16. The first layer is restricted to have $k$ outputs. It has $k(d+1)$ tunable parameters.

Problem 30.16. Find conditions on $k_{n}$ and $\beta_{n}$ that guarantee strong universal consistency in Theorem 30.9.

Problem 30.17. Barron networks. Call a Barron network a network of any number of layers (as in Figure 30.13) with 2 inputs per cell and cells that perform the operation $\alpha+\beta x+\gamma y+\delta x y$ on inputs $x, y \in \mathcal{R}$, with trainable weights $\alpha, \beta, \gamma, \delta$. If $\psi$ is the output
of a network with $d$ inputs and $k$ cells (arranged in any way), compute an upper bound for the vc dimension of the classifier $g(x)=I_{\{\psi(x)>0\}}$, as a function of $d$ and $k$. Note that the structure of the network (i.e., the positions of the cells and the connections) is variable.

Problem 30.18. Continued. Restrict the Barron network to $l$ layers and $k$ cells per layer (for $k l$ cells total), and repeat the previous exercise.

Problem 30.19. Continued. Find conditions on $k$ and $l$ in the previous exercises that would guarantee the universal consistency of the Barron network, if we train to minimize the empirical error.

Problem 30.20. Consider the family of functions $\mathcal{F}_{k, l}$ of the form

$$
\sum_{i=1}^{k} \sum_{j=1}^{l} w_{i j}\left(\psi_{i}(x)\right)^{j}, \quad \psi_{i}(x)=\sum_{i^{\prime}=1}^{d} \sum_{j^{\prime}=1}^{l} w_{i i^{\prime} j^{\prime}}^{\prime}\left(x^{\left(i^{\prime}\right)}\right)^{j^{\prime}}, \quad 1 \leq i \leq k
$$

where all the $w_{i j}$ 's and $w_{i i^{\prime} j^{\prime}}^{\prime}$ 's are tunable parameters. Show that for every $f \in C[0,1]^{d}$ and $\epsilon>0$, there exist $k, l$ large enough so that for some $g \in \mathcal{F}_{k, l}, \sup _{x \in[0,1]^{d}}|f(x)-g(x)| \leq \epsilon$.

Problem 30.21. Continued. Obtain an upper bound on the vc dimension of the above two-hidden-layer network. Hint: The vc dimension is usually about equal to the number of degrees of freedom (which is $(1+d) l k$ here).

Problem 30.22. Continued. If $g_{n}$ is the rule obtained by empirical error minimization over $\mathcal{F}_{k, l}$, then show that $L\left(g_{n}\right) \rightarrow L^{*}$ in probability if $k \rightarrow \infty, l \rightarrow \infty$, and $k l=o(n / \log n)$.

Problem 30.23. How many different monomials of order $r$ in $x^{(1)}, \ldots, x^{(d)}$ are there? How does this grow with $r$ when $d$ is held fixed?

Problem 30.24. Show that $\psi_{1}, \psi_{2}$, and $\phi$ in the bit-interleaving example in the section on Kolmogorov-Lorentz representations are not continuous. Which are the points of discontinuity?

Problem 30.25. Let

$$
\mathcal{F}_{m}^{\prime}=\left\{\sum_{j=1}^{m} c_{j} e^{a_{j}^{T} x} ; c_{j} \in \mathcal{R}, a_{j} \in \mathcal{R}^{d}\right\} .
$$

Pick $\left\{a_{j}, c_{j}\right\}$ by empirical risk minimization for the classifier $g(x)=I_{\{\psi(x)>0\}}, \psi \in \mathcal{F}_{m}^{\prime}$. Show that $\mathbf{E}\left\{L_{n}\right\} \rightarrow L^{*}$ for all distributions of $(X, Y)$ when $m \rightarrow \infty$ and $m=o\left(n^{1 / d} / \log n\right)$.
Problem 30.26. Write $\left(x^{(1)} x^{(2)}\right)^{2}$ as $\sum_{i=1}^{4}\left(a_{i 1} x^{(1)}+a_{i 2} x^{(2)}\right)^{2}$ and identify the coefficients $\left\{a_{i 1}, a_{i 2}\right\}$. Show that there is an entire subspace of solutions (Diaconis and Shahshahani (1984)).

Problem 30.27. Show that $e^{x^{(1)} x^{(2)}}$ and $\sin \left(x^{(1)} x^{(2)}\right)$ cannot be written in the form $\sum_{i=1}^{k}$ $\psi_{i}\left(a_{i}^{T} x\right)$ for any finite $k$, where $x=\left(x^{(1)}, x^{(2)}\right)$ and $a_{i} \in \mathcal{R}^{2}, 1 \leq i \leq k$ (Diaconis and Shahshahani (1984)). Thus, the projection pursuit representation of functions can only at best approximate all continuous functions on bounded sets.

Problem 30.28. Let $\mathcal{F}$ be the class of classifiers of the form $g=I_{\{\psi>0\}}$, where $\psi(x)=$ $a_{1} f_{1}\left(x^{(1)}\right)+a_{2} f_{2}\left(x^{(2)}\right)$, for arbitrary functions $f_{1}, f_{2}$, and coefficients $a_{1}, a_{2} \in \mathcal{R}$. Show that for some distribution of $(X, Y)$ on $\mathcal{R}^{2} \times\{0,1\}$, $\inf _{g \in \mathcal{F}} L(g)>L^{*}$, so that there is no hope of meaningful distribution-free classification based on additive functions only.

Problem 30.29. Continued. Repeat the previous exercise for functions of the form $\psi(x)=$ $a_{1} f_{1}\left(x^{(2)}, x^{(3)}\right)+a_{2} f_{2}\left(x^{(1)}, x^{(3)}\right)+a_{3} f_{3}\left(x^{(1)}, x^{(2)}\right)$ and distributions of $(X, Y)$ on $\mathcal{R}^{3} \times\{0,1\}$. Thus, additive functions of pairs do not suffice either.

Problem 30.30. Let $K: \mathcal{R} \rightarrow \mathcal{R}$ be a nonnegative bounded kernel with $\int K(x) d x<$ $\infty$. Show that for any $\epsilon>0$, any measurable function $f: \mathcal{R}^{d} \rightarrow \mathcal{R}$ and probability measure $\mu$ such that $\int|f(x)| \mu(d x)<\infty$, there exists a function $\psi$ of the form $\psi(x)=$ $\sum_{i=1}^{k} c_{i} K\left(\left(x-b_{i}\right)^{T} A_{i}\left(x-b_{i}\right)\right)+c_{0}$ such that $\int \mid f(x)-\psi(x) \mu(d x)<\epsilon$ (see Poggio and Girosi (1990), Park and Sandberg (1991; 1993), Darken, Donahue, Gurvits, and Sontag (1993), Krzyżak, Linder, and Lugosi (1993) for such denseness results). Hint: Relate the problem to a similar result for kernel estimates.

Problem 30.31. Let $\mathcal{C}^{(k)}$ be the class of classifiers defined by the functions

$$
\sum_{i=1}^{k} c_{i} K\left(\left(x-b_{i}\right)^{T} A_{i}\left(x-b_{i}\right)\right)+c_{0}
$$

Find upper bounds on its vC dimension when $K$ is an indicator of an interval containing the origin.

Problem 30.32. Consider the class of radial basis function networks

$$
\mathcal{F}_{n}=\left\{\sum_{i=1}^{k_{n}} c_{i} K\left(\left(x-b_{i}\right)^{T} A_{i}\left(x-b_{i}\right)\right)+c_{0}: \sum_{i=0}^{k_{n}}\left|c_{i}\right| \leq \beta_{n}\right\},
$$

where $K$ is nonnegative, unimodal, bounded, and continuous. Let $\psi_{n}$ be a function that minimizes $J_{n}(\psi)=n^{-1} \sum_{i=1}^{n}\left|\psi\left(X_{i}\right)-Y_{i}\right|$ over $\psi \in \mathcal{F}_{n}$, and define $g_{n}$ as the corresponding classifier. Prove that if $k_{n} \rightarrow \infty, \beta_{n} \rightarrow \infty$, and $k_{n} \beta_{n}^{2} \log \left(k_{n} \beta_{n}\right)=o(n)$ as $n \rightarrow \infty$, then $g_{n}$ is universally consistent (Krzyżak, Linder, and Lugosi (1993)). Hint: Proceed as in the proof of Theorem 30.9. Use Problem 30.30 to handle the approximation error. Bounding the covering numbers needs a little additional work.

## 31

## Other Error Estimates

In this chapter we discuss some alternative error estimates that have been introduced to improve on the performance of the standard estimates-holdout, resubstitution, and deleted-we have encountered so far. The first group of these estimates-smoothed and posterior probability estimates-are used for their small variance. However, we will give examples that show that classifier selection based on the minimization of these estimates may fail even in the simplest situations. Among other alternatives, we deal briefly with the rich class of bootstrap estimates.

### 31.1 Smoothing the Error Count

The resubstitution, deleted, and holdout estimates of the error probability (see Chapters 22 and 23) are all based on counting the number of errors committed by the classifier to be tested. This is the reason for the relatively large variance inherently present in these estimates. This intuition is based on the following. Most classification rules can be written into the form

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}\left(x, D_{n}\right) \geq 1 / 2 \\ 1 & \text { otherwise },\end{cases}
$$

where $\eta_{n}\left(x, D_{n}\right)$ is either an estimate of the a posteriori probability $\eta(x)$, as in the case of histogram, kernel, or nearest neighbor rules; or something else, as for generalized linear, or neural network classifiers. In any case, if $\eta_{n}\left(x, D_{n}\right)$ is close to $1 / 2$, then we feel that the decision is less robust compared to when the value of $\eta_{n}\left(x, D_{n}\right)$ is far from $1 / 2$. In other words, intuitively, inverting the value of the
decision $g_{n}(x)$ at a point $x$, where $\eta_{n}\left(x, D_{n}\right)$ is close to $1 / 2$ makes less difference in the error probability, than if $\left|\eta_{n}\left(x, D_{n}\right)-1 / 2\right|$ is large. The error estimators based on counting the number of errors do not take the value of $\eta_{n}(x)$ into account: they "penalize" errors with the same amount, no matter what the value of $\eta_{n}$ is. For example, in the case of the resubstitution estimate

$$
L_{n}^{(R)}=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{g_{n}\left(X_{i}\right) \neq Y_{i}\right\}}
$$

each error contributes with $1 / n$ to the overall count. Now, if $\eta_{n}\left(X_{i}, D_{n}\right)$ is close to $1 / 2$, a small perturbation of $X_{i}$ can flip the decision $g_{n}\left(X_{i}\right)$, and therefore change the value of the estimate by $1 / n$, although the error probability of the rule $g_{n}$ probably does not change by this much. This phenomenon is what causes the relatively large variance of error counting estimators.

Glick (1978) proposed a modification of the counting error estimates. The general form of his "smoothed" estimate is

$$
L_{n}^{(S)}=\frac{1}{n} \sum_{i=1}^{n}\left(I_{\left\{Y_{i}=0\right\}} r\left(\eta_{n}\left(X_{i}, D_{n}\right)\right)+I_{\left\{Y_{i}=1\right\}}\left(1-r\left(\eta_{n}\left(X_{i}, D_{n}\right)\right)\right)\right)
$$

where $r$ is a monotone increasing function satisfying $r(1 / 2-u)=1-r(1 / 2+u)$. Possible choices of the smoothing function $r(u)$ are $r(u)=u$, or $r(u)=1 /(1+$ $\left.e^{1 / 2-c u}\right)$, where the parameter $c>0$ may be adjusted to improve the behavior of the estimate (see also Glick (1978), Knoke (1986), or Tutz (1985)). Both of these estimates give less penalty to errors close to the decision boundary, that is, to errors where $\eta_{n}$ is close to $1 / 2$. Note that taking $r(u)=I_{\{u \geq 1 / 2\}}$ corresponds to the resubstitution estimate. We will see in Theorem 31.2 below that if $r$ is smooth, then $L_{n}^{(S)}$ indeed has a very small variance in many situations.

Just like the resubstitution estimate, the estimate $L_{n}^{(S)}$ may be strongly optimistically biased. Just consider the 1-nearest neighbor rule, when $L_{n}^{(S)}=0$ with probability one, whenever $X$ has a density. To combat this defect, one may define the deleted version of the smoothed estimate,

$$
L_{n}^{(S D)}=\frac{1}{n} \sum_{i=1}^{n}\left(I_{\left\{Y_{i}=0\right\}} r\left(\eta_{n-1}\left(X_{i}, D_{n, i}\right)\right)+I_{\left\{Y_{i}=1\right\}}\left(1-r\left(\eta_{n-1}\left(X_{i}, D_{n, i}\right)\right)\right)\right)
$$

where $D_{n, i}$ is the training sequence with the $i$-th pair ( $X_{i}, Y_{i}$ ) deleted. The first thing we notice is that this estimate is still biased, even asymptotically. To illustrate this point, consider $r(u)=u$. In this case,

$$
\begin{aligned}
\mathbf{E} & \left\{L_{n}^{(S D)}\right\} \\
& =\mathbf{E}\left\{I_{\{Y=0\}} \eta_{n-1}\left(X, D_{n-1}\right)+I_{\{Y=1\}}\left(1-\eta_{n-1}\left(X, D_{n-1}\right)\right)\right\} \\
& =\mathbf{E}\left\{(1-\eta(X)) \eta_{n-1}\left(X, D_{n-1}\right)+\eta(X)\left(1-\eta_{n-1}\left(X, D_{n-1}\right)\right)\right\}
\end{aligned}
$$

If the estimate $\eta_{n-1}\left(x, D_{n-1}\right)$ was perfect, that is, equal to $\eta(x)$ for every $x$, then the expected value above would be $2 \mathbf{E}\{\eta(X)(1-\eta(X))\}$, which is the asymptotic
error probability $L_{\mathrm{NN}}$ of the $1-\mathrm{NN}$ rule. In fact,

$$
\begin{aligned}
&\left|\mathbf{E}\left\{L_{n}^{(S D)}\right\}-L_{\mathrm{NN}}\right| \\
&= \mid \mathbf{E}\left\{I_{\{Y=0\}} \eta_{n-1}\left(X, D_{n-1}\right)+I_{\{Y=1\}}\left(1-\eta_{n-1}\left(X, D_{n-1}\right)\right)\right. \\
&\left.-\left(I_{\{Y=0\}} \eta(X)+I_{\{Y=1\}}(1-\eta(X))\right)\right\} \mid \\
& \leq 2\left|\mathbf{E}\left\{\eta_{n-1}\left(X, D_{n-1}\right)-\eta(X)\right\}\right| .
\end{aligned}
$$

This means that when the estimate $\eta_{n}(x)$ of the a posteriori probability of $\eta(x)$ is consistent in $L_{1}(\mu)$, then $L_{n}^{(S D)}$ converges to $L_{\mathrm{NN}}$, and not to $L^{*}$ !

Biasedness of an error estimate is not necessarily a bad property. In most applications, all we care about is how the classification rule selected by minimizing the error estimate works. Unfortunately, in this respect, smoothed estimates perform poorly, even compared to other strongly biased error estimates such as the empirical squared error (see Problem 31.4). The next example illustrates our point.

Theorem 31.1. Let the distribution of $X$ be concentrated on two values such that $\mathbf{P}\{X=a\}=\mathbf{P}\{X=b\}=1 / 2$, and let $\eta(a)=3 / 8$ and $\eta(b)=5 / 8$. Assume that the smoothed error estimate

$$
L_{n}^{(S)}\left(\eta^{\prime}\right)=\frac{1}{n} \sum_{i=1}^{n}\left(I_{\left\{Y_{i}=0\right\}} \eta^{\prime}\left(X_{i}\right)+I_{\left\{Y_{i}=1\right\}}\left(1-\eta^{\prime}\left(X_{i}\right)\right)\right)
$$

is minimized over $\eta^{\prime} \in \mathcal{F}$, to select a classifier from $\mathcal{F}$, where the class $\mathcal{F}$ contains two functions, the true a posteriori probability function $\eta$, and $\tilde{\eta}$, where

$$
\tilde{\eta}(a)=0, \quad \tilde{\eta}(b)=\frac{3}{8} .
$$

Then the probability that $\tilde{\eta}$ is selected converges to one as $n \rightarrow \infty$.

Proof. Straightforward calculation shows that

$$
\mathbf{E}\left\{L_{n}^{(S)}(\eta)\right\}=\frac{15}{32} \text { and } \mathbf{E}\left\{L_{n}^{(S)}(\tilde{\eta})\right\}=\frac{6}{32}
$$

The statement follows from the law of large numbers.

REMARK. The theorem shows that even if the true a posteriori probability function is contained in a finite class of candidates, the smoothed estimate with $r(u)=u$ is unable to select a good discrimination rule. The result may be extended to general smooth $r$ 's. As Theorems 15.1 and 29.2 show, empirical squared error minimization or maximum likelihood never fail in this situation.

Finally we demonstrate that if $r$ is smooth, then the variance of $L_{n}^{(S)}$ is indeed small. Our analysis is based on the work of Lugosi and Pawlak (1994). The bounds
for the variance of $L_{n}^{(S)}$ remain valid for $L_{n}^{(S D)}$. Consider classification rules of the form

$$
g_{n}(x)= \begin{cases}0 & \text { if } \eta_{n}\left(x, D_{n}\right) \geq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

where $\eta_{n}\left(x, D_{n}\right)$ is an estimate of the a posteriori probability $\eta(x)$. Examples include the histogram rule, where

$$
\eta_{n}\left(x, D_{n}\right)=\frac{\sum_{i=1}^{n} Y_{i} I_{\left\{X_{i} \in A_{n}(x)\right\}}}{\sum_{i=1}^{n} I_{\left\{X_{i} \in A_{n}(x)\right\}}}
$$

(see Chapters 6 and 9), the $k$-nearest neighbor rule, where

$$
\eta_{n}\left(x, D_{n}\right)=\frac{1}{k} \sum_{i=1}^{k} Y_{(i)}(x)
$$

(Chapters 5 and 11), or the kernel rule, where

$$
\eta_{n}\left(x, D_{n}\right)=\frac{\sum_{i=1}^{n} Y_{i} K_{h}\left(x-X_{i}\right)}{n \mathbb{E}\left\{K_{h}(x-X)\right\}}
$$

(Chapter 10). In the sequel, we concentrate on the performance of the smoothed estimate of the error probability of these nonparametric rules. The next theorem shows that for these rules, the variance of the smoothed error estimate is $O(1 / n)$, no matter what the distribution is. This is a significant improvement over the variance of the deleted estimate, which, as pointed out in Chapter 23, can be larger than $1 / \sqrt{2 \pi n}$.

Theorem 31.2. Assume that the smoothing function $r(u)$ satisfies $0 \leq r(u) \leq 1$ for $u \in[0,1]$, and is uniformly Lipschitz continuous, that is,

$$
|r(u)-r(v)| \leq c|u-v|
$$

for all $u, v$, and for some constant $c$. Then the smoothed estimate $L_{n}^{(S)}$ of the histogram, $k$-nearest neighbor, and moving window rules (with kernel $K=I_{S_{0,1}}$ ) satisfies

$$
\mathbf{P}\left\{\left|L_{n}^{(S)}-\mathbf{E}\left\{L_{n}^{(S)}\right\}\right| \geq \epsilon\right\} \leq 2 e^{-2 n \epsilon^{2} / C},
$$

and

$$
\operatorname{Var}\left\{L_{n}^{(S)}\right\} \leq \frac{C}{4 n}
$$

where $C$ is a constant depending on the rule only. In the case of the histogram rule the value of $C$ is $C=(1+4 c)^{2}$, for the $k$-nearest neighbor rule $C=\left(1+2 c \gamma_{d}\right)^{2}$, and for the moving window rule $C=\left(1+2 c \beta_{d}\right)^{2}$. Here $c$ is the constant in the Lipschitz condition, $\gamma_{d}$ is the minimal number of cones centered at the origin of angle $\pi / 6$ that cover $\mathcal{R}^{d}$, and $\beta_{d}$ is the minimal number of balls of radius $1 / 2$ that cover the unit ball in $\mathcal{R}^{d}$.

Remark. Notice that the inequalities of Theorem 31.2 are valid for all $n, \epsilon$, and $h>0$ for the histogram and moving window rules, and $k$ for the nearest neighbor rules. Interestingly, the constant $C$ does not change with the dimension in the histogram case, but grows exponentially with $d$ for the $k$-nearest neighbor and moving window rules.

Proof. The probability inequalities follow from appropriate applications of McDiarmid's inequality. The upper bound on the variance follows similarly from Theorem 9.3. We consider each of the three rules in turn.

The histogram rule. Let $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)$ be a fixed training sequence. If we can show that by replacing the value of a pair $\left(x_{i}, y_{i}\right)$ in the training sequence by some $\left(x_{i}^{\prime}, y_{i}^{\prime}\right)$ the value of the estimate $L_{n}^{(S)}$ can change by at most $(1+2 c) / n$, then the inequality follows by applying Theorem 9.2 with $c_{i}=5 / n$.

The $i$-th term of the sum in $L_{n}^{(S)}$ can change by one, causing $1 / n$ change in the average. Obviously, all the other terms in the sum that can change are the ones corresponding to the $x_{j}$ 's that are in the same set of the partition as either $x_{i}$ or $x_{i}^{\prime}$. Denoting the number of points in the same set with $x_{i}$ and $x_{i}^{\prime}$ by $k$ and $k^{\prime}$ respectively, it is easy to see that the estimate of the a posteriori probabilities in these points can change by at most $2 / k$ and $2 / k^{\prime}$, respectively. It means that the overall change in the value of the sum can not exceed $\left(1+k \frac{2 c}{k}+k^{\prime} \frac{2 c}{k^{\prime}}\right) / n=(1+4 c) / n$.
The k-nearest neighbor rule. To avoid difficulties caused by breaking distance ties, assume that $X$ has a density. Then recall that the application of Lemma 11.1 for the empirical distribution implies that no $X_{j}$ can be one of the $k$ nearest neighbors of more than $k \gamma_{d}$ points from $D_{n}$. Thus, changing the value of one pair in the training sequence can change at most $1+2 k \gamma_{d}$ terms in the expression of $L_{n}^{(D)}$, one of them by at most 1 , and all the others by at most $c / k$. Theorem 9.2 yields the result.

The moving window rule. Again, we only have to check the condition of Theorem 9.2 with $c_{i}=\left(1+2 c \beta_{d}\right) / n$. Fix a training sequence $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)$ and replace the pair $\left(x_{i}, y_{i}\right)$ by $\left(x_{i}^{\prime}, y_{i}^{\prime}\right)$. Then the $i$-th term in the sum of the expression of $L_{n}^{(S)}$ can change by at most one. Clearly, the $j$-th term, for which $x_{j} \notin S_{x_{i}, h}$ and $x_{j} \notin S_{x_{i}^{\prime}, h}$, keeps its value. It is easy to see that all the other terms can change by at most $c \cdot \max \left\{1 / k_{j}, 1 / k_{j}^{\prime}\right\}$, where $k_{j}$ and $k_{j}^{\prime}$ are the numbers of points $x_{k}$, $k \neq i, j$, from the training sequence that fall in $S_{x_{i}, h}$ and $S_{x_{i}^{\prime}, h}$, respectively. Thus, the overall change in the sum does not exceed

$$
1+\sum_{x_{j} \in S_{x_{i}, h}} \frac{c}{k_{j}}+\sum_{x_{j} \in S_{x_{i}^{\prime}, n}} \frac{c}{k_{j}^{\prime}} .
$$

It suffices to show by symmetry that $\sum_{x_{j} \in S_{x_{i}, h}} 1 / k_{j} \leq \beta_{d}$. Let $n_{j}=\mid\left\{x_{k}, k \neq\right.$ $\left.i, j: x_{k} \in S_{x_{i}, h} \cap S_{x_{j}, h}\right\} \mid$. Then clearly,

$$
\sum_{x_{j} \in S_{x_{i, h}}} \frac{1}{k_{j}} \leq \sum_{x_{j} \in S_{x_{i}, h}} \frac{1}{n_{j}} .
$$

To bound the right-hand side from above, cover $S_{x_{i}, h}$ by $\beta_{d}$ balls $S_{1}, \ldots, S_{\beta_{d}}$ of radius $h / 2$. Denote the number of points falling in them by $l_{m}\left(m=1, \ldots, \beta_{d}\right)$ :

$$
l_{m}=\left|\left\{x_{k}, \quad k \neq i: x_{k} \in S_{x_{i}, h} \cap S_{m}\right\}\right| .
$$

Then

$$
\sum_{x_{j} \in S_{x_{i}, h}} \frac{1}{n_{j}} \leq \sum_{x_{j} \in S_{x_{i}, h}} \sum_{m: x_{j} \in S_{m}} \frac{1}{l_{m}}=\beta_{d},
$$

and the theorem is proved.

### 31.2 Posterior Probability Estimates

The error estimates discussed in this section improve on the biasedness of smoothed estimates, while preserving their small variance. Still, these estimates are of questionable utility in classifier selection. Considering the formula

$$
L\left(g_{n}\right)=\mathbf{E}\left\{I_{\left\{\eta_{n}\left(X, D_{n}\right) \leq 1 / 2\right\}} \eta(X)+I_{\left\{\eta_{n}\left(X, D_{n}\right)>1 / 2\right\}}(1-\eta(X)) \mid D_{n}\right\}
$$

for the error probability of a classification rule $g_{n}(x)=I_{\left\{\eta_{n}\left(x, D_{n}\right)>1 / 2\right\}}$, it is plausible to introduce the estimate

$$
\begin{aligned}
L_{n}^{(P)}= & \frac{1}{n} \sum_{i=1}^{n}\left(I_{\left\{\eta_{n}\left(X_{i}, D_{n}\right) \leq 1 / 2\right\}} \eta_{n}\left(X_{i}, D_{n}\right)\right. \\
& \left.+I_{\left\{\eta_{n}\left(X_{i}, D_{n}\right)>1 / 2\right\}}\left(1-\eta_{n}\left(X_{i}, D_{n}\right)\right)\right),
\end{aligned}
$$

that is, the expected value is estimated by a sample average, and instead of the (unknown) a posteriori probability $\eta(x)$, its estimate $\eta_{n}\left(x, D_{n}\right)$ is plugged into the formula of $L_{n}$. The estimate $L_{n}^{(P)}$ is usually called the posterior probability error estimate. In the case of nonparametric rules such as histogram, kernel, and $k$-NN rules it is natural to use the corresponding nonparametric estimates of the a posteriori probabilities for plugging in the expression of the error probability. This, and similar estimates of $L_{n}$ have been introduced and studied by Fukunaga and Kessel (1973), Hora and Wilcox (1982), Fitzmaurice and Hand (1987), Ganesalingam and McLachlan (1980), Kittler and Devijver (1981), Matloff and Pruitt (1984), Moore, Whitsitt, and Landgrebe (1976), Pawlak (1988), Schwemer and Dunn (1980), and Lugosi and Pawlak (1994). It is interesting to notice the similarity between the estimates $L_{n}^{(S)}$ and $L_{n}^{(P)}$, although they were developed from different scenarios.

To reduce the bias, we can use the leave-one-out, (or deleted) version of the estimate,

$$
\begin{aligned}
L_{n}^{(P D)}= & \frac{1}{n} \sum_{i=1}^{n}\left(I_{\left\{\eta_{n-1}\left(X_{i}, D_{n, i}\right) \leq 1 / 2\right\}} \eta_{n-1}\left(X_{i}, D_{n, i}\right)\right) \\
& \left.+I_{\left\{\eta_{n-1}\left(X_{i}, D_{n, i}\right)>1 / 2\right\}}\left(1-\eta_{n-1}\left(X_{i}, D_{n, i}\right)\right)\right) .
\end{aligned}
$$

The deleted version $L_{n}^{(P D)}$, has a much better bias than $L_{n}^{(S D)}$. We have the bound

$$
\begin{aligned}
&\left|\mathbf{E}\left\{L_{n}^{(P D)}\right\}-\mathbf{E} L_{n-1}\right| \\
&= \mid \mathbf{E}\left\{I_{\left\{g_{n}(X)=0\right\}} \eta_{n-1}\left(X, D_{n-1}\right)+I_{\left\{g_{n}(X)=1\right\}}\left(1-\eta_{n-1}\left(X, D_{n-1}\right)\right)\right. \\
&\left.-\left(I_{\left\{g_{n}(X)=0\right\}} \eta(X)+I_{\left\{g_{n}(X)=1\right\}}(1-\eta(X))\right)\right\} \mid \\
& \leq 2\left|\mathbf{E}\left\{\eta_{n-1}\left(X, D_{n-1}\right)-\eta(X)\right\}\right|
\end{aligned}
$$

This means that if the estimate $\eta_{n}(x)$ of the a posteriori probability of $\eta(x)$ is consistent in $L_{1}(\mu)$, then $\mathbf{E} L_{n}^{(P D)}$ converges to $L^{*}$. This is the case for all distributions for the histogram and moving window rules if $h \rightarrow 0$ and $n h^{d} \rightarrow \infty$, and the $k$-nearest neighbor rule if $k \rightarrow \infty$ and $k / n \rightarrow 0$, as it is seen in Chapters 9,10 , and 11. For specific cases it is possible to obtain sharper bounds on the bias of $L_{n}^{(P D)}$. For the histogram rule, Lugosi and Pawlak (1994) carried out such analysis. They showed for example that the estimate $L_{n}^{(P D)}$ is optimistically biased (see Problem 31.2).

Posterior probability estimates of $L_{n}$ share the good stability properties of smoothed estimates (Problem 31.1).

Finally, let us select a function $\eta^{\prime}: \mathcal{R}^{d} \rightarrow[0,1]$-and a corresponding rule $g(x)=I_{\left\{\eta^{\prime}(x)>1 / 2\right\}}$-from a class $\mathcal{F}$ based on the minimization of the posterior probability error estimate

$$
L_{n}^{(P)}\left(\eta^{\prime}\right)=\frac{1}{n} \sum_{i=1}^{n}\left(I_{\left\{\eta^{\prime}\left(X_{i}\right) \leq 1 / 2\right\}} \eta^{\prime}\left(X_{i}\right)+I_{\left\{\eta^{\prime}\left(X_{i}\right)>1 / 2\right\}}\left(1-\eta^{\prime}\left(X_{i}\right)\right)\right)
$$

Observe that $L_{n}^{(P)}\left(\eta^{\prime}\right)=0$ when $\eta^{\prime}(x) \in\{0,1\}$ for all $x$, that is, rule selection based on this estimate just does not make sense. The reason is that $L_{n}^{(P)}$ ignores the $Y_{i}$ 's of the data sequence!

Fukunaga and Kessel (1973) argued that efficient posterior probability estimators can be obtained if additional unclassified observations are available. Very often in practice, in addition to the training sequence $D_{n}$, further feature vectors $X_{n+1}, \ldots, X_{n+l}$ are given without their labels $Y_{n+1}, \ldots, Y_{n+l}$, where the $X_{n+i}$ are i.i.d., independent from $X$ and $D_{n}$, and they have the same distribution as that of $X$. This situation is typical in medical applications, when large sets of medical records are available, but it is usually very expensive to get their correct diagnosis. These unclassified samples can be efficiently used for testing the performance of a classifier designed from $D_{n}$ by using the estimate

$$
\begin{aligned}
L_{n, l}^{(U)}= & \frac{1}{l} \sum_{i=1}^{l}\left(I_{\left\{\eta_{n}\left(X_{n+i}, D_{n}\right) \leq 1 / 2\right\}} \eta_{n}\left(X_{n+i}, D_{n}\right)\right. \\
& \left.+I_{\left\{\eta_{n}\left(X_{n+i}, D_{n}\right)>1 / 2\right\}}\left(1-\eta_{n}\left(X_{n+i}, D_{n}\right)\right)\right) .
\end{aligned}
$$

Again, using $L_{n, l}^{(U)}$ for rule selection is meaningless.

### 31.3 Rotation Estimate

This method, suggested by Toussaint and Donaldson (1970), is a combination of the holdout and deleted estimates. It is sometimes called the $\Pi$-estimate. Let $s<n$ be a positive integer (typically much smaller than $n$ ), and assume, for the sake of simplicity, that $q=n / s$ is an integer. The rotation method forms the holdout estimate, by holding the first $s$ pairs of the training data out, then the second $s$ pairs, and so forth. The estimate is defined by averaging the $q$ numbers obtained this way. To formalize, denote by $D_{n, j}^{(s)}$ the training data, with the $j$-th $s$-block held out $(j=1, \ldots, q)$ :

$$
D_{n, j}^{(s)}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{s(j-1)}, Y_{s(j-1)}\right),\left(X_{s j+1}, Y_{s j+1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right) .
$$

The estimate is defined by

$$
L_{n, s}^{(\Pi)}=\frac{1}{q} \sum_{j=1}^{q} \frac{1}{s} \sum_{i=s(j-1)+1}^{s j} I_{\left\{g_{n-s}\left(X_{i}, D_{n, j}^{(s)}\right) \neq Y_{i}\right\}} .
$$

$s=1$ yields the deleted estimate. If $s>1$, then the estimate is usually more biased than the deleted estimate, as

$$
\mathbf{E} L_{n, s}^{(\Pi)}=\mathbf{E} L_{n-s},
$$

but usually exhibits smaller variance.

### 31.4 Bootstrap

Bootstrap methods for estimating the misclassification error became popular following the revolutionary work of Efron (1979; 1983). All bootstrap estimates introduce artificial randomization. The bootstrap sample

$$
D_{m}^{(b)}=\left(\left(X_{1}^{(b)}, Y_{1}^{(b)}\right), \ldots,\left(X_{m}^{(b)}, Y_{m}^{(b)}\right)\right)
$$

is a sequence of random variable pairs drawn randomly with replacement from the set $\left\{\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right\}$. In other words, conditionally on the training sample $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$, the pairs $\left(X_{i}^{(b)}, Y_{i}^{(b)}\right)$ are drawn independently from $v_{n}$, the empirical distribution based on $D_{n}$ in $\mathcal{R}^{d} \times\{0,1\}$.

One of the standard bootstrap estimates aims to compensate the (usually optimistic) bias

$$
B\left(g_{n}\right)=\mathbb{E}\left\{L\left(g_{n}\right)\right\}-L_{n}^{(R)}\left(g_{n}\right)
$$

of the resubstitution estimate $L_{n}^{(R)}$. To estimate $B\left(g_{n}\right)$, a bootstrap sample of size $m=n$ may be used:

$$
\widehat{B}_{n}^{(b)}\left(g_{n}\right)=\sum_{i=1}^{n}\left(\frac{1}{n}-\frac{1}{n} \sum_{j=1}^{n} I_{\left\{X_{j}^{(b)} \neq X_{i}\right\}}\right) I_{\left\{g_{n}\left(X_{i}, D_{n}^{(b)}\right) \neq Y_{i}\right\}} .
$$

Often, bootstrap sampling is repeated several times to average out effects of the additional randomization. In our case,

$$
\bar{B}_{n}\left(g_{n}\right)=\frac{1}{B} \sum_{b=1}^{B} \widehat{B}_{n}^{(b)}\left(g_{n}\right)
$$

yields the estimator

$$
L_{n}^{(B)}\left(g_{n}\right)=L_{n}^{(R)}\left(g_{n}\right)-\bar{B}_{n}\left(g_{n}\right)
$$

Another instance of a bootstrap sample of size $n$, the so-called E0 estimator, uses resubstitution on the training pairs not appearing in the bootstrap sample. The estimator is defined by

$$
\frac{\sum_{j=1}^{n} I_{\left\{X_{j} \notin\left\{X_{1}^{(b)}, \ldots, X_{n}^{(b)}\right\}\right.} I_{\left\{g_{n}\left(X_{i}, D_{n}^{(b)}\right) \neq Y_{i}\right\}}}{\left.\sum_{j=1}^{n} I_{\left\{X_{j} \notin\left\{X_{1}^{(b)}, \ldots, X_{n}^{(b)}\right\}\right.}\right\}} .
$$

Here, too, averages may be taken after generating the bootstrap sample several times.

Many other versions of bootstrap estimates have been reported, such as the "0.632 estimate," "double bootstrap," and "randomized bootstrap" (see Hand (1986), Jain, Dubes, and Chen (1987), and McLachlan (1992) for surveys and additional references). Clearly, none of these estimates provides a universal remedy, but for several specific classification rules, bootstrap estimates have been experimentally found to outperform the deleted and resubstitution estimates. However, one point has to be made clear: we always lose information with the additional randomization. We summarize this in the following simple general result:

Theorem 31.3. Let $X_{1}, \ldots, X_{n}$ be drawn from an unknown distribution $\mu$, and let $\alpha(\mu)$ be a functional to be estimated. Let $r(\cdot)$ be a convex risk function (such as $r(u)=u^{2}$ or $\left.r(u)=|u|\right)$. Let $X_{1}^{(b)}, \ldots X_{m}^{(b)}$ be a bootstrap sample drawn from $X_{1}, \ldots, X_{n}$. Then

$$
\inf _{\psi} \mathbf{E}\left\{r\left(\psi\left(X_{1}^{(b)}, \ldots X_{m}^{(b)}\right)-\alpha(\mu)\right)\right\} \geq \inf _{\psi} \mathbf{E}\left\{r\left(\psi\left(X_{1}, \ldots, X_{n}\right)-\alpha(\mu)\right)\right\}
$$

The theorem states that no matter how large $m$ is, the class of estimators that are functions of the original sample is always at least as good as the class of all estimators that are based upon bootstrap samples. In our case, $\alpha(\mu)$ plays the role of the expected error probability $\mathbf{E} L_{n}=\mathbf{P}\left\{g_{n}(X) \neq Y\right\}$. If we take $r(u)=$ $u^{2}$, then it follows from the theorem that there is no estimator $\widehat{L}_{m}$ based on the bootstrap sample $D_{m}^{(b)}$ whose squared error $\mathbf{E}\left\{\left(\mathbb{E} L_{n}-\widehat{L}_{m}\right)^{2}\right\}$ is smaller than that of some nonbootstrap estimate. In the proof of the theorem we construct such a non-bootstrap estimator. It is clear, however, that, in general, the latter estimator is too complex to have any practical value. The randomization of bootstrap methods may provide a useful tool to overcome the computational difficulties in finding good estimators.

Proof. Let $\psi$ be any mapping taking $m$ arguments. Then

$$
\begin{aligned}
\mathbf{E} & \left\{r\left(\psi\left(X_{1}^{(b)}, \ldots X_{m}^{(b)}\right)-\alpha(\mu)\right) \mid X_{1}, \ldots, X_{n}\right\} \\
& =\frac{1}{n^{m}} \sum_{\left(i_{1}, \ldots, i_{m}\right) \subset\{1, \ldots, n\}^{m}} r\left(\psi\left(X_{i_{1}}, \ldots, X_{i_{m}}\right)-\alpha(\mu)\right) \\
& \geq r\left(\frac{1}{n^{m}} \sum_{\left(i_{1}, \ldots, i_{m}\right) \subset\{1, \ldots, n\}^{m}}\left(\psi\left(X_{i_{1}}, \ldots, X_{i_{m}}\right)-\alpha(\mu)\right)\right)
\end{aligned}
$$

(by Jensen's inequality and the convexity of $r$ )

$$
\begin{aligned}
& =r\left(\frac{1}{n^{m}} \sum_{\left(i_{1}, \ldots, i_{m}\right) \subset\{1, \ldots, n\}^{m}} \psi\left(X_{i_{1}}, \ldots, X_{i_{m}}\right)-\alpha(\mu)\right) \\
& =r\left(\psi^{*}\left(X_{1}, \ldots, X_{n}\right)-\alpha(\mu)\right)
\end{aligned}
$$

where

$$
\psi^{*}\left(x_{1}, \ldots, x_{n}\right)=\frac{1}{n^{m}} \sum_{\left(i_{1}, \ldots, i_{m}\right) \subset\{1, \ldots, n\}^{m}} \psi\left(x_{i_{1}}, \ldots, x_{i_{m}}\right) .
$$

Now, after taking expectations with respect to $X_{1}, \ldots, X_{n}$, we see that for every $\psi$ we start out with, there is a $\psi^{*}$ that is at least as good.

If $m=O(n)$, then the bootstrap has an additional problem related to the coupon collector problem. Let $N$ be the number of different pairs in the bootstrap sample $\left(X_{1}^{(b)}, Y_{1}^{(b)}\right), \ldots,\left(X_{m}^{(b)}, Y_{m}^{(b)}\right)$. Then, if $m \sim c n$ for some constant $c$, then $N / n \rightarrow$ $1-e^{-c}$ with probability one. To see this, note that

$$
\mathbf{E}\{n-N\}=n\left(1-\frac{1}{n}\right)^{c n} \sim n e^{-c}
$$

so $\mathbf{E}\{N / n\} \rightarrow 1-e^{-c}$. Furthermore, if one of the $m$ drawings is varied, $N$ changes by at most one. Hence, by McDiarmid's inequality, for $\epsilon>0$,

$$
\mathbf{P}\{|N-\mathbf{E}\{N\}|>n \epsilon\} \leq 2 e^{-2 n \epsilon^{2}},
$$

from which we conclude that $N / n \rightarrow 1-e^{-c}$ with probability one. As $n\left(1-e^{-c}\right)$ of the original data pairs do not appear in the bootstrap sample, a considerable loss of information takes place that will be reflected in the performance. This phenomenon is well-known, and motivated several modifications of the simplest bootstrap estimate. For more information, see the surveys by Hand (1986), Jain, Dubes, and Chen (1987), and McLachlan (1992).

## Problems and Exercises

Problem 31.1. Show that the posterior probability estimate $L_{n}^{(P)}$ of the histogram, $k$-nearest neighbor, and moving window rules satisfies

$$
\mathbf{P}\left\{\left|L_{n}^{(P)}-\mathbf{E}\left\{L_{n}^{(P)}\right\}\right| \geq \epsilon\right\} \leq 2 e^{-2 n \epsilon^{2} / C}
$$

and

$$
\operatorname{Var}\left\{L_{n}^{(P)}\right\} \leq \frac{C}{4 n}
$$

where $C$ is a constant, depending on the rule. In the case of the histogram rule the value of $C$ is $C=25$, for the $k$-nearest neighbor rule $C=\left(1+2 \gamma_{d}\right)^{2}$, and for the moving window rule $C=\left(1+2 \beta_{d}\right)^{2}$. Also show that the deleted version $L_{n}^{(P D)}$ of the estimate satisfies the same inequalities (Lugosi and Pawlak (1994)). Hint: Proceed as in the proof of Theorem 31.2.

Problem 31.2. Show that the deleted posterior probability estimate of the error probability of a histogram rule is always optimistically biased, that is, for all $n$, and all distributions, $\mathbf{E}\left\{L_{n}^{(P D)}\right\} \leq \mathbf{E} L_{n-1}$.

Problem 31.3. Show that for any classification rule and any estimate $0 \leq \eta_{n}\left(x, D_{n}\right) \leq 1$ of the a posteriori probabilities, for all distributions of $(X, Y)$ for all $l, n$, and $\epsilon>0$

$$
\mathbf{P}\left\{\left|L_{n, l}^{(U)}-\mathbf{E}\left\{L_{n, l}^{(U)}\right\}\right| \geq \epsilon \mid D_{n}\right\} \leq 2 e^{-2 l \epsilon^{2}},
$$

and $\operatorname{Var}\left\{L_{n, l}^{(U)} \mid D_{n}\right\} \leq 1 / l$. Further, show that for all $l, \mathbf{E}\left\{L_{n, l}^{(U)}\right\}=\mathbf{E}\left\{L_{n+1}^{(P D)}\right\}$.
Problem 31.4. Empirical squared error. Consider the deleted empirical squared error

$$
e_{n}=\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\eta_{n-1}\left(X_{i}, D_{n, i}\right)\right)^{2} .
$$

Show that

$$
\mathbf{E}\left\{e_{n}\right\}=\frac{L_{\mathrm{NN}}}{2}+\mathbf{E}\left\{\left(\eta(X)-\eta_{n-1}\left(X, D_{n-1}\right)\right)^{2}\right\}
$$

where $L_{\mathrm{NN}}$ is the asymptotic error probability on the 1-nearest neighbor rule. Show that if $\eta_{n-1}$ is the histogram, kernel, or $k$-NN estimate, then $\operatorname{Var}\left\{e_{n}\right\} \leq c / n$ for some constant depending on the dimension only. We see that $e_{n}$ is an asymptotically optimistically biased estimate of $L\left(g_{n}\right)$ when $\eta_{n-1}$ is an $L_{2}(\mu)$-consistent estimate of $\eta$. Still, this estimate is useful in classifier selection (see Theorem 29.2).

## 32

## Feature Extraction

### 32.1 Dimensionality Reduction

So far, we have not addressed the question of how the components of the feature vector $X$ are obtained. In general, these components are based on $d$ measurements of the object to be classified. How many measurements should be made? What should these measurements be? We study these questions in this chapter. General recipes are hard to give as the answers depend on the specific problem. However, there are some rules of thumb that should be followed. One such rule is that noisy measurements, that is, components that are independent of $Y$, should be avoided. Also, adding a component that is a function of other components is useless. A necessary and sufficient condition for measurements providing additional information is given in Problem 32.1.

Our goal, of course, is to make the error probability $L\left(g_{n}\right)$ as small as possible. This depends on many things, such as the joint distribution of the selected components and the label $Y$, the sample size, and the classification rule $g_{n}$. To make things a bit simpler, we first investigate the Bayes errors corresponding to the selected components. This approach makes sense, since the Bayes error is the theoretical limit of the performance of any classifier. As Problem 2.1 indicates, collecting more measurements cannot increase the Bayes error. On the other hand, having too many components is not desirable. Just recall the curse of dimensionality that we often faced: to get good error rates, the number of training samples should be exponentially large in the number of components. Also, computational and storage limitations may prohibit us from working with many components.

We may formulate the feature selection problem as follows: Let $X^{(1)}, \ldots, X^{(d)}$
be random variables representing $d$ measurements. For a set $A \subseteq\{1, \ldots, d\}$ of indices, let $X_{A}$ denote the $|A|$-dimensional random vector, whose components are the $X^{(i)}$ 's with $i \in A$ (in the order of increasing indices). Define

$$
L^{*}(A)=\inf _{g: \mathcal{R}^{|A|} \rightarrow\{0,1\}} \mathbf{P}\left\{g\left(X_{A}\right) \neq Y\right\}
$$

as the Bayes error corresponding to the pair $\left(X_{A}, Y\right)$.


Figure 32.1. A possible arrangement of $L^{*}(A)$ 's for $d=3$.

Obviously, $L^{*}(A) \leq L^{*}(B)$ whenever $B \subset A$, and $L^{*}(\emptyset)=\min (\mathbf{P}\{Y=$ $0\}, \mathbf{P}\{Y=1\})$. The problem is to find an efficient way of selecting an index set $A$ with $|A|=k$, whose corresponding Bayes error is the smallest. Here $k<d$ is a fixed integer. Exhaustive search through the $\binom{d}{k}$ possibilities is often undesirable because of the imposed computational burden. Many attempts have been made to find fast algorithms to obtain the best subset of features. See Fu, Min, and Li (1970), Kanal (1974), Ben-Bassat (1982), and Devijver and Kittler (1982) for surveys. It is easy to see that the best $k$ individual features-that is, components corresponding to the $k$ smallest values of $L^{*}(\{i\})$-do not necessarily constitute the best $k$-dimensional vector: just consider a case in which $X^{(1)} \equiv X^{(2)} \equiv \cdots \equiv X^{(k)}$. Cover and Van Campenhout (1977) showed that any ordering of the $2^{d}$ subsets of $\{1, \ldots, d\}$ consistent with the obvious requirement $L^{*}(A) \leq L^{*}(B)$ if $B \subseteq A$ is possible. More precisely, they proved the following surprising result:

Theorem 32.1. (Cover and Van Campenhout (1977)). Let $A_{1}, A_{2}, \ldots, A_{2^{d}}$ be an ordering of the $2^{d}$ subsets of $\{1, \ldots, d\}$, satisfying the consistency property $i<j$ if $A_{i} \subset A_{j}$. (Therefore, $A_{1}=\emptyset$ and $A_{2^{d}}=\{1, \ldots, d\}$.) Then there exists a distribution of the random variables $(X, Y)=\left(X^{(1)}, \ldots, X^{(d)}, Y\right)$ such that

$$
L^{*}\left(A_{1}\right)>L^{*}\left(A_{2}\right)>\cdots>L^{*}\left(A_{2^{d}}\right) .
$$

The theorem shows that every feature selection algorithm that finds the best $k$ element subset has to search exhaustively through all $k$-element subsets for some distributions. Any other method is doomed to failure for some distribution. Many suboptimal, heuristic algorithms have been introduced trying to avoid the computational demand of exhaustive search (see, e.g., Sebestyen (1962), Meisel (1972),

Chang (1973), Vilmansen (1973), and Devijver and Kittler (1982)). Narendra and Fukunaga (1977) introduced an efficient branch-and-bound method that finds the optimal set of features. Their method avoids searching through all subsets in many cases by making use of the monotonicity of the Bayes error with respect to the partial ordering of the subsets. The key of our proof of the theorem is the following simple lemma:

Lemma 32.1. Let $A_{1}, \ldots, A_{2^{d}}$ be as in Theorem 32.1. Let $1<i<2^{d}$. Assume that the distribution of $(X, Y)$ on $\mathcal{R}^{d} \times\{0,1\}$ is such that the distribution of $X$ is concentrated on a finite set, $L^{*}\left(A_{2^{d}}\right)=L^{*}\left(\left\{1, \ldots, 2^{d}\right\}\right)=0$, and $L^{*}\left(A_{j}\right)<1 / 2$ for each $i<j \leq 2^{d}$. Then there exists another finite distribution such that $L^{*}\left(A_{j}\right)$ remains unchanged for each $j>i$, and

$$
L^{*}\left(A_{j}\right)<L^{*}\left(A_{i}\right)<\frac{1}{2} \text { for each } j>i
$$

Proof. We denote the original distribution of $X$ by $\mu$ and the a posteriori probability function by $\eta$. We may assume without loss of generality that every atom of the distribution of $X$ is in $[0, M)^{d}$ for some $M>0$. Since $L^{*}(\{1, \ldots, d\})=0$, the value of $\eta(x)$ is either zero or one at each atom. We construct the new distribution by duplicating each atom in a special way. We describe the new distribution by defining a measure $\mu^{\prime}$ on $\mathcal{R}^{d}$ and an a posteriori function $\eta^{\prime}: \mathcal{R}^{d} \rightarrow\{0,1\}$.

Define the vector $v_{A_{i}} \in \mathcal{R}^{d}$ such that its $m$-th component equals $M$ if $m \notin A_{i}$, and zero if $m \in A_{i}$. The new measure $\mu^{\prime}$ has twice as many atoms as $\mu$. For each atom $x \in \mathcal{R}^{d}$ of $\mu$, the new measure $\mu^{\prime}$ has two atoms, namely, $x_{1}=x$ and $x_{2}=x+v_{A_{i}}$. The new distribution is specified by
$\mu^{\prime}\left(x_{1}\right)=q \mu(x), \quad \mu^{\prime}\left(x_{2}\right)=(1-q) \mu(x), \quad \eta^{\prime}\left(x_{1}\right)=\eta(x)$, and $\eta^{\prime}\left(x_{2}\right)=1-\eta(x)$,
where $q \in(0,1 / 2)$ is specified later. It remains for us to verify that this distribution satisfies the requirements of the lemma for some $q$. First observe that the values $L^{*}\left(A_{j}\right)$ remain unchanged for all $j>i$. This follows from the fact that there is at least one component in $A_{j}$ along which the new set of atoms is strictly separated from the old one, leaving the corresponding contribution to the Bayes error unchanged. On the other hand, as we vary $q$ from zero to $1 / 2$, the new value of $L^{*}\left(A_{i}\right)$ grows continuously from the old value of $L^{*}\left(A_{i}\right)$ to $1 / 2$. Therefore, since by assumption $\max _{j>i} L^{*}\left(A_{j}\right)<1 / 2$, there exists a value of $q$ such that the new $L^{*}\left(A_{i}\right)$ satisfies $\max _{j>i} L^{*}\left(A_{j}\right)<L^{*}\left(A_{i}\right)<1 / 2$ as desired.

Proof of Theorem 32.1. We construct the desired distribution in $2^{d}-2$ steps, applying Lemma 32.1 in each step. The procedure for $d=3$ is illustrated in Figure 32.2.


Figure 32.2. Construction of a prespecified ordering. In this three-dimensional example, the first four steps of the procedure are shown, when the desired ordering is $L^{*}(\{1,2,3\}) \leq L^{*}(\{2,3\}) \leq L^{*}(\{1,2\}) \leq L^{*}(\{2\}) \leq L^{*}(\{1,3\}) \leq \cdots$. Black circles represent atoms with $\eta=1$ and white circles are those with $\eta=0$.

We start with a monoatomic distribution concentrated at the origin, with $\eta(0)=$ 1. Then clearly, $L^{*}\left(A_{2^{d}}\right)=0$. By Lemma 32.1, we construct a distribution such that $L^{*}\left(A_{2^{d}}\right)=0$ and $0=L^{*}\left(A_{2^{d}}\right)<L^{*}\left(A_{2^{d}-1}\right)<1 / 2$. By applying the lemma again, we can construct a distribution with $0=L^{*}\left(A_{2^{d}}\right)<L^{*}\left(A_{2^{d}-1}\right)<L^{*}\left(A_{2^{d}-2}\right)<$ $1 / 2$. After $i$ steps, we have a distribution satisfying the last $i$ inequalities of the desired ordering. The construction is finished after $2^{d}-2$ steps.

Remark. The original example of Cover and Van Campenhout (1977) uses the multidimensional gaussian distribution. Van Campenhout (1980) developed the idea further by showing that not only all possible orderings, but all possible values of the $L^{*}\left(A_{i}\right)$ 's can be achieved by some distributions. The distribution constructed in the above proof is discrete. It has $2^{2^{d}-2}$ atoms.

One may suspect that feature extraction is much easier if given $Y$, the components $X^{(1)}, \ldots, X^{(d)}$ are conditionally independent. However, three and fourdimensional examples given by Elashoff, Elashoff, and Goldman (1967), Toussaint (1971), and Cover (1974) show that even the individually best two independent components are not the best pair of components. We do not know if Theorem 32.1
generalizes to the case when the components are conditionally independent. In the next example, the pair of components consisting of the two worst single features is the best pair, and vice versa.

Theorem 32.2. (Toussaint (1971)). There exist binary-valued random variables $X_{1}, X_{2}, X_{3}, Y \in\{0,1\}$ such that $X_{1}, X_{2}$, and $X_{3}$ are conditionally independent (given $Y$ ), and

$$
L^{*}(\{1\})<L^{*}(\{2\})<L^{*}(\{3\})
$$

but

$$
L^{*}(\{1,2\})>L^{*}(\{1,3\})>L^{*}(\{2,3\}) .
$$

Proof. Let $\mathbf{P}\{Y=1\}=1 / 2$. Then the joint distribution of $X_{1}, X_{2}, X_{3}, Y$ is specified by the conditional probabilities $\mathbf{P}\left\{X_{i}=1 \mid Y=0\right\}$ and $\mathbf{P}\left\{X_{i}=1 \mid Y=1\right\}$, $i=1,2,3$. Straightforward calculation shows that the values

$$
\begin{aligned}
\mathbf{P}\left\{X_{1}=1 \mid Y=0\right\}=0.1, & \mathbf{P}\left\{X_{1}=1 \mid Y=1\right\}=0.9, \\
\mathbf{P}\left\{X_{2}=1 \mid Y=0\right\}=0.05, & \mathbf{P}\left\{X_{2}=1 \mid Y=1\right\}=0.8, \\
\mathbf{P}\left\{X_{3}=1 \mid Y=0\right\}=0.01, & \mathbf{P}\left\{X_{3}=1 \mid Y=1\right\}=0.71
\end{aligned}
$$

satisfy the stated inequalities.

As our ultimate goal is to minimize the error probability, finding the feature set minimizing the Bayes error is not the best we can do. For example, if we know that we will use the 3 -nearest neighbor rule, then it makes more sense to select the set of features that minimizes the asymptotic error probability $L_{3 \mathrm{NN}}$ of the 3 -nearest neighbor rule. Recall from Chapter 5 that

$$
L_{3 \mathrm{NN}}=\mathbf{E}\{\eta(X)(1-\eta(X))(1+4 \eta(X)(1-\eta(X)))\} .
$$

The situation here is even messier than for the Bayes error. As the next example shows, it is not even true that $A \subset B$ implies $L_{3 \mathrm{NN}}(A) \geq L_{3 \mathrm{NN}}(B)$, where $A, B \subseteq$ $\{1, \ldots, d\}$ are two subsets of components, and $L_{3 \mathrm{NN}}(A)$ denotes the asymptotic error probability of the 3 -nearest neighbor rule for $\left(X_{A}, Y\right)$. In other words, adding components may increase $L_{3 N \mathrm{NN}}$ ! This can never happen to the Bayes error-and in fact, to any $f$-error (Theorem 3.3). The anomaly is due to the fact that the function $x(1-x)(1+4 x(1-x))$ is convex near zero and one.

Example. Let the joint distribution of $X=\left(X_{1}, X_{2}\right)$ be uniform on $[0,2]^{2}$. The joint distribution of $(X, Y)$ is defined by the a posteriori probabilities given by

$$
\eta(x)= \begin{cases}0.1 & \text { if } x \in[0,1 / 2) \times[0,1 / 2) \\ 0 & \text { if } x \in[1 / 2,1] \times[0,1 / 2) \\ 1 & \text { if } x \in[0,1 / 2) \times[1 / 2,1] \\ 0.9 & \text { if } x \in[1 / 2 \times 1] \times[1 / 2,1]\end{cases}
$$

(Figure 32.3). Straightforward calculations show that $L_{3 \mathrm{NN}}(\{1,2\})=0.0612$, while $L_{3 \mathrm{NN}}(\{2\})=0.056525$, a smaller value! $\square$


Figure 32.3. An example when an additionalfeature increases the error probability of the 3nN rule.

Of course, the real measure of the goodness of the selected feature set is the error probability $L\left(g_{n}\right)$ of the classifier designed by using training data $D_{n}$. If the classification rule $g_{n}$ is not known at the stage of feature selection, then the best one can do is to estimate the Bayes errors $L^{*}(A)$ for each set $A$ of features, and select a feature set by minimizing the estimate. Unfortunately, as Theorem 8.5 shows, no method of estimating the Bayes errors can guarantee good performance. If we know what classifier will be used after feature selection, then the best strategy is to select a set of measurements based on comparing estimates of the error probabilities. We do not pursue this question further.

For special cases, we do not need to mount a big search for the best features. Here is a simple example: given $Y=i$, let $X=\left(X^{(1)}, \ldots, X^{(d)}\right)$ have $d$ independent components, where given $Y=i, X^{(j)}$ is normal $\left(m_{j i}, \sigma_{j}^{2}\right)$. It is easy to verify (Problem 32.2) that if $\mathbf{P}\{Y=1\}=1 / 2$, then

$$
L^{*}=\mathbf{P}\left\{N>\frac{r}{2}\right\}
$$

where $N$ is a standard normal random variable, and

$$
r^{2}=\sum_{j=1}^{d}\left(\frac{m_{j 1}-m_{j 0}}{\sigma_{j}}\right)^{2}
$$

is the square of the Mahalanobis distance (see also Duda and Hart (1973, pp. 66-67)). For this case, the quality of the $j$-th feature is measured by

$$
\left(\frac{m_{j 1}-m_{j 0}}{\sigma_{j}}\right)^{2}
$$

We may as well rank these values, and given that we need only $d^{\prime}<d$ features, we are best off taking the $d^{\prime}$ features with the highest quality index.

It is possible to come up with analytic solutions in other special cases as well. For example, Raudys (1976) and Raudys and Pikelis (1980) investigated the dependence of $\mathbf{E}\left\{L\left(g_{n}\right)\right\}$ on the dimension of the feature space in the case of certain linear classifiers and normal distributions. They point out that for a fixed $n$, by increasing the number of features, the expected error probability $\mathbf{E}\left\{L\left(g_{n}\right)\right\}$ first decreases, and then, after reaching an optimum, grows again.

### 32.2 Transformations with Small Distortion

One may view the problem of feature extraction in general as the problem of finding a transformation (i.e., a function) $T: \mathcal{R}^{d} \rightarrow \mathcal{R}^{d}$ so that the Bayes error $L_{T}^{*}$ corresponding to the pair $(T(X), Y)$ is close to the Bayes error $L^{*}$ corresponding to the pair $(X, Y)$. One typical example of such transformations is fine quantization (i.e., discretization) of $X$, when $T$ maps the observed values into a set of finitely, or countably infinitely many values. Reduction of dimensionality of the observations can be put in this framework as well. In the following result we show that small distortion of the observation cannot cause large increase in the Bayes error probability.

Theorem 32.3. (FARAGÓ AND GyÖrfi (1975)). Assume that for a sequence of transformations $T_{n}, n=1,2, \ldots$

$$
\left\|T_{n}(X)-X\right\| \rightarrow 0
$$

in probability, where $\|\cdot\|$ denotes the Euclidean norm in $\mathcal{R}^{d}$. Then, if $L^{*}$ is the Bayes error for $(X, Y)$ and $L_{T_{n}}^{*}$ is the Bayes error for $\left(T_{n}(X), Y\right)$,

$$
L_{T_{n}}^{*} \rightarrow L^{*}
$$

Proof. For arbitrarily small $\epsilon>0$ we can choose a uniformly continuous function $0 \leq \tilde{\eta}(x) \leq 1$ such that

$$
2 \mathbf{E}\{|\eta(X)-\tilde{\eta}(X)|\}<\epsilon .
$$

For any transformation $T$, consider now the decision problem of $(T(X), \tilde{Y})$, where the random variable $\tilde{Y}$ satisfies $\mathbf{P}\{\tilde{Y}=1 \mid X=x\}=\tilde{\eta}(x)$ for every $x \in \mathcal{R}^{d}$. Denote the corresponding Bayes error by $\tilde{L}_{T}^{*}$, and the Bayes error corresponding to the pair $(X, \tilde{Y})$ by $\tilde{L}^{*}$. Obviously

$$
\begin{equation*}
0 \leq L_{T_{n}}^{*}-L^{*} \leq\left|L_{T_{n}}^{*}-\tilde{L}_{T_{n}}^{*}\right|+\left|\tilde{L}_{T_{n}}^{*}-\tilde{L}^{*}\right|+\left|\tilde{L}^{*}-L^{*}\right| \tag{32.1}
\end{equation*}
$$

To bound the first and third terms on the right-hand side, observe that for any transformation $T$,

$$
\mathbf{P}\{Y=1 \mid T(X)\}=\mathbf{E}\{\mathbf{P}\{Y=1 \mid X\} \mid T(X)\}=\mathbf{E}\{\eta(X) \mid T(X)\} .
$$

Therefore, the Bayes error corresponding to the pair $(T(X), Y)$ equals

$$
L_{T}^{*}=\mathbf{E}\{\min (\mathbf{E}\{\eta(X) \mid T(X)\}, 1-\mathbf{E}\{\eta(X) \mid T(X)\})\}
$$

Thus,

$$
\begin{aligned}
\left|L_{T}^{*}-\tilde{L}_{T}^{*}\right|= & \mid \mathbf{E}\{\min (\mathbf{E}\{\eta(X) \mid T(X)\}, 1-\mathbf{E}\{\eta(X) \mid T(X)\})\} \\
& -\mathbf{E}\{\min (\mathbf{E}\{\tilde{\eta}(X) \mid T(X)\}, 1-\mathbf{E}\{\tilde{\eta}(X) \mid T(X)\})\} \mid \\
\leq & \mathbf{E}\{|\mathbf{E}\{\tilde{\eta}(X) \mid T(X)\}-\mathbf{E}\{\eta(X) \mid T(X)\}|\} \\
\leq & \mathbf{E}\{|\tilde{\eta}(X)-\eta(X)|\} \\
& \text { (by Jensen's inequality) } \\
< & \epsilon
\end{aligned}
$$

so the first and third terms of (32.1) are less than $\epsilon$. For the second term, define the decision function

$$
g_{n}(x)= \begin{cases}0 & \text { if } \tilde{\eta}(x) \leq 1 / 2 \\ 1 & \text { otherwise }\end{cases}
$$

which has error probability $\tilde{L}\left(g_{n}\right)=\mathbf{P}\left\{g_{n}\left(T_{n}(X)\right) \neq \tilde{Y}\right\}$. Then $\tilde{L}\left(g_{n}\right) \geq \tilde{L}_{T_{n}}^{*}$, and we have

$$
0 \leq \tilde{L}_{T_{n}}^{*}-\tilde{L}^{*} \leq \tilde{L}\left(g_{n}\right)-\tilde{L}^{*} \leq 2 \mathbf{E}\left\{\left|\tilde{\eta}\left(T_{n}(X)\right)-\tilde{\eta}(X)\right|\right\}
$$

by Theorem 2.2. All we have to show now is that the limit supremum of the above quantity does not exceed $\epsilon$. Let $\delta(\epsilon)$ be the inverse modulus of continuity of the a posteriori probability $\tilde{\eta}$, that is,

$$
\delta(\epsilon)=\sup \{\|x-y\|: 2|\tilde{\eta}(x)-\tilde{\eta}(y)|<\epsilon\} .
$$

For every $\epsilon>0$, we have $\delta(\epsilon)>0$ by the uniform continuity of $\tilde{\eta}$. Now, we have

$$
\begin{aligned}
& 2 \mathbf{E}\left\{\left|\tilde{\eta}\left(T_{n}(X)\right)-\tilde{\eta}(X)\right|\right\} \\
& \quad \leq 2 \mathbf{E}\left\{I_{\left\{\left\|X-T_{n}(X)\right\|>\delta(\epsilon)\right\}}\right\}+2 \mathbf{E}\left\{I_{\left\{\left\|X-T_{n}(X)\right\| \leq \delta(\epsilon)\right\}}\left|\tilde{\eta}\left(T_{n}(X)\right)-\tilde{\eta}(X)\right|\right\} .
\end{aligned}
$$

Clearly, the first term on the right-hand side converges to zero by assumption, while the second term does not exceed $\epsilon$ by the definition of $\delta(\epsilon)$.

REmark. It is clear from the proof of Theorem 32.3 that everything remains true if the observation $X$ takes its values from a separable metric space with metric $\rho$, and the condition of the theorem is modified to $\rho(T(X), X) \rightarrow 0$ in probability. This generalization has significance in curve recognition, when $X$ is a stochastic process. Then Theorem 32.3 asserts that one does not lose much information by using usual discretization methods, such as, for example, Karhunen-Loève series expansion (see Problems 32.3 to 32.5 ).

### 32.3 Admissible and Sufficient Transformations

Sometimes the cost of guessing zero while the true value of $Y$ is one is different from the cost of guessing one, while $Y=0$. These situations may be handled as follows. Define the costs

$$
C(m, l), \quad m, l=0,1 .
$$

Here $C(Y, g(X))$ is the cost of deciding on $g(X)$ when the true label is $Y$. The risk of a decision function $g$ is defined as the expected value of the cost:

$$
R_{g}=\mathbf{E}\{C(Y, g(X))\} .
$$

Note that if

$$
C(m, l)=\left\{\begin{array}{cc}
1 & \text { if } m \neq l \\
0 & \text { otherwise },
\end{array}\right.
$$

then the risk is just the probability of error. Introduce the notation

$$
Q_{m}(x)=\eta(x) C(1, m)+(1-\eta(x)) C(0, m), \quad m=0,1 .
$$

Then we have the following extension of Theorem 2.1:
Theorem 32.4. Define

$$
\widehat{g}(x)=\left\{\begin{array}{cc}
1 & \text { if } Q_{1}(x)>Q_{0}(x) \\
0 & \text { otherwise } .
\end{array}\right.
$$

Then for all decision functions $g$ we have

$$
R_{\widehat{g}} \leq R_{g} .
$$

$R_{\hat{g}}$ is called the Bayes risk. The proof is left as an exercise (see Problem 32.7). Which transformations preserve all the necessary information in the sense that the Bayes error probability corresponding to the pair $(T(X), Y)$ equals that of $(X, Y)$ ? Clearly, every invertible mapping $T$ has this property. However, the practically interesting transformations are the ones that provide some compression of the data. The most efficient of such transformations is the Bayes decision $g^{*}$ itself: $g^{*}$ is specifically designed to minimize the error probability. If the goal is to minimize the Bayes risk with respect to some other cost function than the error probability, then $g^{*}$ generally fails to preserve the Bayes risk. It is natural to ask what transformations preserve the Bayes risk for all possible cost functions. This question has a practical significance, when collecting data and construction of the decision are separated in space or in time. In such cases the data should be transmitted via a communication channel (or should be stored). In both cases there is a need for an efficient data compression rule. In this problem formulation, when getting the data, one may not know the final cost function. Therefore a desirable data compression (transformation) does not increase the Bayes risk for any cost function $C(\cdot, \cdot)$. Here $T$ is a measurable function mapping from $\mathcal{R}^{d}$ to $\mathcal{R}^{d^{\prime}}$ for some positive integer $d^{\prime}$.

DEFINITION 32.1. Let $R_{C, T}^{*}$ denote the Bayes risk for the cost function $C$ and transformed observation $T(X)$. A transformation $T$ is called admissible iffor any cost function $C$,

$$
R_{C, T}^{*}=R_{C}^{*},
$$

where $R_{C}^{*}$ denotes the Bayes risk for the original observation.
Obviously each invertible transformation $T$ is admissible. A nontrivial example of an admissible transformation is

$$
T^{*}(X)=\eta(X)
$$

since according to Theorem 32.4, the Bayes decision for any cost function can be constructed by the a posteriori probability $\eta(x)$ and by the cost function. Surprisingly, this is basically the only such transformation in the following sense:

Theorem 32.5. A transformation $T$ is admissible if and only if there is a mapping $G$ such that

$$
G(T(X))=\eta(X) \quad \text { with probability one. }
$$

Proof. The converse is easy since for such $G$

$$
R_{C}^{*} \leq R_{C, T(X)}^{*} \leq R_{C, G(T(X))}^{*}=R_{C}^{*}
$$

Assume now that $T$ is admissible but such function $G$ does not exist. Then there is a set $A \subset \mathcal{R}^{d}$ such that $\mu(A)>0, T(x)$ is constant on $A$, while all values of $\eta(x)$ are different on $A$, that is, if $x, y \in A$, then $x \neq y$ implies $\eta(x) \neq \eta(y)$. Then there are real numbers $0<c<1$ and $\epsilon>0$, and sets $B, D \subset A$ such that $\mu(B), \mu(D)>0$, and

$$
\begin{aligned}
& 1-\eta(x)>c+\epsilon \text { if } x \in B \\
& 1-\eta(x)<c-\epsilon \text { if } x \in D .
\end{aligned}
$$

Now, choose a cost function with the following values:

$$
c_{0,0}=c_{1,1}=0, \quad c_{0,1}=1, \quad \text { and } \quad c_{1,0}=\frac{1-c}{c}
$$

Then,

$$
\begin{aligned}
Q_{0}(x) & =\eta(x) \\
Q_{1}(x) & =c_{1,0}(1-\eta(x))
\end{aligned}
$$

and the Bayes decision on $B \cup D$ is given by

$$
g^{*}(x)= \begin{cases}0 & \text { if } c<1-\eta(x) \\ 1 & \text { otherwise }\end{cases}
$$

Now, let $g(T(x))$ be an arbitrary decision. Without loss of generality we can assume that $g(T(x))=0$ if $x \in A$. Then the difference between the risk of $g(T(x))$ and the Bayes risk is

$$
\begin{aligned}
R_{g(T(X))}-R_{g^{*}(X)} & =\int_{\mathcal{R}^{d}} Q_{g(T(x))}(x) \mu(d x)-\int_{\mathcal{R}^{d}} Q_{g^{*}(x)}(x) \mu(d x) \\
& =\int_{\mathcal{R}^{d}}\left(Q_{g(T(x))}(x)-Q_{g^{*}(x)}(x)\right) \mu(d x) \\
& \geq \int_{D}\left(Q_{0}(x)-Q_{1}(x)\right) \mu(d x) \\
& =\int_{D}\left(\eta(x)-c_{1,0}(1-\eta(x)) \mu(d x)\right. \\
& =\int_{D}\left(1-\frac{1-\eta(x)}{c}\right) \mu(d x) \geq \frac{\epsilon}{c} \mu(D)>0 .
\end{aligned}
$$

We can give another characterization of admissible transformations by virtue of a well-known concept of mathematical statistics:

Definition 32.2. $T(X)$ is called a sufficient statistic if the random variables $X$, $T(X)$, and $Y$ form a Markov chain in this order. That is, for any set $A, \mathbf{P}\{Y \in$ $A \mid T(X), X\}=\mathbf{P}\{Y \in A \mid T(X)\}$.

Theorem 32.6. A transformation $T$ is admissible if and only if $T(X)$ is a sufficient statistic.

Proof. Assume that $T$ is admissible. Then according to Theorem 32.5 there is a mapping $G$ such that

$$
G(T(X))=\eta(X) \quad \text { with probability one. }
$$

Then

$$
\mathbf{P}\{Y=1 \mid T(X), X\}=\mathbf{P}\{Y=1 \mid X\}=\eta(X)=G(T(X))
$$

and

$$
\begin{aligned}
\mathbf{P}\{Y=1 \mid T(X)\} & =\mathbf{E}\{\mathbf{P}\{Y=1 \mid T(X), X\} \mid T(X)\} \\
& =\mathbf{E}\{G(T(X)) \mid T(X)\} \\
& =G(T(X)),
\end{aligned}
$$

thus, $\mathbf{P}\{Y=1 \mid T(X), X\}=\mathbf{P}\{Y=1 \mid T(X)\}$, therefore $T(X)$ is sufficient. On the other hand, if $T(X)$ is sufficient, then

$$
\mathbf{P}\{Y=1 \mid X\}=\mathbf{P}\{Y=1 \mid T(X), X\}=\mathbf{P}\{Y=1 \mid T(X)\}
$$

so for the choice

$$
G(T(X))=\mathbf{P}\{Y=1 \mid T(X)\}
$$

we have the desired function $G(\cdot)$, and therefore $T$ is admissible. $\square$

Theorem 32.6 states that we may replace $X$ by any sufficient statistic $T(X)$ without altering the Bayes error. The problem with this, in practice, is that we do not know the sufficient statistics because we do not know the distribution. If the distribution of $(X, Y)$ is known to some extent, then Theorem 32.6 may be useful.
EXAMPLE. Assume that it is known that $\eta(x)=e^{-c\|x\|}$ for some unknown $c>0$. Then $\|X\|$ is a sufficient statistic. Thus, for discrimination, we may replace the $d$-dimensional vector $X$ by the 1 -dimensional random variable $\|X\|$ without loss of discrimination power.
EXAMPLE. If $\eta\left(x^{(1)}, x^{(2)}, x^{(3)}\right)=\eta_{0}\left(x^{(1)} x^{(2)}, x^{(2)} x^{(3)}\right)$ for some function $\eta_{0}$, then $\left(X^{(1)} X^{(2)}, X^{(2)} X^{(3)}\right)$ is a 2-dimensional sufficient statistic. For discrimination, there is no need to deal with $X^{(1)}, X^{(2)}, X^{(3)}$. It suffices to extract the features $X^{(1)} X^{(2)}$ and $X^{(2)} X^{(3)}$.

Example. If given $Y=i, X$ is normal with unknown mean $m_{i}$ and diagonal covariance matrix $\sigma^{2} I$ (for unknown $\sigma$ ), then $\eta(x)$ is a function of $\left\|x-m_{1}\right\|^{2}-$ $\left\|x-m_{0}\right\|^{2}$ only for unknown $m_{0}, m_{1}$. Here we have no obvious sufficient statistic. However, if $m_{0}$ and $m_{1}$ are both known, then a quick inspection shows that $x^{T}\left(m_{1}-\right.$ $m_{0}$ ) is a 1-dimensional sufficient statistic. Again, it suffices to look for the simplest possible argument for $\eta$. If $m_{0}=m_{1}=0$ but the covariance matrices are $\sigma_{0}^{2} I$ and $\sigma_{1}^{2} I$ given that $Y=0$ or $Y=1$, then $\|X\|^{2}$ is a sufficient statistic.

In summary, the results of this section are useful for picking out features when some theoretical information is available regarding the distribution of $(X, Y)$.

## Problems and Exercises

Problem 32.1. Consider the pair $(X, Y) \in \mathcal{R}^{d} \times\{0,1\}$ of random variables, and let $X^{(d+1)} \in$ $\mathcal{R}$ be an additional component. Define the augmented random vector $X^{\prime}=\left(X, X^{(d+1)}\right)$. Denote the Bayes errors corresponding to the pairs $(X, Y)$ and $\left(X^{\prime}, Y\right)$ by $L_{X}^{*}$ and $L_{X^{\prime}}^{*}$ respectively. Clearly, $L_{X}^{*} \geq L_{X^{\prime}}^{*}$. Prove that equality holds if and only if

$$
\mathbf{P}\left\{I_{\left\{\eta^{\prime}\left(X^{\prime}\right)>1 / 2\right\}} \neq I_{\{\eta(X)>1 / 2\}}\right\}=0,
$$

where the a posteriori probability functions $\eta: \mathcal{R}^{d} \rightarrow\{0,1\}$ and $\eta^{\prime}: \mathcal{R}^{d+1} \rightarrow\{0,1\}$ are defined as $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}$ and $\eta^{\prime}\left(x^{\prime}\right)=\mathbf{P}\left\{Y=1 \mid X^{\prime}=x^{\prime}\right)$. Hint: Consult with the proof of Theorem 32.5.

Problem 32.2. Let $\mathbf{P}\{Y=1\}=1 / 2$, and given $Y=i, i=0$, 1 , let $X=\left(X^{(1)}, \ldots, X^{(d)}\right)$ have $d$ independent components, where $X^{(j)}$ is normal $\left(m_{j i}, \sigma_{j}^{2}\right)$. Prove that $L^{*}=\mathbf{P}\{N>r / 2\}$, where $N$ is a standard normal random variable, and $r^{2}$ is the square of the Mahalanobis distance: $r^{2}=\sum_{j=1}^{d}\left(\left(m_{j 1}-m_{j 0}\right) / \sigma_{j}\right)^{2}$ (Duda and Hart (1973, pp. 66-67)).

Problem 32.3. Sampling of a stochastic process. Let $X(t), t \in[0,1]$, be a stochastic process (i.e., a collection of real-valued random variables indexed by $t$ ), and let $Y$ be a binary random variable. For integer $N>0$, define $X_{N}^{(i)}=X(i / N), i=1, \ldots, N$. Find sufficient conditions on the function $m(t)=\mathbf{E}\{X(t)\}$ and on the covariance function $K(t, s)=\mathbf{E}\{(X(t)-\mathbf{E}\{X(t)\})(X(s)-\mathbf{E}\{X(s)\})\}$ such that

$$
\lim _{N \rightarrow \infty} L^{*}\left(X_{N}^{(\cdot)}\right)=L^{*}(X(\cdot))
$$

where $L^{*}\left(X_{N}^{(\cdot)}\right)$ is the Bayes error corresponding to $\left(X_{N}^{(1)}, \ldots, X_{N}^{(N)}\right)$, and $L^{*}(X(\cdot))$ is the infimum of the error probabilities of decision functions that map measurable functions into $\{0,1\}$. Hint: Introduce the stochastic process $X_{N}(t)$ as the linear interpolation of $X_{n}^{(1)}, \ldots, X^{(N)}$. Find conditions under which

$$
\lim _{N \rightarrow \infty} \mathbf{E}\left\{\int_{0}^{1}\left(X_{N}(t)-X(t)\right)^{2} d t\right\}=0
$$

and use Theorem 32.3, and the remark following it.
Problem 32.4. Expansion of a stochastic process. Let $X(t), m(t)$, and $K(t, s)$ be as in the previous problem. Let $\psi_{1}, \psi_{2}, \ldots$ be a complete orthonormal system of functions on [0, 1]. Define

$$
X^{(i)}=\int_{0}^{1} X(t) \psi_{i}(t) d t
$$

Find conditions under which

$$
\lim _{N \rightarrow \infty} L^{*}\left(X^{(1)}, \ldots, X^{(N)}\right)=L^{*}(X(\cdot))
$$

Problem 32.5. Extend Theorem 32.3 such that the transformations $T_{n}\left(X, D_{n}\right)$ are allowed to depend on the training data. This extension has significance, because feature extraction algorithms use the training data.

Problem 32.6. For discrete $X$ prove that $T(X)$ is sufficient iff $Y, T(X), X$ form a Markov chain (in this order).

Problem 32.7. Prove Theorem 32.4.
Problem 32.8. Recall the definition of $F$-errors from Chapter 3. Let $F$ be a strictly concave function. Show that $d_{F}(X, Y)=d_{F}(T(X), Y)$ if and only if the transformation $T$ is admissible. Conclude that $L_{\mathrm{NN}}(X, Y)=L_{\mathrm{NN}}(T(X), Y)$ Construct a $T$ and a distribution of ( $X, Y$ ) such that $L^{*}(X, Y)=L^{*}(T(X), Y)$, but $T$ is not admissible.

Problem 32.9. Find sufficient statistics of minimal dimension for the following discrimination problems:
(1) It is known that for two given sets $A$ and $B$ with $A \cap B=\emptyset$, if $Y=1$, we have $X \in A$ and if $Y=0$, then $X \in B$, or vice versa.
(2) Given $Y, X$ is a vector of independent gamma random variables with common unknown shape parameter $a$ and common unknown scale parameter $b$ (i.e., the marginal density of each component of $X$ is of the form $x^{a-1} e^{-x / b} /\left(\Gamma(a) b^{a}\right)$, $x>0$ ). The parameters $a$ and $b$ depend upon $Y$.

Problem 32.10. Let $X$ have support on the surface of a ball of $\mathcal{R}^{d}$ centered at the origin of unknown radius. Find a sufficient statistic for discrimination of dimension smaller than $d$.

Problem 32.11. Assume that the distribution of $X=\left(X^{(1)}, X^{(2)}, X^{(3)}, X^{(4)}\right)$ is such that $X^{(2)} X^{(4)}=1$ and $X^{(1)}+X^{(2)}+X^{(3)}=0$ with probability one. Find a simple sufficient statistic.

## Appendix

In this appendix we summarize some basic definitions and results from the theory of probability. Most proofs are omitted as they may be found in standard textbooks on probability, such as Ash (1972), Shiryayev (1984), Chow and Teicher (1978), Durrett (1991), Grimmett and Stirzaker (1992), and Zygmund (1959). We also give a list of useful inequalities that are used in the text.

## A. 1 Basics of Measure Theory

Definition A.1. Let $S$ be a set, and let $\mathcal{F}$ be a family of subsets of $S . \mathcal{F}$ is called a $\sigma$-algebra if
(i) $\emptyset \in \mathcal{F}$,
(ii) $A \in \mathcal{F}$ implies $A^{c} \in \mathcal{F}$,
(iii) $A_{1}, A_{2}, \ldots \in \mathcal{F}$ implies $\cup_{i=1}^{\infty} A_{i} \in \mathcal{F}$.

A $\sigma$-algebra is closed under complement and union of countably infinitely many sets. Conditions (i) and (ii) imply that $S \in \mathcal{F}$. Moreover, (ii) and (iii) imply that a $\sigma$-algebra is closed under countably infinite intersections.

Definition A.2. Let $S$ be a set, and let $\mathcal{F}$ be a $\sigma$-algebra of subsets of $S$. Then $(S, \mathcal{F})$ is called a measurable space. The elements of $\mathcal{F}$ are called measurable sets.

DEFINITION A.3. If $S=\mathcal{R}^{d}$ and $\mathcal{B}$ is the smallest $\sigma$-algebra containing all rectangles, then $\mathcal{B}$ is called the Borel $\sigma$-algebra. The elements of $\mathcal{B}$ are called Borel sets.

Definition A.4. Let $(S, \mathcal{F})$ be a measurable space and let $f: S \rightarrow \mathcal{R}$ be a function. $f$ is called measurable if for all $B \in \mathcal{B}$

$$
f^{-1}(B)=\{s: f(s) \in B\} \in \mathcal{F}
$$

that is, the inverse image of any Borel set $B$ is in $\mathcal{F}$.
Obviously, if $A$ is a measurable set, then the indicator variable $I_{A}$ is a measurable function. Moreover, finite linear combinations of indicators of measurable sets (called simple functions) are also measurable functions. It can be shown that the set of measurable functions is closed under addition, subtraction, multiplication, and division. Moreover, the supremum and infimum of a sequence of measurable functions, as well as its pointwise limit supremum and limit infimum are measurable.

Definition A.5. Let $(S, \mathcal{F})$ be a measurable space and let $v: \mathcal{F} \rightarrow[0, \infty)$ be a function. $v$ is a measure on $\mathcal{F}$ if
(i) $v(\emptyset)=0$,
(ii) $v$ is $\sigma$-additive, that is, $A_{1}, A_{2}, \ldots \in \mathcal{F}$, and $A_{i} \cap A_{j}=0, i \neq j$ imply that $v\left(\cup_{i=1}^{\infty} A_{i}\right)=\sum_{i=1}^{\infty} v\left(A_{i}\right)$.

In other words, a measure is a nonnegative, $\sigma$-additive set function.
DEfinition A.6. $v$ is a finite measure if $v(S)<\infty$. v is a $\sigma$-finite measure if there are countably many measurable sets $A_{1}, A_{2}, \ldots$ such that $\cup_{i=1}^{\infty} A_{i}=S$ and $v\left(A_{i}\right)<\infty, i=1,2, \ldots$.

DEfinition A.7. The triple $(S, \mathcal{F}, v)$ is a measure space if $(S, \mathcal{F})$ is a measurable space and $v$ is a measure on $\mathcal{F}$.

Definition A.8. The Lebesgue measure $\lambda$ on $\mathcal{R}^{d}$ is a measure on the Borel $\sigma$-algebra of $\mathcal{R}^{d}$ such that the $\lambda$ measure of each rectangle equals to its volume.

## A. 2 The Lebesgue Integral

DEFINITION A.9. Let $(S, \mathcal{F}, v)$ be a measure space and $f=\sum_{i=1}^{n} x_{i} I_{A_{i}}$ a simple function such that the measurable sets $A_{1}, \ldots, A_{n}$ are disjoint. Then the (Lebesgue) integral of $f$ with respect to $v$ is defined by

$$
\int f d v=\int_{S} f(s) v(d s)=\sum_{i=1}^{n} x_{i} v\left(A_{i}\right)
$$

If $f$ is a nonnegative-valued measurable function, then introduce a sequence of simple functions as follows:

$$
f_{n}(s)= \begin{cases}(k-1) / n & \text { if }(k-1) / n \leq f(s)<k / n, \quad k=1,2, \ldots n 2^{n} \\ 2^{n} & \text { if } f(s) \geq 2^{n} .\end{cases}
$$

Then the $f_{n}$ 's are simple functions, and $f_{n}(s) \rightarrow f(s)$ in a monotone nondecreasing fashion. Therefore, the sequence of integrals $\int f_{n} d v$ is monotone nondecreasing, with a limit. The integral $f$ is then defined by

$$
\int f d v=\int_{S} f(s) v(d s)=\lim _{n \rightarrow \infty} \int f_{n} d v
$$

If $f$ is an arbitrary measurable function, then decompose it as a difference of its positive and negative parts,

$$
f(s)=f(s)^{+}-f(s)^{-}=f(s)^{+}-(-f(s))^{+}
$$

where $f^{+}$and $f^{-}$are both nonnegative functions. Define the integral of $f$ by

$$
\int f d v=\int f^{+} d v-\int f^{-} d v
$$

if at least one term on the right-hand side is finite. Then we say that the integral exists. If the integral is finite then $f$ is integrable.

Definition A.10. If $\int f d v$ exists and $A$ is a measurable function: then $\int_{A} f d v$ is defined by

$$
\int_{A} f d \nu=\int f I_{A} d v
$$

Definition A.11. We say that $f_{n} \rightarrow f(\bmod v)$ if

$$
v\left(\left\{s: \lim _{n \rightarrow \infty} f_{n}(s) \neq f(s)\right\}\right)=0
$$

Theorem A.1. (BEPPO-LEVY THEOREM). If $f_{n}(s) \rightarrow f(s)$ in a monotone increasing way for some nonnegative integrable function $f$, then

$$
\int \lim _{n \rightarrow \infty} f_{n} d v=\lim _{n \rightarrow \infty} \int f_{n} d v
$$

Theorem A.2. (Lebesgue's dominated convergence theorem). Assume that $f_{n} \rightarrow f(\bmod v)$ and $\left|f_{n}(s)\right| \leq g(s)$ for $s \in S, n=1,2, \ldots$, where $\int g d v<\infty$. Then

$$
\int \lim _{n \rightarrow \infty} f_{n} d v=\lim _{n \rightarrow \infty} \int f_{n} d v
$$

Theorem A.3. (Fatou's Lemma). Let $f_{1}, f_{2}, \ldots$ be measurable functions.
(i) If there exists a measurable function $g$ with $\int g d v>-\infty$ such that for every $n, f_{n}(s) \geq g(s)$, then

$$
\liminf _{n \rightarrow \infty} \int f_{n} d v \geq \int \liminf _{n \rightarrow \infty} f_{n} d v
$$

(ii) If there is a a measurable function $g$ with $\int g d v<\infty$, such that $f_{n}(s) \leq$ $g(s)$ for every $n$, then

$$
\limsup _{n \rightarrow \infty} \int f_{n} d v \leq \int \limsup _{n \rightarrow \infty} f_{n} d v
$$

Definition A.12. Let $v_{1}$ and $\nu_{2}$ be measures on a measurable space $(S, \mathcal{F})$. We say that $v_{1}$ is absolutely continuous with respect to $\nu_{2}$ if and only if $v_{2}(A)=0$ implies $\nu_{1}(A)=0(A \in \mathcal{F})$. We denote this relation by $\nu_{1} \ll \nu_{2}$.

Theorem A.4. (RADON-NIKODYM THEOREM). Let $\nu_{1}$ and $\nu_{2}$ be measures on the measurable space $(S, \mathcal{F})$ such that $\nu_{1} \ll \nu_{2}$ and $\nu_{2}$ is $\sigma$-finite. Then there exists a measurable function $f$ such that for all $A \in \mathcal{F}$

$$
v_{1}(A)=\int_{A} f d v_{2}
$$

$f$ is unique $\left(\bmod \nu_{2}\right)$. If $\nu_{1}$ is finite, then $f$ has a finite integral.
DEfinition A.13. $f$ is called the density, or Radon-Nikodym derivative of $\nu_{1}$ with respect to $\nu_{2}$. We use the notation $f=d \nu_{1} / d \nu_{2}$.

DEFINITION A.14. Let $v_{1}$ and $v_{2}$ be measures on a measurable space $(S, \mathcal{F})$. If there exists a set $A \in \mathcal{F}$ such that $\nu_{1}\left(A^{c}\right)=0$ and $\nu_{2}(A)=0$, then $\nu_{1}$ is singular with respect to $v_{2}$ (and vice versa).

Theorem A.5. (LEBESGUE DECOMPOSITION THEOREM). If $\mu$ is a $\sigma$-finite measure on a measurable space $(S, \mathcal{F})$, then there exist two unique measures $v_{1}, v_{2}$ such that $\mu=\nu_{1}+\nu_{2}$, where $\nu_{1} \ll \mu$ and $\nu_{2}$ is singular with respect to $\mu$.

Definition A.15. Let $(S, \mathcal{F}, v)$ be a measure space and let $f$ be a measurable function. Then $f$ induces a measure $\mu$ on the Borel $\sigma$-algebra as follows:

$$
\mu(B)=\nu\left(f^{-1}(B)\right), \quad B \in \mathcal{B}
$$

Theorem A.6. Let $v$ be a measure on the Borel $\sigma$-algebra $\mathcal{B}$ of $\mathcal{R}$, and let $f$ and $g$ be measurable functions. Then for all $B \in \mathcal{B}$,

$$
\int_{B} g(x) \mu(d x)=\int_{f^{-1}(B)} g(f(s)) v(d s)
$$

where $\mu$ is induced by $f$.
Definition A.16. Let $v_{1}$ and $v_{2}$ be measures on the measurable spaces $\left(S_{1}, \mathcal{F}_{1}\right)$ and $\left(S_{2}, \mathcal{F}_{2}\right)$, respectively. Let $(S, \mathcal{F})$ be a measurable space such that $S=S_{1} \times S_{2}$, and $F_{1} \times F_{2} \in \mathcal{F}$ whenever $F_{1} \in \mathcal{F}_{1}$ and $F_{2} \in \mathcal{F}_{2}$. v is called the product measure of $\nu_{1}$ and $\nu_{2}$ on $\mathcal{F}$ if for $F_{1} \in \mathcal{F}_{1}$ and $F_{2} \in \mathcal{F}_{2}, \nu\left(F_{1} \times F_{2}\right)=\nu_{1}\left(F_{1}\right) \nu_{2}\left(F_{2}\right)$. The product of more than two measures can be defined similarly.

Theorem A.7. (Fubini's theorem). Let h be a measurable function on the product space $(S, \mathcal{F})$. Then

$$
\begin{aligned}
\int_{S} h(u, v) v(d(u, v)) & =\int_{S_{1}}\left(\int_{S_{2}} h(u, v) \nu_{2}(d v)\right) \nu_{1}(d u) \\
& =\int_{S_{2}}\left(\int_{S_{1}} h(u, v) \nu_{1}(d u)\right) \nu_{2}(d v)
\end{aligned}
$$

assuming that one of the three integrals is finite.

## A. 3 Denseness Results

Lemma A.1. (Cover and Hart (1967)). Let $\mu$ be a probability measure on $\mathcal{R}^{d}$, and define its support set by

$$
A=\operatorname{support}(\mu)=\left\{x: \text { for all } r>0, \mu\left(S_{x, r}\right)>0\right\}
$$

Then $\mu(A)=1$.

Proof. By the definition of $A$,

$$
A^{c}=\left\{x: \mu\left(S_{x, r_{x}}\right)=0 \text { for some } r_{x}>0\right\}
$$

Let $Q$ denote the set of vectors in $\mathcal{R}^{d}$ with rational components (or any countable dense set). Then for each $x \in A^{c}$, there is a $y_{x} \in Q$ with $\left\|x-y_{x}\right\| \leq r_{x} / 3$. This implies $S_{y_{x}, r_{x} / 2} \subset S_{x, r_{x}}$. Therefore, $\mu\left(S_{y_{x}, r_{x} / 2}\right)=0$, and

$$
A^{c} \subset \bigcup_{x \in A^{c}} S_{y_{x}, r_{x} / 2}
$$

The right-hand side is a union of countably many sets of zero measure, and therefore $\mu\left(A^{\varphi}\right)=1$.

DEFINITION A.17. Let $(S, \mathcal{F}, v)$ be a measure space. For a fixed number $p \geq 1$, $L_{p}(\nu)$ denotes the set of all measurable functions satisfying $\int|f|^{p} d \nu<\infty$.

Theorem A.8. For every probability measure $v$ on $\mathcal{R}^{d}$, the set of continuous functions with bounded support is dense in $L_{p}(v)$. In other words, for every $\epsilon>0$ and $f \in L_{p}$ there is a continuous function with bounded support $g \in L_{p}$ such that

$$
\int|f-g|^{p} d v<\epsilon
$$

The following theorem is a rich source of denseness results:

Theorem A.9. (Stone-Weierstrass theorem). Let $\mathcal{F}$ be a family of real-valued continuous functions on a closed bounded subset $B$ of $\mathcal{R}^{d}$. Assume that $\mathcal{F}$ is an algebra, that is, for any $f_{1}, f_{2} \in \mathcal{F}$ and $a, b \in \mathcal{R}$, we have $a f_{1}+b f_{2} \in \mathcal{F}$, and $f_{1} f_{2} \in \mathcal{F}$. Assume furthermore that if $x \neq y$ then there is an $f \in \mathcal{F}$ such that $f(x) \neq f(y)$, and that for each $x \in B$ there exists an $f \in \mathcal{F}$ with $f(x) \neq 0$. Then for every $\epsilon>0$ and continuous function $g: B \rightarrow \mathcal{R}$, there exists an $f \in \mathcal{F}$ such that

$$
\sup _{x \in B}|g(x)-f(x)|<\epsilon
$$

The following two theorems concern differentiation of integrals. Good general references are Wheeden and Zygmund (1977) and de Guzmán (1975):

Theorem A.10. (The Lebesgue density theorem). Let $f$ be a density on $\mathcal{R}^{d}$. Let $\left\{Q_{k}(x)\right\}$ be a sequence of closed cubes centered at $x$ and contracting to $x$. Then

$$
\lim _{k \rightarrow \infty} \frac{\int_{Q_{k}(x)}|f(x)-f(y)| d y}{\lambda\left(Q_{k}(x)\right)}=0
$$

at almost all $x$, where $\lambda$ denotes the Lebesgue measure. Note that this implies

$$
\lim _{k \rightarrow \infty} \frac{\int_{Q_{k}(x)} f(y) d y}{\lambda\left(Q_{k}(x)\right)}=f(x)
$$

at almost all $x$.
Corollary A.1. Let $\mathcal{A}$ be a collection of subsets of $S_{0,1}$ with the property that for all $A \in \mathcal{A}, \lambda(A) \geq c \lambda\left(S_{0,1}\right)$ for some fixed $c>0$. Then for almost all $x$, if $x+r A=\{y:(y-x) / r \in A\}$,

$$
\lim _{r \rightarrow 0} \sup _{A \in \mathcal{A}}\left|\frac{\int_{x+r A} f(y) d y}{\lambda(x+r A)}-f(x)\right|=0
$$

The Lebesgue density theorem also holds if $\left\{Q_{k}(x)\right\}$ is replaced by a sequence of contracting balls centered at $x$, or indeed by any sequence of sets that satisfy

$$
x+b r_{k} S \subseteq Q_{k}(x) \subseteq x+a r_{k} S
$$

where $S$ is the unit ball of $\mathcal{R}^{d}, r_{k} \downarrow 0$, and $0<b \leq a<\infty$ are fixed constants. This follows from the Lebesgue density theorem. It does not hold in general when $\left\{Q_{k}(x)\right\}$ is a sequence of hyperrectangles containing $x$ and contracting to $x$. For that, an additional restriction is needed:

Theorem A.11. (The Jessen-Marcinkiewicz-Zygmund theorem). Let $f$ be $a$ density on $\mathcal{R}^{d}$ with $\int f \log ^{d-1}(1+f) d x<\infty$. Let $\left\{Q_{k}(x)\right\}$ be a sequence of hyperrectangles containing $x$ and for which $\operatorname{diam}\left(Q_{k}(x)\right) \rightarrow 0$. Then

$$
\lim _{k \rightarrow \infty} \frac{\int_{Q_{k}(x)}|f(x)-f(y)| d y}{\lambda\left(Q_{k}(x)\right)}=0
$$

at almost all $x$.

Corollary A.2. If $f$ is a density and $\left\{Q_{k}(x)\right\}$ is as in Theorem A.11, then

$$
\liminf _{k \rightarrow \infty} \frac{1}{\lambda\left(Q_{k}(x)\right)} \int_{Q_{k}(x)} f(y) d y \geq f(x)
$$

at almost all $x$.
To see this, take $g=\min (f, M)$ for large fixed $M$. As $\int g \log ^{d-1}(1+g)<\infty$, by the Jessen-Marcinkiewicz-Zygmund theorem,

$$
\liminf _{k \rightarrow \infty} \frac{1}{\lambda\left(Q_{k}(x)\right)} \int_{Q_{k}(x)} g(y) d y \geq g(x)
$$

at almost all $x$. Conclude by letting $M \rightarrow \infty$ along the integers.

## A. 4 Probability

DEFINITION A.18. A measure space $(\Omega, \mathcal{F}, \mathbf{P})$ is called a probability space if $\mathbf{P}\{\Omega\}=1 . \Omega$ is the sample space or sure event, the measurable sets are called events, and the measurable functions are called random variables. If $X_{1}, \ldots, X_{n}$ are random variables then $X=\left(X_{1}, \ldots, X_{n}\right)$ is a vector-valued random variable.

Definition A.19. Let $X$ be a random variable, then $X$ induces the measure $\mu$ on the Borel $\sigma$-algebra of $\mathcal{R}$ by

$$
\mu(B)=\mathbf{P}\{\{\omega: X(\omega) \in B\}\}=\mathbf{P}\{X \in B\}, \quad B \in \mathcal{B} .
$$

The probability measure $\mu$ is called the distribution of the random variable $X$.
Definition A.20. Let $X$ be a random variable. The expectation of $X$ is the integral of $x$ with respect to the distribution $\mu$ of $X$ :

$$
\mathbf{E}\{X\}=\int_{\mathcal{R}} x \mu(d x)
$$

if it exists.
Definition A.21. Let $X$ be a random variable. The variance of $X$ is

$$
\operatorname{Var}\{X\}=\mathbb{E}\left\{(X-\mathbb{E}\{X\})^{2}\right\}
$$

if $\mathbf{E}\{X\}$ is finite, and $\infty$ if $\mathbf{E}\{X\}$ is not finite or does not exist.
Definition A.22. Let $X_{1}, \ldots, X_{n}$ be random variables. They induce the measure $\mu^{(n)}$ on the Borel $\sigma$-algebra of $\mathcal{R}^{n}$ with the property

$$
\mu^{(n)}\left(B_{1} \times \ldots \times B_{n}\right)=\mathbf{P}\left\{X_{1} \in B_{1}, \ldots, X_{n} \in B_{n}\right\}, \quad B_{1}, \ldots, B_{n} \in \mathcal{B}
$$

$\mu^{(n)}$ is called the joint distribution of the random variables $X_{1}, \ldots, X_{n}$. Let $\mu_{i}$ be the distribution of $X_{i}(i=1, \ldots, n)$. The random variables $X_{1}, \ldots, X_{n}$ are independent if their joint distribution $\mu^{(n)}$ is the product measure of $\mu_{1}, \ldots, \mu_{n}$. The events $A_{1}, \ldots, A_{n} \in \mathcal{F}$ are independent if the random variables $I_{A_{1}}, \ldots, I_{A_{n}}$ are independent.

Fubini's theorem implies the following:
Theorem A.12. If the random variables $X_{1}, \ldots, X_{n}$ are independent and have finite expectations then

$$
\mathbf{E}\left\{X_{1} X_{2} \ldots X_{n}\right\}=\mathbf{E}\left\{X_{1}\right\} \mathbf{E}\left\{X_{2}\right\} \cdots \mathbf{E}\left\{X_{n}\right\}
$$

## A. 5 Inequalities

Theorem A.13. (Cauchy-Schwarz inequality). If the random variables $X$ and $Y$ have finite second moments $\left(\mathbf{E}\left\{X^{2}\right\}<\infty\right.$ and $\left.\mathbf{E}\left\{Y^{2}\right\}<\infty\right)$, then

$$
|\mathbf{E}\{X Y\}| \leq \sqrt{\mathbf{E}\left\{X^{2}\right\} \mathbf{E}\left\{Y^{2}\right\}}
$$

Theorem A.14. (HÖLDER'S inequality). Let $p, q \in(1, \infty)$ such that $(1 / p)+$ $(1 / q)=1$. Let $X$ and $Y$ be random variables such that $\left(\mathbf{E}\left\{\left|X^{p}\right|\right\}\right)^{1 / p}<\infty$ and $\left(\mathbf{E}\left\{\left|Y^{q}\right|\right\}\right)^{1 / q}<\infty$. Then

$$
\mathbf{E}\{|X Y|\} \leq\left(\mathbf{E}\left\{\left|X^{p}\right|\right\}\right)^{1 / p}\left(\mathbf{E}\left\{\left|Y^{q}\right|\right\}\right)^{1 / q}
$$

Theorem A.15. (MARKOv's inequality). Let $X$ be a nonnegative-valued random variable. Then for each $t>0$,

$$
\mathbf{P}\{X \geq t\} \leq \frac{\mathbf{E}\{X\}}{t}
$$

Theorem A.16. (Chebyshev's inequality). Let $X$ be a random variable. Then for each $t>0$,

$$
\mathbf{P}\{|X-\mathbf{E}\{X\}| \geq t\} \leq \frac{\operatorname{Var}\{X\}}{t^{2}}
$$

Theorem A.17. (Chebyshev-Cantelli inequality). Let $t \geq 0$. Then

$$
\mathbf{P}\{X-\mathbf{E} X>t\} \leq \frac{\operatorname{Var}\{X\}}{\operatorname{Var}\{X\}+t^{2}}
$$

Proof. We may assume without loss of generality that $\mathbf{E} X=0$. Then for all $t$

$$
t=\mathbf{E}\{t-X\} \leq \mathbf{E}\left\{(t-X) I_{\{X \leq t\}}\right\} .
$$

Thus for $t \geq 0$ from the Cauchy-Schwarz inequality,

$$
\begin{aligned}
t^{2} & \leq \mathbf{E}\left\{(t-X)^{2}\right\} \mathbf{E}\left\{I_{\{X \leq t\}}^{2}\right\} \\
& =\mathbf{E}\left\{(t-X)^{2}\right\} \mathbf{P}\{X \leq t\} \\
& =\left(\operatorname{Var}\{X\}+t^{2}\right) \mathbf{P}\{X \leq t\}
\end{aligned}
$$

that is,

$$
\mathbf{P}\{X \leq t\} \geq \frac{t^{2}}{\operatorname{Var}\{X\}+t^{2}}
$$

and the claim follows.

Theorem A.18. (JEnSEn's inequality). If $f$ is a real-valued convex function on a finite or infinite interval of $\mathcal{R}$, and $X$ is a random variable with finite expectation, taking its values in this interval, then

$$
f(\mathbf{E}\{X\}) \leq \mathbf{E}\{f(X)\}
$$

Theorem A.19. (Association inequalities). Let $X$ be a real-valued random variable and let $f(x)$ and $g(x)$ be monotone nondecreasing real-valued functions. Then

$$
\mathbf{E}\{f(X) g(X)\} \geq \mathbf{E}\{f(X)\} \mathbf{E}\{g(X)\}
$$

provided that all expectations exist and are finite. If $f$ is monotone increasing and $g$ is monotone decreasing, then

$$
\mathbf{E}\{f(X) g(X)\} \leq \mathbf{E}\{f(X)\} \mathbf{E}\{g(X)\}
$$

Proof. We prove the first inequality. The second follows by symmetry. Let $X$ have distribution $\mu$. Then we write

$$
\begin{aligned}
\mathbf{E} & \{f(X) g(X)\}-\mathbf{E}\{f(X)\} \mathbf{E}\{g(X)\} \\
& =\int f(x) g(x) \mu(d x)-\int f(y) \mu(d y) \int g(x) \mu(d x) \\
& =\int\left(\int[f(x)-f(y)] g(x) \mu(d x)\right) \mu(d y) \\
& =\int\left(\int h(x, y) g(x) \mu(d x)\right) \mu(d y),
\end{aligned}
$$

where $h(x, y)=f(x)-f(y)$. By Fubini's theorem the last integral equals

$$
\begin{aligned}
& \int_{\mathcal{R}^{2}} h(x, y) g(x) \mu^{2}(d x d y) \\
& \quad=\int_{x>y} h(x, y) g(x) \mu^{2}(d x d y)+\int_{x<y} h(x, y) g(x) \mu^{2}(d x d y)
\end{aligned}
$$

since $h(x, x)=0$ for all $x$. Here $\mu^{2}(d x d y)=\mu(d x) \cdot \mu(d y)$. The second integral on the right-hand side is just

$$
\int_{x}\left(\int_{y>x} h(x, y) g(x) \mu(d y)\right) \mu(d x)=\int_{y}\left(\int_{x>y} h(y, x) g(y) \mu(d x)\right) \mu(d y) .
$$

Thus, we have

$$
\begin{aligned}
\mathbf{E}\{ & f(X) g(X)\}-\mathbf{E}\{f(X)\} \mathbf{E}\{g(X)\} \\
& =\int_{\mathcal{R}^{2}} h(x, y) g(x) \mu^{2}(d x d y) \\
& =\int_{y}\left(\int_{x>y}[h(x, y) g(x)+h(y, x) g(y)] \mu(d x)\right) \mu(d y) \\
& =\int_{y}\left(\int_{x>y} h(x, y)[g(x)-g(y)] \mu(d x)\right) \mu(d y) \\
& \geq 0
\end{aligned}
$$

since $h(y, x)=-h(x, y)$, and by the fact that $h(x, y) \geq 0$ and $g(x)-g(y) \geq 0$ if $x>y$.

## A. 6 Convergence of Random Variables

Definition A.23. Let $\left\{X_{n}\right\}, n=1,2, \ldots$, be a sequence of random variables. We say that

$$
\lim _{n \rightarrow \infty} X_{n}=X \quad \text { in probability }
$$

iffor each $\epsilon>0$

$$
\lim _{n \rightarrow \infty} \mathbb{P}\left\{\left|X_{n}-X\right| \geq \epsilon\right\}=0
$$

We say that

$$
\lim _{n \rightarrow \infty} X_{n}=X \quad \text { with probability one (or almost surely), }
$$

if $X_{n} \rightarrow X(\bmod \mathbf{P})$, that is,

$$
\mathbb{P}\left\{\omega: \lim _{n \rightarrow \infty} X_{n}(\omega)=X(\omega)\right\}=1
$$

For a fixed number $p \geq 1$ we say that

$$
\lim _{n \rightarrow \infty} X_{n}=X \quad \text { in } L_{p}
$$

if

$$
\lim _{n \rightarrow \infty} \mathbb{E}\left\{\left|X_{n}-X\right|^{p}\right\}=0
$$

Theorem A.20. Convergence in $L_{p}$ implies convergence in probability.
Theorem A.21. $\lim _{n \rightarrow \infty} X_{n}=X$ with probability one if and only if

$$
\lim _{n \rightarrow \infty} \sup _{n \leq m}\left|X_{m}-X\right|=0
$$

in probability. Thus, convergence with probability one implies convergence in probability.

Theorem A.22. (Borel-CANTELLI LEMMA). Let $A_{n}, n=1,2, \ldots$, be a sequence of events. Introduce the notation

$$
\left[\begin{array}{ll}
A_{n} & \text { i.o. }]
\end{array}=\limsup _{n \rightarrow \infty} A_{n}=\cap_{n=1}^{\infty} \cup_{m=n}^{\infty} A_{m} .\right.
$$

("i.o." stands for "infinitely often.") If

$$
\sum_{n=1}^{\infty} \mathbf{P}\left\{A_{n}\right\}<\infty
$$

then

$$
\mathbf{P}\left\{\left[\begin{array}{ll}
A_{n} & \text { i.o. }]
\end{array}\right]=0 .\right.
$$

By Theorems A. 21 and A.22, we have
Theorem A.23. If for each $\epsilon>0$

$$
\sum_{n=1}^{\infty} \mathbf{P}\left\{\left|X_{n}-X\right| \geq \epsilon\right\}<\infty
$$

then $\lim _{n \rightarrow \infty} X_{n}=X$ with probability one.

## A. 7 Conditional Expectation

If $Y$ is a random variable with finite expectation and $A$ is an event with positive probability, then the conditional expectation of $Y$ given $A$ is defined by

$$
\mathbf{E}\{Y \mid A\}=\frac{\mathbf{E}\left\{Y I_{A}\right\}}{\mathbf{P}\{A\}}
$$

The conditional probability of an event $B$ given $A$ is

$$
\mathbf{P}\{B \mid A\}=\mathbf{E}\left\{I_{B} \mid A\right\}=\frac{\mathbf{P}\{A \cap B\}}{\mathbf{P}\{A\}} .
$$

DEfinition A.24. Let $Y$ be a random variable with finite expectation and $X$ be a $d$-dimensional vector-valued random variable. Let $\mathcal{F}_{X}$ be the $\sigma$-algebra generated by $X$ :

$$
\mathcal{F}_{X}=\left\{X^{-1}(B) ; B \in \mathcal{B}^{n}\right\}
$$

The conditional expectation $\mathbf{E}\{Y \mid X\}$ of $Y$ given $X$ is a random variable with the property that for all $A \in \mathcal{F}_{X}$

$$
\int_{A} Y d \mathbf{P}=\int_{A} \mathbf{E}\{Y \mid X\} d \mathbf{P}
$$

The existence and uniqueness (with probability one) of $\mathbf{E}\{Y \mid X\}$ is a consequence of the Radon-Nikodym theorem if we apply it to the measures

$$
v^{+}(A)=\int_{A} Y^{+} d \mathbf{P}, \quad v^{-}(A)=\int_{A} Y^{-} d \mathbf{P}, \quad A \in \mathcal{F}_{X},
$$

such that

$$
\mathbf{E}\left\{Y^{+} \mid X\right\}=\frac{d v^{+}}{d \mathbf{P}}, \quad \mathbf{E}\left\{Y^{-} \mid X\right\}=\frac{d v^{-}}{d \mathbf{P}}
$$

and

$$
\mathbf{E}\{Y \mid X\}=\mathbf{E}\left\{Y^{+} \mid X\right\}-\mathbf{E}\left\{Y^{-} \mid X\right\} .
$$

Definition A.25. Let $C$ be an event and $X$ be a d-dimensional vector-valued random variable. Then the conditional probability of $C$ given $X$ is $\mathbf{P}\{C \mid X\}=$ $\mathbf{E}\left\{I_{C} \mid X\right\}$.

Theorem A.24. Let $Y$ be a random variable with finite expectation. Let $C$ be an event, and let $X$ and $Z$ be vector-valued random variables. Then
(i) There is a measurable function $g$ on $\mathcal{R}^{d}$ such that $\mathbf{E}\{Y \mid X\}=g(X)$ with probability one.
(ii) $\mathbf{E}\{Y\}=\mathbf{E}\{\mathbf{E}\{Y \mid X\}\}, \quad \mathbf{P}\{C\}=\mathbf{E}\{\mathbf{P}\{C \mid X\}\}$.
(iii) $\mathbf{E}\{Y \mid X\}=\mathbf{E}\{\mathbf{E}\{Y \mid X, Z\} \mid X\}, \quad \mathbf{P}\{C \mid X\}=\mathbf{E}\{\mathbf{P}\{C \mid X, Y\} \mid X\}$.
(iv) If $Y$ is a function of $X$ then $\mathbf{E}\{Y \mid X\}=Y$.
(v) If $(Y, X)$ and $Z$ are independent, then $\mathbf{E}\{Y \mid X, Z\}=\mathbf{E}\{Y \mid X\}$.
(vi) If $Y=f(X, Z)$ for a measurable function $f$, and $X$ and $Z$ are independent, then $\mathbf{E}\{Y \mid X\}=g(X)$, where $g(x)=\mathbf{E}\{f(x, Z)\}$.

## A. 8 The Binomial Distribution

An integer-valued random variable $X$ is said to be binomially distributed with parameters $n$ and $p$ if

$$
\mathbf{P}\{X=k\}=\binom{n}{k} p^{k}(1-p)^{n-k}, \quad k=0,1, \ldots, n
$$

If $A_{1}, \ldots, A_{n}$ are independent events with $\mathbf{P}\left\{A_{i}\right\}=p$, then $X=\sum_{i=1}^{n} I_{A_{i}}$ is binomial $(n, p) . I_{A_{i}}$ is called a Bernoulli random variable with parameter $p$.

Lemma A.2. Let the random variable $B(n, p)$ be binomially distributed with $p a$ rameters $n$ and $p$. Then
(i)

$$
\mathbf{E}\left\{\frac{1}{1+B(n, p)}\right\} \leq \frac{1}{(n+1) p}
$$

and
(ii)

$$
\mathbf{E}\left\{\frac{1}{B(n, p)} I_{\{B(n, p)>0\}}\right\} \leq \frac{2}{(n+1) p}
$$

Proof. (i) follows from the following simple calculation:

$$
\begin{aligned}
\mathbf{E}\left\{\frac{1}{1+B(n, p)}\right\} & =\sum_{k=0}^{n} \frac{1}{k+1}\binom{n}{k} p^{k}(1-p)^{n-k} \\
& =\frac{1}{(n+1) p} \sum_{k=0}^{n}\binom{n+1}{k+1} p^{k+1}(1-p)^{n-k} \\
& \leq \frac{1}{(n+1) p} \sum_{k=0}^{n+1}\binom{n+1}{k} p^{k}(1-p)^{n-k+1}=\frac{1}{(n+1) p}
\end{aligned}
$$

For (ii) we have

$$
\mathbf{E}\left\{\frac{1}{B(n, p)} I_{\{B(n, p)>0\}}\right\} \leq \mathbf{E}\left\{\frac{2}{1+B(n, p)}\right\} \leq \frac{2}{(n+1) p}
$$

by (i).

Lemma A.3. Let $B$ be a binomial random variable with parameters $n$ and $p$. Then for every $0 \leq p \leq 1$,

$$
\mathbf{P}\left\{B=\left\lfloor\frac{n}{2}\right\rfloor\right\} \leq \sqrt{\frac{1}{2 n \pi}}
$$

and for $p=1 / 2$

$$
\mathbf{P}\left\{B=\left\lfloor\frac{n}{2}\right\rfloor\right\} \geq \sqrt{\frac{1}{2 n \pi}} \frac{1}{e^{1 / 12}}
$$

Proof. The lemma follows from Stirling's formula:

$$
\sqrt{2 \pi n}\left(\frac{n}{e}\right)^{n} e^{1 /(12 n)} \leq n!\leq \sqrt{2 \pi n}\left(\frac{n}{e}\right)^{n} e^{1 /(12 n+1)}
$$

(see, e.g., Feller (1968)).

Lemma A.4. (Devroye and Györfi (1985), p. 194). For any random variable $X$ with finite fourth moment,

$$
\mathbf{E}\{|X|\} \geq \frac{\left(\mathbf{E}\left\{X^{2}\right\}\right)^{3 / 2}}{\left(\mathbf{E}\left\{X^{4}\right\}\right)^{1 / 2}}
$$

Proof. Fix $a>0$. The function $1 / x+a x^{2}$ is minimal on $(0, \infty)$ when $x^{3}=1 /(2 a)$. Thus,

$$
\frac{x+a x^{4}}{x^{2}} \geq(2 a)^{1 / 3}+\frac{a}{(2 a)^{2 / 3}}=\frac{3}{2}(2 a)^{1 / 3}
$$

Replace $x$ by $|X|$ and take expectations:

$$
\mathbf{E}\{|X|\} \geq \frac{3}{2}(2 a)^{1 / 3} \mathbf{E}\left\{X^{2}\right\}-a \mathbf{E}\left\{X^{4}\right\}
$$

The lower bound, considered as a function of $a$, is maximized if we take $a=$ $\frac{1}{2}\left(\mathbb{E}\left\{X^{2}\right\} / \mathbb{E}\left\{X^{4}\right\}\right)^{3 / 2}$. Resubstitution yields the given inequality.

Lemma A.5. Let $B$ be a binomial $(n, 1 / 2)$ random variable. Then

$$
\mathbf{E}\left\{\left|B-\frac{n}{2}\right|\right\} \geq \sqrt{\frac{n}{8}}
$$

Proof. This bound is a special case of Khintchine's inequality (see Szarek (1976), Haagerup (1978), and also Devroye and Györf (1985), p. 139). Rather than proving the given inequality, we will show how to apply the previous lemma to get (without further work) the inequality

$$
\mathbf{E}\left\{\left|B-\frac{n}{2}\right|\right\} \geq \sqrt{\frac{n}{12}}
$$

Indeed, $\mathbb{E}\left\{(B-n / 2)^{2}\right\}=n / 4$ and $\mathbf{E}\left\{(B-n / 2)^{4}\right\}=3 n^{2} / 16-n / 8 \leq 3 n^{2} / 16$. Thus,

$$
\mathbf{E}\left\{\left|B-\frac{n}{2}\right|\right\} \geq \frac{(n / 4)^{3 / 2}}{\left(3 n^{2} / 16\right)^{1 / 2}}=\sqrt{\frac{n}{12}}
$$

Lemma A.6. (SLud (1977)). Let $B$ be a binomial ( $n, p$ ) random variable with $p \leq 1 / 2$. Then for $n(1-p) \geq k \geq n p$,

$$
\mathbf{P}\{B \geq k\} \geq \mathbf{P}\left\{N \geq \frac{k-n p}{\sqrt{n p(1-p)}}\right\}
$$

where $N$ is normal $(0,1)$.

## A. 9 The Hypergeometric Distribution

Let $N, b$, and $n$ be positive integers with $N>n$ and $N>b$. A random variable $X$ taking values on the integers $0,1, \ldots, b$ is hypergeometric with parameters $N, b$ and $n$, if

$$
\mathbf{P}\{X=k\}=\frac{\binom{b}{k}\binom{N-b}{n-k}}{\binom{N}{n}}, \quad k=1, \ldots, b
$$

$X$ models the number of blue balls in a sample of $n$ balls drawn without replacement from an urn containing $b$ blue and $N-b$ red balls.

Theorem A.25. (Hoeffding (1963)). Let the set A consist of $N$ numbers $a_{1}, \ldots$, $a_{N}$. Let $Z_{1}, \ldots, Z_{n}$ denote a random sample taken without replacement from $A$, where $n \leq N$. Denote

$$
m=\frac{1}{N} \sum_{i=1}^{N} a_{i} \quad \text { and } \quad c=\max _{i, j \leq N}\left|a_{i}-a_{j}\right|
$$

Then for any $\epsilon>0$ we have

$$
\mathbf{P}\left\{\left|\frac{1}{n} \sum_{i=1}^{n} Z_{i}-m\right| \geq \epsilon\right\} \leq 2 e^{-2 n \epsilon^{2} / c^{2}}
$$

Specifically, if $X$ is hypergeometrically distributed with parameters $N, b$, and $n$, then

$$
\mathbf{P}\left\{\left|X-\frac{b n}{N}\right| \geq n \epsilon\right\} \leq 2 e^{-2 n \epsilon^{2}}
$$

For more inequalities of this type, see Hoeffding (1963) and Serfling (1974).

## A. 10 The Multinomial Distribution

A vector $\left(N_{1}, \ldots, N_{k}\right)$ of integer-valued random variables is multinomially distributed with parameters $\left(n, p_{1}, \ldots, p_{k}\right)$ if

$$
\mathbf{P}\left\{N_{1}=i_{1}, \ldots, N_{k}=i_{k}\right\}= \begin{cases}\frac{n!}{i_{1}!\cdots i_{k}!} p_{1}^{i_{1}} \cdots p_{k}^{i_{k}} & \text { if } \sum_{j=1}^{k} i_{j}=k, \quad i_{j} \geq 0 \\ 0 & \text { otherwise }\end{cases}
$$

Lemma A.7. The moment-generating function of a multinomial $\left(k, p_{1}, \ldots, p_{k}\right)$ vector is

$$
\mathbf{E}\left\{e^{t_{1} N_{1}+\cdots+t_{k} N_{k}}\right\}=\left(p_{1} e^{t_{1}}+\cdots+p_{k} e^{t_{k}}\right)^{n}
$$

## A. 11 The Exponential and Gamma Distributions

A nonnegative random variable has exponential distribution with parameter $\lambda>0$ if it has a density

$$
f(x)=\lambda e^{-\lambda x}, \quad x \geq 0 .
$$

A nonnegative-valued random variable has the gamma distribution with parameters $a, b \geq 0$ if it has density

$$
f(x)=\frac{x^{a-1} e^{-x / b}}{\Gamma(a) b^{a}}, \quad x \geq 0, \quad \text { where } \quad \Gamma(a)=\int_{0}^{\infty} x^{a-1} e^{-x} d x .
$$

The sum of $n$ i.i.d. exponential $(\lambda)$ random variables has gamma distribution with parameters $n$ and $1 / \lambda$.

## A. 12 The Multivariate Normal Distribution

A $d$-dimensional random variable $X=\left(X^{(1)}, \ldots, X^{(d)}\right)$ has the multivariate normal distribution if it has a density

$$
f(x)=\frac{1}{\sqrt{(2 \pi)^{d} \operatorname{det}(\Sigma)}} e^{-\frac{1}{2}(x-m)^{T} \Sigma^{-1}(x-m)},
$$

where $m \in \mathcal{R}^{d}, \Sigma$ is a positive definite symmetric $d \times d$ matrix with entries $\sigma_{i j}$, and $\operatorname{det}(\Sigma)$ denotes the determinant of $\Sigma$. Then $\mathbf{E} X=m$, and for all $i, j=1, \ldots, d$,

$$
\mathbf{E}\left\{\left(X^{(i)}-\mathbf{E} X^{(i)}\right)\left(X^{(j)}-\mathbf{E} X^{(j)}\right)\right\}=\sigma_{i j} .
$$

$\Sigma$ is called the covariance matrix of $X$.

## Notation

- $I_{A}$ indicator of an event $A$.
- $I_{B}(x)=I_{\{x \in B\}}$ indicator function of a set $B$.
- $|A|$ cardinality of a finite set $A$.
- $A^{c}$ complement of a set $A$.
- $A \triangle B$ symmetric difference of sets $A, B$.
- $f \circ g$ composition of functions $f, g$.
- log natural logarithm (base e).
- $\lfloor x\rfloor$ integer part of the real number $x$.
- $\lceil x\rceil$ upper integer part of the real number $x$.
- $X \stackrel{\mathcal{L}}{=} Z$ if $X$ and $Z$ have the same distribution.
- $x^{(1)}, \ldots, x^{(d)}$ components of the $d$-dimensional column vector $x$.
- $\|x\|=\sqrt{\sum_{i=1}^{d}\left(x^{(i)}\right)^{2}} \quad L_{2}$-norm of $x \in \mathcal{R}^{d}$.
- $X \in \mathcal{R}^{d}$ observation, vector-valued random variable.
- $Y \in\{0,1\}$ label, binary random variable.
- $D_{n}=\left(\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)\right)$ training data, sequence of i.i.d. pairs that are independent of $(X, Y)$, and have the same distribution as that of $(X, Y)$.
- $\eta(x)=\mathbf{P}\{Y=1 \mid X=x\}, 1-\eta(x)=\mathbf{P}\{Y=0 \mid X=x\} \quad$ a posteriori probabilities.
- $p=\mathbf{P}\{Y=1\}, 1-p=\mathbf{P}\{Y=0\}$ class probabilities.
- $g^{*}: \mathcal{R}^{d} \rightarrow\{0,1\}$ Bayes decision function.
- $\phi: \mathcal{R}^{d} \rightarrow\{0,1\}, g_{n}: \mathcal{R}^{d} \times\left\{\mathcal{R}^{d} \times\{0,1\}\right\}^{n} \rightarrow\{0,1\} \quad$ classification functions. The short notation $g_{n}(x)=g_{n}\left(x, D_{n}\right)$ is also used.
- $L^{*}=\mathbf{P}\left\{g^{*}(X) \neq Y\right\}$ Bayes risk, the error probability of the Bayes decision.
- $L_{n}=L\left(g_{n}\right)=\mathbf{P}\left\{g_{n}\left(X, D_{n}\right) \neq Y \mid D_{n}\right\}$ error probability of a classification function $g_{n}$.
- $\hat{L}_{n}(\phi)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{\phi\left(X_{i}\right) \neq Y_{i}\right\}} \quad$ empirical error probability of a classifier $\phi$.
- $\mu(A)=\mathbf{P}\{X \in A\}$ probability measure of $X$.
- $\mu_{n}(A)=\frac{1}{n} \sum_{i=1}^{n} I_{\left\{X_{i} \in A\right\}}$ empirical measure corresponding to $X_{1}, \ldots, X_{n}$.
- $\lambda$ Lebesgue measure on $\mathcal{R}^{d}$.
- $f(x)$ density of $X$, Radon-Nikodym derivative of $\mu$ with respect to $\lambda$ (if it exists).
- $f_{0}(x), f_{1}(x)$ conditional densities of $X$ given $Y=0$ and $Y=1$, respectively (if they exist).
- $\mathcal{P}$ partition of $\mathcal{R}^{d}$.
- $X_{(k)}(x), X_{(k)} k$-th nearest neighbor of $x$ among $X_{1}, \ldots, X_{n}$.
- $K: \mathcal{R}^{d} \rightarrow \mathcal{R}$ kernel function.
- $h, h_{n}>0$ smoothing factor for a kernel rule
- $K_{h}(x)=(1 / h) K(x / h)$ scaled kernel function.
- $T_{m}=\left(\left(X_{n+1}, Y_{n+1}\right), \ldots,\left(X_{n+m}, Y_{n+m}\right)\right)$ testing data, sequence of i.i.d. pairs that are independent of $(X, Y)$ and $D_{n}$, and have the same distribution as that of ( $X, Y$ ).
- $\mathcal{A}$ class of sets.
- $\mathcal{C}, \mathcal{C}_{n}$ classes of classification functions.
- $s(\mathcal{A}, n) \quad n$-th shatter coefficient of the class of sets $\mathcal{A}$.
- $V_{\mathcal{A}}$ Vapnik-Chervonenkis dimension of the class of sets $\mathcal{A}$.
- $\mathcal{S}(\mathcal{C}, n) \quad n$-th shatter coefficient of the class of classifiers $\mathcal{C}$.
- $V_{\mathcal{C}}$ Vapnik-Chervonenkis dimension of the class of classifiers $\mathcal{C}$.
- $S_{x, r}=\left\{y \in \mathcal{R}^{d}:\|y-x\| \leq r\right\} \quad$ closed Euclidean ball in $\mathcal{R}^{d}$ centered at $x \in \mathcal{R}^{d}$, with radius $r>0$.


## References

Abou-Jaoude, S. (1976a). Conditions nécessaires et suffisantes de convergence $L_{1}$ en probabilité de l'histogramme pour une densité. Annales de l'Institut Henri Poincaré, 12:213231.

Abou-Jaoude, S. (1976b). La convergence $L_{1}$ et $L_{\infty}$ de l'estimateur de la partition aléatoire pour une densité. Annales de l'Institut Henri Poincaré, 12:299-317.
Abou-Jaoude, S. (1976c). Sur une condition nécessaire et suffisante de $L_{1}$-convergence presque complète de l'estimateur de la partition fixe pour une densité. Comptes Rendus de l'Académie des Sciences de Paris, 283:1107-1110.
Aitchison, J. and Aitken, C. (1976). Multivariate binary discrimination by the kernel method. Biometrika, 63:413-420.
Aizerman, M., Braverman, E., and Rozonoer, L. (1964a). The method of potential functions for the problem of restoring the characteristic of a function converter from randomly observed points. Automation and Remote Control, 25:1546-1556.
Aizerman, M., Braverman, E., and Rozonoer, L. (1964b). The probability problem of pattern recognition learning and the method of potential functions. Automation and Remote Control, 25:1307-1323.
Aizerman, M., Braverman, E., and Rozonoer, L. (1964c). Theoretical foundations of the potential function method in pattern recognition learning. Automation and Remote Control, 25:917-936.
Aizerman, M., Braverman, E., and Rozonoer, L. (1970). Extrapolative problems in automatic control and the method of potential functions. American Mathematical Society Translations, 87:281-303.
Akaike, H. (1954). An approximation to the density function. Annals of the Institute of Statistical Mathematics, 6:127-132.
Akaike, H. (1974). A new look at the statistical model identification. IEEE Transactions on Automatic Control, 19:716-723.

Alexander, K. (1984). Probability inequalities for empirical processes and a law of the iterated logarithm. Annals of Probability, 4:1041-1067.
Anderson, A. and Fu, K. (1979). Design and development of a linear binary tree classifier for leukocytes. Technical Report TR-EE-79-31, Purdue University, Lafayette, IN.
Anderson, J. (1982). Logistic discrimination. In Handbook of Statistics, Krishnaiah, P. and Kanal, L., editors, volume 2, pages 169-191. North-Holland, Amsterdam.
Anderson, M. and Benning, R. (1970). A distribution-free discrimination procedure based on clustering. IEEE Transactions on Information Theory, 16:541-548.
Anderson, T. (1958). An Introduction to Multivariate Statistical Analysis. John Wiley, New York.
Anderson, T. (1966). Some nonparametric multivariate procedures based on statistically equivalent blocks. In Multivariate Analysis, Krishnaiah, P., editor, pages 5-27. Academic Press, New York.
Angluin, D. and Valiant, L. (1979). Fast probabilistic algorithms for Hamiltonian circuits and matchings. Journal of Computing System Science, 18:155-193.
Anthony, M. and Holden, S. (1993). On the power of polynomial discriminators and radial basis function networks. In Proceedings of the Sixth Annual ACM Conference on Computational Learning Theory, pages 158-164. Association for Computing Machinery, New York.
Anthony, M. and Shawe-Taylor, J. (1990). A result of Vapnik with applications. Technical Report CSD-TR-628, University of London, Surrey.
Argentiero, P., Chin, R., and Beaudet, P. (1982). An automated approach to the design of decision tree classifiers. IEEE Transactions on Pattern Analysis and Machine Intelligence, 4:51-57.
Arkadjew, A. and Braverman, E. (1966). Zeichenerkennung und Maschinelles Lernen. Oldenburg Verlag, München, Wien.
Ash, R. (1972). Real Analysis and Probability. Academic Press, New York.
Assouad, P. (1983a). Densité et dimension. Annales de l'Institut Fourier, 33:233-282.
Assouad, P. (1983b). Deux remarques sur l'estimation. Comptes Rendus de l'Académie des Sciences de Paris, 296:1021-1024.
Azuma, K. (1967). Weighted sums of certain dependent random variables. Tohoku Mathematical Journal, 68:357-367.
Bahadur, R. (1961). A representation of the joint distribution of responses to $n$ dichotomous items. In Studies in Item Analysis and Prediction, Solomon, H., editor, pages 158-168. Stanford University Press, Stanford, CA.
Bailey, T. and Jain, A. (1978). A note on distance-weighted $k$-nearest neighbor rules. IEEE Transactions on Systems, Man, and Cybernetics, 8:311-313.
Barron, A. (1985). Logically smooth density estimation. Technical Report TR 56, Department of Statistics, Stanford University, Stanford, CA.
Barron, A. (1989). Statistical properties of artificial neural networks. In Proceedings of the 28th Conference on Decision and Control, pages 280-285. Tampa, FL.
Barron, A. (1991). Complexity regularization with application to artificial neural networks. In Nonparametric Functional Estimation and Related Topics, Roussas, G., editor, pages 561-576. NATO ASI Series, Kluwer Academic Publishers, Dordrecht.
Barron, A. (1993). Universal approximation bounds for superpositions of a sigmoidal function. IEEE Transactions on Information Theory, 39:930-944.
Barron, A. (1994). Approximation and estimation bounds for artificial neural networks. Machine Learning, 14:115-133.

Barron, A. and Barron, R. (1988). Statistical learning networks: a unifying view. In Proceedings of the 20th Symposium on the Interface: Computing Science and Statistics, Wegman, E., Gantz, D., and Miller, J., editors, pages 192-203. AMS, Alexandria, VA.
Barron, A. and Cover, T. (1991). Minimum complexity density estimation. IEEE Transactions on Information Theory, 37:1034-1054.
Barron, A., Györfi, L., and van der Meulen, E. (1992). Distribution estimation consistent in total variation and in two types of information divergence. IEEE Transactions on Information Theory, 38:1437-1454.
Barron, R.(1975). Learning networks improve computer-aided prediction and control. Computer Design, 75:65-70.
Bartlett, P. (1993). Lower bounds on the Vapnik-Chervonenkis dimension of multi-layer threshold networks. In Proceedings of the Sixth annual ACM Conference on Computational Learning Theory, pages 144-150. Association for Computing Machinery, New York.
Bashkirov, O., Braverman, E., and Muchnik, I. (1964). Potential function algorithms for pattern recognition learning machines. Automation and Remote Control, 25:692-695.
Baum, E. (1988). On the capabilities of multilayer perceptrons. Journal of Complexity, 4:193-215.
Baum, E. and Haussler, D. (1989). What size net gives valid generalization? Neural Computation, 1:151-160.
Beakley, G. and Tuteur, F. (1972). Distribution-free pattern verification using statistically equivalent blocks. IEEE Transactions on Computers, 21:1337-1347.
Beck, J. (1979). The exponential rate of convergence of error for $k_{n}$-NN nonparametric regression and decision. Problems of Control and Information Theory, 8:303-311.
Becker, P. (1968). Recognition of Patterns. Polyteknisk Forlag, Copenhagen.
Ben-Bassat, M. (1982). Use of distance measures, information measures and error bounds in feature evaluation. In Handbook of Statistics, Krishnaiah, P. and Kanal, L., editors, volume 2, pages 773-792. North-Holland, Amsterdam.
Benedek, G. and Itai, A. (1988). Learnability by fixed distributions. In Computational Learning Theory: Proceedings of the 1988 Workshop, pages 80-90. Morgan Kaufman, San Mateo, CA.
Benedek, G. and Itai, A. (1994). Nonuniform learnability. Journal of Computer and Systems Sciences, 48:311-323.
Bennett, G. (1962). Probability inequalities for the sum of independent random variables. Journal of the American Statistical Association, 57:33-45.
Beran, R. (1977). Minimum Hellinger distance estimates for parametric models. Annals of Statistics, 5:445-463.
Beran, R. (1988). Comments on "a new theoretical and algorithmical basis for estimation, identification and control" by P. Kovanic. Automatica, 24:283-287.
Bernstein, S. (1946). The Theory of Probabilities. Gastehizdat Publishing House, Moscow.
Bhattacharya, P. and Mack, Y. (1987). Weak convergence of $k$-nN density and regression estimators with varying $k$ and applications. Annals of Statistics, 15:976-994.
Bhattacharyya, A. (1946). On a measure of divergence between two multinomial populations. Sankhya, Series A, 7:401-406.
Bickel, P. and Breiman, L. (1983). Sums of functions of nearest neighbor distances, moment bounds, limit theorems and a goodness of fit test. Annals of Probability, 11:185-214.
Birgé, L. (1983). Approximation dans les espaces métriques et théorie de l'estimation. Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete, 65:181-237.

Birgé, L. (1986). On estimating a density using Hellinger distance and some other strange facts. Probability Theory and Related Fields, 71:271-291.
Blumer, A., Ehrenfeucht, A., Haussler, D., and Warmuth, M. (1989). Learnability and the Vapnik-Chervonenkis dimension. Journal of the ACM, 36:929-965.
Braverman, E. (1965). The method of potential functions. Automation and Remote Control, 26:2130-2138.
Braverman, E. and Pyatniskii, E. (1966). Estimation of the rate of convergence of algorithms based on the potential function method. Automation and Remote Control, 27:80-100.
Breiman, L., Friedman, J., Olshen, R., and Stone, C. (1984). Classification and Regression Trees. Wadsworth International, Belmont, CA.
Breiman, L., Meisel, W., and Purcell, E. (1977). Variable kernel estimates of multivariate densities. Technometrics, 19:135-144.
Brent, R. (1991). Fast training algorithms for multilayer neural nets. IEEE Transactions on Neural Networks, 2:346-354.
Broder, A. (1990). Strategies for efficient incremental nearest neighbor search. Pattern Recognition, 23:171-178.
Broomhead, D. and Lowe, D. (1988). Multivariable functional interpolation and adaptive networks. Complex Systems, 2:321-323.
Buescher, K. and Kumar, P. (1996a). Learning by canonical smooth estimation, Part I: Simultaneous estimation. IEEE Transactions on Automatic Control, 41:545-556.
Buescher, K. and Kumar, P. (1996b). Learning by canonical smooth estimation, Part II: Learning and choice of model complexity. IEEE Transactions on Automatic Control, 41:557-569.
Burbea, J. (1984). The convexity with respect to gaussian distributions of divergences of order $\alpha$. Utilitas Mathematica, 26:171-192.
Burbea, J. and Rao, C. (1982). On the convexity of some divergence measures based on entropy functions. IEEE Transactions on Information Theory, 28:48-495.
Burshtein, D., Della Pietra, V., Kanevsky, D., and Nádas, A. (1992). Minimum impurity partitions. Annals of Statistics, 20:1637-1646.
Cacoullos, T. (1965). Estimation of a multivariate density. Annals of the Institute of Statistical Mathematics, 18:179-190.
Carnal, H. (1970). Die konvexe Hülle von $n$ rotationssymmetrisch verteilten Punkten. Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete, 15:168-176.
Casey, R. and Nagy, G. (1984). Decision tree design using a probabilistic model. IEEE Transactions on Information Theory, 30:93-99.
Cencov, N. (1962). Evaluation of an unknown distribution density from observations. Soviet Math. Doklady, 3:1559-1562.
Chang, C. (1974). Finding prototypes for nearest neighbor classifiers. IEEE Transactions on Computers, 26:1179-1184.
Chang, C. (1973). Dynamic programming as applied to feature selection in pattern recognition systems. IEEE Transactions on Systems, Man, and Cybernetics, 3:166-171.
Chen, T., Chen, H., and Liu, R. (1990). A constructive proof and an extension of Cybenko's approximation theorem. In Proceedings of the 22nd Symposium of the Interface: Computing Science and Statistics, pages 163-168. American Statistical Association, Alexandria, VA.
Chen, X. and Zhao, L. (1987). Almost sure $L_{1}$-norm convergence for data-based histogram density estimates. Journal of Multivariate Analysis, 21:179-188.

Chen, Z. and Fu, K. (1973). Nonparametric Bayes risk estimation for pattern classification. In Proceedings of the IEEE Conference on Systems, Man, and Cybernetics. Boston, MA.
Chernoff, H. (1952). A measure of asymptotic efficiency of tests of a hypothesis based on the sum of observations. Annals of Mathematical Statistics, 23:493-507.
Chernoff, H. (1971). A bound on the classification error for discriminating between populations with specified means and variances. In Studi di probabilità, statistica e ricerca operativa in onore di Giuseppe Pompilj, pages 203-211. Oderisi, Gubbio.
Chou, P. (1991). Optimal partitioning for classification and regression trees. IEEE Transactions on Pattern Analysis and Machine Intelligence, 13:340-354.
Chou, W. and Chen, Y. (1992). A new fast algorithm for effective training of neural classifiers. Pattern Recognition, 25:423-429.
Chow, C. (1965). Statistical independence and threshold functions. IEEE Transactions on Computers, E-14:66-68.
Chow, C. (1970). On optimum recognition error and rejection tradeoff. IEEE Transactions on Information Theory, 16:41-46.
Chow, Y. and Teicher, H. (1978). Probability Theory, Independence, Interchangeability, Martingales. Springer-Verlag, New York.
Ciampi, A. (1991). Generalized regression trees. Computational Statistics and Data Analysis, 12:57-78.
Collomb, G. (1979). Estimation de la regression par la méthode des $k$ points les plus proches: propriétés de convergence ponctuelle. Comptes Rendus de l'Académie des Sciences de Paris, 289:245-247.
Collomb, G. (1980). Estimation de la regression par la méthode des $k$ points les plus proches avec noyau. Lecture Notes in Mathematics \#821, Springer-Verlag, Berlin. 159-175.
Collomb, G. (1981). Estimation non parametrique de la regression: revue bibliographique. International Statistical Review, 49:75-93.
Conway, J. and Sloane, N. (1993). Sphere-Packings, Lattices and Groups. Springer-Verlag, Berlin.
Coomans, D. and Broeckaert, I. (1986). Potential Pattern Recognition in Chemical and Medical Decision Making. Research Studies Press, Letchworth, Hertfordshire, England.
Cormen, T., Leiserson, C., and Rivest, R. (1990). Introduction to Algorithms. MIT Press, Boston, MA.
Cover, T. (1965). Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. IEEE Transactions on Electronic Computers, 14:326-334.
Cover, T. (1968a). Estimation by the nearest neighbor rule. IEEE Transactions on Information Theory, 14:50-55.
Cover, T. (1968b). Rates of convergence for nearest neighbor procedures. In Proceedings of the Hawaii International Conference on Systems Sciences, pages 413-415. Honolulu.
Cover, T. (1969). Learning in pattern recognition. In Methodologies of Pattern Recognition, Watanabe, S., editor, pages 111-132. Academic Press, New York.
Cover, T. (1974). The best two independent measurements are not the two best. IEEE Transactions on Systems, Man, and Cybernetics, 4:116-117.
Cover, T. and Hart, P. (1967). Nearest neighbor pattern classification. IEEE Transactions on Information Theory, 13:21-27.
Cover, T. and Thomas, J. (1991). Elements of Information Theory. John Wiley, New York.
Cover, T. and Van Campenhout, J. (1977). On the possible orderings in the measurement selection problem. IEEE Transactions on Systems, Man, and Cybernetics, 7:657-661.

Cover, T. and Wagner, T. (1975). Topics in statistical pattern recognition. Communication and Cybernetics, 10:15-46.
Cramér, H. and Wold, H. (1936). Some theorems on distribution functions. Journal of the London Mathematical Society, 11:290-294.
Csibi, S. (1971). Simple and compound processes in iterative machine learning. Technical Report, CISM Summer Course, Udine, Italy.
Csibi, S. (1975). Using indicators as a base for estimating optimal decision functions. In Colloquia Mathematica Societatis János Bolyai: Topics in Information Theory, pages 143-153. Keszthely, Hungary.
Csiszár, I. (1967). Information-type measures of difference of probability distributions and indirect observations. Studia Scientiarium Mathematicarum Hungarica, 2:299-318.
Csiszár, I. (1973). Generalized entropy and quantization problems. In Transactions of the Sixth Prague Conference on Information Theory, Statistical Decision Functions, Random Processes, pages 159-174. Academia, Prague.
Csiszár, I. and Körner, J. (1981). Information Theory: Coding Theorems for Discrete Memoryless Systems. Academic Press, New York.
Cybenko, G. (1989). Approximations by superpositions of sigmoidal functions. Math. Control, Signals, Systems, 2:303-314.
Darken, C., Donahue, M., Gurvits, L., and Sontag, E. (1993). Rate of approximation results motivated by robust neural network learning. In Proceedings of the Sixth ACM Workshop on Computational Learning Theory, pages 303-309. Association for Computing Machinery, New York.
Das Gupta, S. (1964). Nonparametric classification rules. Sankhya Series A, 26:25-30.
Das Gupta, S. and Lin, H. (1980). Nearest neighbor rules of statistical classification based on ranks. Sankhya Series A, 42:419-430.
Dasarathy, B. (1991). Nearest Neighbor Pattern Classification Techniques. IEEE Computer Society Press, Los Alamitos, CA.
Day, N. and Kerridge, D. (1967). A general maximum likelihood discriminant. Biometrics, 23:313-324.
de Guzman, M. (1975). Differentiation of Integrals in $R^{n}$. Lecture Notes in Mathematics \#481, Springer-Verlag, Berlin.
Deheuvels, P. (1977). Estimation nonparamétrique de la densité par histogrammes generalisés. Publications de l'Institut de Statistique de l'Université de Paris, 22:1-23.
Devijver, P. (1978). A note on ties in voting with the $k$-nN rule. Pattern Recognition, 10:297298.

Devijver, P. (1979). New error bounds with the nearest neighbor rule. IEEE Transactions on Information Theory, 25:749-753.
Devijver, P. (1980). An overview of asymptotic properties of nearest neighbor rules. In Pattern Recognition in Practice, Gelsema, E. and Kanal, L., editors, pages 343-350. Elsevier Science Publishers, Amsterdam.
Devijver, P. and Kittler, J. (1980). On the edited nearest neighbor rule. In Proceedings of the Fifth International Conference on Pattern Recognition, pages 72-80. Pattern Recognition Society, Los Alamitos, CA.
Devijver, P. and Kittler, J. (1982). Pattern Recognition: A Statistical Approach. PrenticeHall, Englewood Cliffs, NJ.
Devroye, L. (1978). A universal $k$-nearest neighbor procedure in discrimination. In Proceedings of the 1978 IEEE Computer Society Conference on Pattern Recognition and Image Processing, pages 142-147. IEEE Computer Society, Long Beach, CA.

Devroye, L. (1981a). On the almost everywhere convergence of nonparametric regression function estimates. Annals of Statistics, 9:1310-1309.
Devroye, L. (1981b). On the asymptotic probability of error in nonparametric discrimination. Annals of Statistics, 9:1320-1327.
Devroye, L. (1981c). On the inequality of Cover and Hart in nearest neighbor discrimination. IEEE Transactions on Pattern Analysis and Machine Intelligence, 3:75-78.
Devroye, L. (1982a). Bounds for the uniform deviation of empirical measures. Journal of Multivariate Analysis, 12:72-79.
Devroye, L. (1982b). Necessary and sufficient conditions for the almost everywhere convergence of nearest neighbor regression function estimates. Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete, 61:467-481.
Devroye, L. (1983). The equivalence of weak, strong and complete convergence in $L_{1}$ for kernel density estimates. Annals of Statistics, 11:896-904.
Devroye, L. (1987). A Course in Density Estimation. Birkhäuser, Boston, MA.
Devroye, L. (1988a). Applications of the theory of records in the study of random trees. Acta Informatica, 26:123-130.
Devroye, L. (1988b). Automatic pattern recognition: A study of the probability of error. IEEE Transactions on Pattern Analysis and Machine Intelligence, 10:530-543.
Devroye, L. (1988c). The expected size of some graphs in computational geometry. Computers and Mathematics with Applications, 15:53-64.
Devroye, L. (1988d). The kernel estimate is relatively stable.Probability Theory and Related Fields, 77:521-536.
Devroye, L. (1991a). Exponential inequalities in nonparametric estimation. In Nonparametric Functional Estimation and Related Topics, Roussas, G., editor, pages 31-44. NATO ASI Series, Kluwer Academic Publishers, Dordrecht.
Devroye, L. (1991b). On the oscillation of the expected number of points on a random convex hull. Statistics and Probability Letters, 11:281-286.
Devroye, L. and Györfi, L. (1983). Distribution-free exponential bound on the $L_{1}$ error of partitioning estimates of a regression function. In Proceedings of the Fourth Pannonian Symposium on Mathematical Statistics, Konecny, F., Mogyoródi, J., and Wertz, W., editors, pages 67-76. Akadémiai Kiadó, Budapest, Hungary.
Devroye, L. and Györfi, L. (1985). Nonparametric Density Estimation: The L ${ }_{1}$ View. John Wiley, New York.
Devroye, L. and Györfi, L. (1992). No empirical probability measure can converge in the total variation sense for all distributions. Annals of Statistics, 18:1496-1499.
Devroye, L., Györfi, L., Krzyżak, A., and Lugosi, G. (1994). On the strong universal consistency of nearest neighbor regression function estimates. Annals of Statistics, 22:13711385.

Devroye, L. and Krzyżak, A. (1989). An equivalence theorem for $L_{1}$ convergence of the kernel regression estimate. Journal of Statistical Planning and Inference, 23:71-82.
Devroye, L. and Laforest, L. (1990). An analysis of random d-dimensional quadtrees. SIAM Journal on Computing, 19:821-832.
Devroye, L. and Lugosi, G. (1995). Lower bounds in pattern recognition and learning. Pattern Recognition, 28:1011-1018.
Devroye, L. and Machell, F. (1985). Data structures in kernel density estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 7:360-366.
Devroye, L. and Wagner, T. (1976a). A distribution-free performance bound in error estimation. IEEE Transactions Information Theory, 22:586-587.

Devroye, L. and Wagner, T. (1976b). Nonparametric discrimination and density estimation. Technical Report 183, Electronics Research Center, University of Texas.
Devroye, L. and Wagner, T. (1979a). Distribution-free inequalities for the deleted and holdout error estimates. IEEE Transactions on Information Theory, 25:202-207.
Devroye, L. and Wagner, T. (1979b). Distribution-free performance bounds for potential function rules. IEEE Transactions on Information Theory, 25:601-604.
Devroye, L. and Wagner, T. (1979c). Distribution-free performance bounds with the resubstitution error estimate. IEEE Transactions on Information Theory, 25:208-210.
Devroye, L. and Wagner, T. (1980a). Distribution-free consistency results in nonparametric discrimination and regression function estimation. Annals of Statistics, 8:231-239.
Devroye, L. and Wagner, T. (1980b). On the $L_{1}$ convergence of kernel estimators of regression functions with applications in discrimination. Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete, 51:15-21.
Devroye, L. and Wagner, T. (1982). Nearest neighbor methods in discrimination. In Handbook of Statistics, Krishnaiah, P. and Kanal, L., editors, volume 2, pages 193-197. North Holland, Amsterdam.
Devroye, L. and Wise, G. (1980). Consistency of a recursive nearest neighbor regression function estimate. Journal of Multivariate Analysis, 10:539-550.
Diaconis, P. and Shahshahani, M. (1984). On nonlinear functions of linear combinations. SIAM Journal on Scientific and Statistical Computing, 5:175-191.
Do-Tu, H. and Installé, M. (1975). On adaptive solution of piecewise linear approximation problem-application to modeling and identification. In Milwaukee Symposium on Automatic Computation and Control, pages 1-6. Milwaukee, WI.
Duda, R. and Hart, P. (1973). Pattern Classification and Scene Analysis. John Wiley, New York.
Dudani, S. (1976). The distance-weighted $k$-nearest-neighbor rule. IEEE Transactions on Systems, Man, and Cybernetics, 6:325-327.
Dudley, R. (1978). Central limit theorems for empirical measures. Annals of Probability, 6:899-929.
Dudley, R. (1979). Balls in $R^{k}$ do not cut all subsets of $k+2$ points. Advances in Mathematics, 31 (3):306-308.
Dudley, R. (1984). Empirical processes. In Ecole de Probabilité de St. Flour 1982. Lecture Notes in Mathematics \#1097, Springer-Verlag, New York.
Dudley, R. (1987). Universal Donsker classes and metric entropy. Annals of Probability, 15:1306-1326.
Dudley, R., Kulkarni, S., Richardson, T., and Zeitouni, O. (1994). A metric entropy bound is not sufficient for learnability. IEEE Transactions on Information Theory, 40:883-885.
Durrett, R. (1991). Probability: Theory and Examples. Wadsworth and Brooks/Cole, Pacific Grove, CA.
Dvoretzky, A. (1956). On stochastic approximation. In Proceedings of the First Berkeley Symposium on Mathematical Statistics and Probability, Neyman, J., editor, pages 39-55. University of California Press, Berkeley, Los Angeles.
Dvoretzky, A., Kiefer, J., and Wolfowitz, J. (1956). Asymptotic minimax character of a sample distribution function and of the classical multinomial estimator. Annals of Mathematical Statistics, 33:642-669.
Edelsbrunner, H. (1987). Algorithms for Computational Geometry. Springer-Verlag, Berlin.
Efron, B. (1979). Bootstrap methods: another look at the jackknife. Annals of Statistics, 7:1-26.

Efron, B. (1983). Estimating the error rate of a prediction rule: improvement on cross validation. Journal of the American Statistical Association, 78:316-331.
Efron, B. and Stein, C. (1981). The jackknife estimate of variance. Annals of Statistics, 9:586-596.
Ehrenfeucht, A., Haussler, D., Kearns, M., and Valiant, L. (1989). A general lower bound on the number of examples needed for learning. Information and Computation, 82:247261.

Elashoff, J., Elashoff, R., and Goldmann, G. (1967). On the choice of variables in classification problems with dichotomous variables. Biometrika, 54:668-670.
Fabian, V. (1971). Stochastic approximation. In Optimizing Methods in Statistics, Rustagi, J., editor, pages 439-470. Academic Press, New York, London.

Fano, R. (1952). Class Notes for Transmission of Information. Course 6.574. MIT, Cambridge, MA.
Faragó, A., Linder, T., and Lugosi, G. (1993). Fast nearest neighbor search in dissimilarity spaces. IEEE Transactions on Pattern Analysis and Machine Intelligence, 15:957-962.
Faragó, A. and Lugosi, G. (1993). Strong universal consistency of neural network classifiers. IEEE Transactions on Information Theory, 39:1146-1151.
Faragó, T. and Györfi, L. (1975). On the continuity of the error distortion function for multiple hypothesis decisions. IEEE Transactions on Information Theory, 21:458-460.
Feinholz, L. (1979). Estimation of the performance of partitioning algorithms in pattern classification. Master's Thesis, Department of Mathematics, McGill University, Montreal.

Feller, W. (1968). An Introduction to Probability Theory and its Applications, Vol.1. John Wiley, New York.
Finkel, R. and Bentley, J. (1974). Quad trees: A data structure for retrieval on composite keys. Acta Informatica, 4:1-9.
Fisher, R. (1936). The case of multiple measurements in taxonomic problems. Annals of Eugenics, 7, part II:179-188.
Fitzmaurice, G. and Hand, D. (1987). A comparison of two average conditional error rate estimators. Pattern Recognition Letters, 6:221-224.
Fix, E. and Hodges, J. (1951). Discriminatory analysis. Nonparametric discrimination: Consistency properties. Technical Report 4, Project Number 21-49-004, USAF School of Aviation Medicine, Randolph Field, TX.
Fix, E. and Hodges, J. (1952). Discriminatory analysis: small sample performance. Technical Report 21-49-004, USAF School of Aviation Medicine, Randolph Field, TX.
Fix, E. and Hodges, J. (1991a). Discriminatory analysis, nonparametric discrimination, consistency properties. In Nearest Neighbor Pattern Classification Techniques, Dasarathy, B., editor, pages 32-39. IEEE Computer Society Press, Los Alamitos, CA.

Fix, E. and Hodges, J. (1991b). Discriminatory analysis: small sample performance. In Nearest Neighbor Pattern Classification Techniques, Dasarathy, B., editor, pages 4056. IEEE Computer Society Press, Los Alamitos, CA.

Flick, T., Jones, L., Priest, R., and Herman, C. (1990). Pattern classification using protection pursuit. Pattern Recognition, 23:1367-1376.
Forney, G. (1968). Exponential error bounds for erasure, list, and decision feedback schemes. IEEE Transactions on Information Theory, 14:41-46.
Friedman, J. (1977). A recursive partitioning decision rule for nonparametric classification. IEEE Transactions on Computers, 26:404-408.

Friedman, J., Baskett, F., and Shustek, L. (1975). An algorithm for finding nearest neighbor. IEEE Transactions on Computers, 24:1000-1006.
Friedman, J., Bentley, J., and Finkel, R. (1977). An algorithm for finding best matches in logarithmic expected time. ACM Transactions on Mathematical Software, 3:209-226.
Friedman, J. and Silverman, B. (1989). Flexible parsimonious smoothing and additive modeling. Technometrics, 31:3-39.
Friedman, J. and Stuetzle, W. (1981). Projection pursuit regression. Journal of the American Statistical Association, 76:817-823.
Friedman, J., Stuetzle, W., and Schroeder, A. (1984). Projection pursuit density estimation. Journal of the American Statistical Association, 79:599-608.
Friedman, J. and Tukey, J. (1974). A projection pursuit algorithm for exploratory data analysis. IEEE Transactions on Computers, 23:881-889.
Fritz, J. and Györfi, L. (1976). On the minimization of classification error probability in statistical pattern recognition. Problems of Control and Information Theory, 5:371-382.
Fu, K., Min, P., and Li, T. (1970). Feature selection in pattern recognition. IEEE Transactions on Systems Science and Cybernetics, 6:33-39.
Fuchs, H., Abram, G., and Grant, E. (1983). Near real-time shaded display of rigid objects. Computer Graphics, 17:65-72.
Fuchs, H., Kedem, Z., and Naylor, B. (1980). On visible surface generation by a priori tree structures. In Proceedings SIGGRAPH '80, pages 124-133. Published as Computer Graphics, volume 14.
Fukunaga, K. and Flick, T. (1984). An optimal global nearest neighbor metric. IEEE Transactions on Pattern Analysis and Machine Intelligence, 6:314-318.
Fukunaga, K. and Hostetler, L. (1973). Optimization of $k$-nearest neighbor density estimates. IEEE Transactions on Information Theory, 19:320-326.
Fukunaga, K. and Hummels, D. (1987). Bayes error estimation using Parzen and $k$-NN procedures. IEEE Transactions on Pattern Analysis and Machine Intelligence, 9:634643.

Fukunaga, K. and Kessel, D. (1971). Estimation of classification error. IEEE Transactions on Computers, 20:1521-1527.
Fukunaga, K. and Kessel, D. (1973). Nonparametric Bayes error estimation using unclassified samples. IEEE Transactions Information Theory, 19:434-440.
Fukunaga, K. and Mantock, J. (1984). Nonparametric data reduction. IEEE Transactions on Pattern Analysis and Machine Intelligence, 6:115-118.
Fukunaga, K. and Narendra, P. (1975). A branch and bound algorithm for computing knearest neighbors. IEEE Transactions on Computers, 24:750-753.
Funahashi, K. (1989). On the approximate realization of continuous mappings by neural networks. Neural Networks, 2:183-192.
Gabor, D. (1961). A universal nonlinear filter, predictor, and simulator which optimizes itself. Proceedings of the Institute of Electrical Engineers, 108B:422-438.
Gabriel, K. and Sokal, R. (1969). A new statistical approach to geographic variation analysis. Systematic Zoology, 18:259-278.
Gaenssler, P. (1983). Empirical processes. Lecture Notes-Monograph Series, Institute of Mathematical Statistics, Hayward, CA.
Gaenssler, P. and Stute, W. (1979). Empirical processes: A survey of results for independent identically distributed random variables. Annals of Probability, 7:193-243.
Gallant, A. (1987). Nonlinear Statistical Models. John Wiley, New York.

Ganesalingam, S. and McLachlan, G. (1980). Error rate estimation on the basis of posterior probabilities. Pattern Recognition, 12:405-413.
Garnett, J. and Yau, S. (1977). Nonparametric estimation of the Bayes error of feature extractors using ordered nearest neighbour sets. IEEE Transactions on Computers, 26:46-54.
Gates, G. (1972). The reduced nearest neighbor rule. IEEE Transactions on Information Theory, 18:431-433.
Gelfand, S. and Delp, E. (1991). On tree structured classifiers. In Artificial Neural Networks and Statistical Pattern Recognition, Old and New Connections, Sethi, I. and Jain, A., editors, pages 71-88. Elsevier Science Publishers, Amsterdam.
Gelfand, S., Ravishankar, C., and Delp, E. (1989). An iterative growing and pruning algorithm for classification tree design. In Proceedings of the 1989 IEEE International Conference on Systems, Man, and Cybernetics, pages 818-823. IEEE Press, Piscataway, NJ.
Gelfand, S., Ravishankar, C., and Delp, E. (1991). An iterative growing and pruning algorithm for classification tree design. IEEE Transactions on Pattern Analysis and Machine Intelligence, 13:163-174.
Geman, S. and Hwang, C. (1982). Nonparametric maximum likelihood estimation by the method of sieves. Annals of Statistics, 10:401-414.
Gessaman, M. (1970). A consistent nonparametric multivariate density estimator based on statistically equivalent blocks. Annals of Mathematical Statistics, 41:1344-1346.
Gessaman, M. and Gessaman, P. (1972). A comparison of some multivariate discrimination procedures. Journal of the American Statistical Association, 67:468-472.
Geva, S. and Sitte, J. (1991). Adaptive nearest neighbor pattern classification. IEEE Transactions on Neural Networks, 2:318-322.
Giné, E. and Zinn, J. (1984). Some limit theorems for empirical processes. Annals of Probability, 12:929-989.
Glick, N. (1973). Sample-based multinomial classification. Biometrics, 29:241-256.
Glick, N. (1976). Sample-based classification procedures related to empiric distributions. IEEE Transactions on Information Theory, 22:454-461.
Glick, N. (1978). Additive estimators for probabilities of correct classification. Pattern Recognition, 10:211-222.
Goldberg, P. and Jerrum, M. (1993). Bounding the Vapnik-Chervonenkis dimension of concept classes parametrized by real numbers. In Proceedings of the Sixth ACM Workshop on Computational Learning Theory, pages 361-369. Association for Computing Machinery, New York.
Goldstein, M. (1977). A two-group classification procedure for multivariate dichotomous responses. Multivariate Behavorial Research, 12:335-346.
Goldstein, M. and Dillon, W. (1978). Discrete Discriminant Analysis. John Wiley, New York.
Golea, M. and Marchand, M. (1990). A growth algorithm for neural network decision trees. Europhysics Letters, 12:205-210.
Goodman, R. and Smyth, P. (1988). Decision tree design from a communication theory viewpoint. IEEE Transactions on Information Theory, 34:979-994.
Gordon, L. and Olshen, R. (1984). Almost surely consistent nonparametric regression from recursive partitioning schemes. Journal of Multivariate Analysis, 15:147-163.
Gordon, L. and Olshen, R. (1978). Asymptotically efficient solutions to the classification problem. Annals of Statistics, 6:515-533.

Gordon, L. and Olshen, R. (1980). Consistent nonparametric regression from recursive partitioning schemes. Journal of Multivariate Analysis, 10:611-627.
Gowda, K. and Krishna, G. (1979). The condensed nearest neighbor rule using the concept of mutual nearest neighborhood. IEEE Transactions on Information Theory, 25:488-490.
Greblicki, W. (1974). Asymptotically optimal probabilistic algorithms for pattern recognition and identification. Technical Report, Monografie No. 3, Prace Naukowe Instytutu Cybernetyki Technicznej Politechniki Wroclawsjiej No. 18, Wroclaw, Poland.
Greblicki, W. (1978a). Asymptotically optimal pattern recognition procedures with density estimates. IEEE Transactions on Information Theory, 24:250-251.
Greblicki,W.(1978b). Pattern recognition procedures with nonparametric density estimates. IEEE Transactions on Systems, Man and Cybernetics, 8:809-812.
Greblicki, W. (1981). Asymptotic efficiency of classifying procedures using the Hermite series estimate of multivariate probability densities. IEEE Transactions on Information Theory, 27:364-366.
Greblicki, W., Krzyżak, A., and Pawlak, M. (1984). Distribution-free pointwise consistency of kernel regression estimate. Annals of Statistics, 12:1570-1575.
Greblicki, W. and Pawlak, M. (1981). Classification using the Fourier series estimate of multivariate density functions. IEEE Transactions on Systems, Man and Cybernetics, 11:726-730.
Greblicki, W. and Pawlak, M. (1982). A classification procedure using the multiple Fourier series. Information Sciences, 26:115-126.
Greblicki, W. and Pawlak, M. (1983). Almost sure convergence of classification procedures using Hermite series density estimates. Pattern Recognition Letters, 2:13-17.
Greblicki, W. and Pawlak, M. (1985). Pointwise consistency of the Hermite series density estimate. Statistics and Probability Letters, 3:65-69.
Greblicki, W. and Pawlak, M. (1987). Necessary and sufficient conditions for Bayes risk consistency of a recursive kernel classification rule. IEEE Transactions on Information Theory, 33:408-412.
Grenander, U. (1981). Abstract Inference. John Wiley, New York.
Grimmett, G. and Stirzaker, D. (1992). Probability and Random Processes. Oxford University Press, Oxford.
Guo, H. and Gelfand, S. (1992). Classification trees with neural network feature extraction. IEEE Transactions on Neural Networks, 3:923-933.
Györfi, L. (1975). An upper bound of error probabilities for multihypothesis testing and its application in adaptive pattern recognition. Problems of Control and Information Theory, 5:449-457.
Györfi, L. (1978). On the rate of convergence of nearest neighbor rules. IEEE Transactions on Information Theory, 29:509-512.
Györfi, L. (1981). Recent results on nonparametric regression estimate and multiple classification. Problems of Control and Information Theory, 10:43-52.
Györfi, L. (1984). Adaptive linear procedures under general conditions. IEEE Transactions on Information Theory, 30:262-267.
Györfi, L. and Györfi, Z. (1975). On the nonparametric estimate of a posteriori probabilities of simple statistical hypotheses. In Colloquia Mathematica Societatis János Bolyai: Topics in Information Theory, pages 299-308. Keszthely, Hungary.
Györfi, L. and Györfi, Z. (1978). An upper bound on the asymptotic error probability of the $k$-nearest neighbor rule for multiple classes. IEEE Transactions on Information Theory, 24:512-514.

Györfi, L., Györfi, Z., and Vajda, I. (1978). Bayesian decision with rejection. Problems of Control and Information Theory, 8:445-452.
Györfi, L. and Vajda, I. (1980). Upper bound on the error probability of detection in nongaussian noise. Problems of Control and Information Theory, 9:215-224.
Haagerup, U. (1978). Les meilleures constantes de l'inégalite de Khintchine. Comptes Rendus de l'Académie des Sciences de Paris A, 286:259-262.
Habbema, J., Hermans, J., and van den Broek, K. (1974). A stepwise discriminant analysis program using density estimation. In COMPSTAT 1974, Bruckmann, G., editor, pages 101-110. Physica Verlag, Wien.

Hagerup, T. and Rüb, C. (1990). A guided tour of Chernoff bounds. Information Processing Letters, 33:305-308.
Hall, P. (1981). On nonparametric multivariate binary discrimination. Biometrika, 68:287294.

Hall, P. (1989). On projection pursuit regression. Annals of Statistics, 17:573-588.
Hall, P. and Wand, M. (1988). On nonparametric discrimination using density differences. Biometrika, 75:541-547.

Hand, D. (1981). Discrimination and Classification. John Wiley, Chichester, U.K.
Hand, D. (1986). Recent advances in error rate estimation. Pattern Recognition Letters, 4:335-346.
Härdle, W. and Marron, J. (1985). Optimal bandwidth selection in nonparametric regression function estimation. Annals of Statistics, 13:1465-1481.
Hart, P. (1968). The condensed nearest neighbor rule. IEEE Transactions on Information Theory, 14:515-516.
Hartigan, J. (1975). Clustering Algorithms. John Wiley, New York.
Hartmann, C., Varshney, P., Mehrotra, K., and Gerberich, C. (1982). Application of information theory to the construction of efficient decision trees. IEEE Transactions on Information Theory, 28:565-577.
Hashlamoun, W., Varshney, P., and Samarasooriya, V. (1994). A tight upper bound on the Bayesian probability of error. IEEE Transactions on Pattern Analysis and Machine Intelligence, 16:220-224.
Hastie, T. and Tibshirani, R. (1990). Generalized Additive Models. Chapman and Hall, London, U.K.
Haussler, D. (1991). Sphere packing numbers for subsets of the boolean $n$-cube with bounded Vapnik-Chervonenkis dimension. Technical Report, Computer Research Laboratory, University of California, Santa Cruz.
Haussler, D. (1992). Decision theoretic generalizations of the PAC model for neural net and other learning applications. Information and Computation, 100:78-150.
Haussler, D., Littlestone, N., and Warmuth, M. (1988). Predicting \{0, 1\} functions from randomly drawn points. In Proceedings of the 29th IEEE Symposium on the Foundations of Computer Science, pages 100-109. IEEE Computer Society Press, Los Alamitos, CA.
Hecht-Nielsen, R. (1987). Kolmogorov's mapping network existence theorem. In IEEE First International Conference on Neural Networks, volume 3, pages 11-13. IEEE, Piscataway, NJ.
Hellman, M. (1970). The nearest neighbor classification rule with a reject option. IEEE Transactions on Systems, Man and Cybernetics, 2:179-185.
Henrichon, E. and Fu, K. (1969). A nonparametric partitioning procedure for pattern classification. IEEE Transactions on Computers, 18:614-624.

Hertz, J., Krogh, A., and Palmer, R. (1991). Introduction to the Theory of Neural Computation. Addison-Wesley, Redwood City, CA.
Hills, M. (1966). Allocation rules and their error rates. Journal of the Royal Statistical Society, B28:1-31.
Hjort, N. (1986a). Contribution to the discussion of a paper by P. Diaconis and Freeman. Annals of Statistics, 14:49-55.
Hjort, N. (1986b). Notes on the theory of statistical symbol recognition. Technical Report 778, Norwegian Computing Centre, Oslo.
Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association, 58:13-30.
Holmström, L. and Klemelä, J. (1992). Asymptotic bounds for the expected $l_{1}$ error of a multivariate kernel density estimator. Journal of Multivariate Analysis, 40:245-255.
Hora, S. and Wilcox, J. (1982). Estimation of error rates in several-population discriminant analysis. Journal of Marketing Research, 19:57-61.
Horibe, Y. (1970). On zero error probability of binary decisions. IEEE Transactions on Information Theory, 16:347-348.
Horne, B. and Hush, D. (1990). On the optimality of the sigmoid perceptron. In International Joint Conference on Neural Networks, volume 1, pages 269-272. Lawrence Erlbaum Associates, Hillsdale, NJ.
Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. Neural Networks, 4:251-257.
Hornik, K. (1993). Some new results on neural network approximation. Neural Networks, 6:1069-1072.
Hornik, K., Stinchcombe, M., and White, H. (1989). Multi-layer feedforward networks are universal approximators. Neural Networks, 2:359-366.
Huber, P. (1985). Projection pursuit. Annals of Statistics, 13:435-525.
Hudimoto, H. (1957). A note on the probability of the correct classification when the distributions are not specified. Annals of the Institute of Statistical Mathematics, 9:3136.

Ito, T. (1969). Note on a class of statistical recognition functions. IEEE Transactions on Computers, C-18:76-79.
Ito, T. (1972). Approximate error bounds in pattern recognition. In Machine Intelligence, Meltzer, B. and Mitchie, D., editors, pages 369-376. Edinburgh University Press, Edinburgh.
Ivakhnenko, A. (1968). The group method of data handling-a rival of the method of stochastic approximation. Soviet Automatic Control, 1:43-55.
Ivakhnenko, A. (1971). Polynomial theory of complex systems. IEEE Transactions on Systems, Man, and Cybernetics, 1:364-378.
Ivakhnenko, A., Konovalenko, V., Tulupchuk, Y., and Tymchenko, I. (1968). The group method of data handling in pattern recognition and decision problems. Soviet Automatic Control, 1:31-41.
Ivakhnenko, A., Petrache, G., and Krasyts'kyy, M. (1968). A GMDH algorithm with random selection of pairs. Soviet Automatic Control, 5:23-30.
Jain, A., Dubes, R., and Chen, C. (1987). Bootstrap techniques for error estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 9:628-633.
Jeffreys, H. (1948). Theory of Probability. Clarendon Press, Oxford.
Johnson, D. and Preparata, F. (1978). The densest hemisphere problem. Theoretical Computer Science, 6:93-107.

Kůrková, V. (1992). Kolmogorov's theorem and multilayer neural networks. Neural Networks, 5:501-506.
Kailath, T. (1967). The divergence and Bhattacharyya distance measures in signal detection. IEEE Transactions on Communication Technology, 15:52-60.
Kanal, L. (1974). Patterns in pattern recognition.IEEE Transactions on Information Theory, 20:697-722.
Kaplan, M. (1985). The uses of spatial coherence in ray tracing. SIGGRAPH '85 Course Notes, 11:22-26.
Karlin, A. (1968). Total Positivity, Volume 1. Stanford University Press, Stanford, CA.
Karp, R. (1988).Probabilistic Analysis of Algorithms. Class Notes, University of California, Berkeley.
Karpinski, M. and Macintyre, A. (1994). Quadratic bounds for vc dimension of sigmoidal neural networks. Submitted to the ACM Symposium on Theory of Computing.
Kazmierczak, H. and Steinbuch, K. (1963). Adaptive systems in pattern recognition. IEEE Transactions on Electronic Computers, 12:822-835.
Kemp, R. (1984). Fundamentals of the Average Case Analysis of Particular Algorithms. B.G. Teubner, Stuttgart.

Kemperman, J. (1969). On the optimum rate of transmitting information. In Probability and Information Theory, pages 126-169. Springer Lecture Notes in Mathematics, SpringerVerlag, Berlin.
Kiefer, J. and Wolfowitz, J. (1952). Stochastic estimation of the maximum of a regression function. Annals of Mathematical Statistics, 23:462-466.
Kim, B. and Park, S. (1986). A fast $k$-nearest neighbor finding algorithm based on the ordered partition. IEEE Transactions on Pattern Analysis and Machine Intelligence, 8:761-766.
Kittler, J. and Devijver, P. (1981). An efficient estimator of pattern recognition system error probability. Pattern Recognition, 13:245-249.
Knoke, J. (1986). The robust estimation of classification error rates. Computers and Mathematics with Applications, 12A:253-260.
Kohonen, T. (1988). Self-Organization and Associative Memory. Springer-Verlag, Berlin.
Kohonen, T. (1990). Statistical pattern recognition revisited. In Advanced Neural Computers, Eckmiller, R., editor, pages 137-144. North-Holland, Amsterdam.
Kolmogorov, A. (1957). On the representation of continuous functions of many variables by superposition of continuous functions of one variable and addition. Doklady Akademii Nauk USSR, 114:953-956.
Kolmogorov, A. and Tikhomirov, V. (1961). $\epsilon$-entropy and $\epsilon$-capacity of sets in function spaces. Translations of the American Mathematical Society, 17:277-364.
Koutsougeras, C. and Papachristou, C. (1989). Training of a neural network for pattern classification based on an entropy measure. In Proceedings of IEEE International Conference on Neural Networks, volume 1, pages 247-254. IEEE San Diego Section, San Diego, CA.
Kraaijveld, M. and Duin, R. (1991). Generalization capabilities of minimal kernel-based methods. In International Joint Conference on Neural Networks, volume 1, pages 843848. Piscataway, NJ.

Kronmal, R. and Tarter, M. (1968). The estimation of probability densities and cumulatives by Fourier series methods. Journal of the American Statistical Association, 63:925-952.
Krzanowski, W. (1987). A comparison between two distance-based discriminant principles. Journal of Classification, 4:73-84.

Krzyżak, A. (1983). Classification procedures using multivariate variable kernel density estimate. Pattern Recognition Letters, 1:293-298.
Krzyżak, A. (1986). The rates of convergence of kernel regression estimates and classification rules. IEEE Transactions on Information Theory, 32:668-679.
Krzyżak, A. (1991). On exponential bounds on the bayes risk of the kernel classification rule. IEEE Transactions on Information Theory, 37:490-499.
Krzyżak, A., Linder, T., and Lugosi, G. (1993). Nonparametric estimation and classification using radial basis function nets and empirical risk minimization. IEEE Transactions on Neural Networks. To appear.
Krzyżak, A. and Pawlak, M. (1983). Universal consistency results for Wolverton-Wagner regression function estimate with application in discrimination. Problems of Control and Information Theory, 12:33-42.
Krzyżak, A. and Pawlak, M. (1984a). Almost everywhere convergence of recursive kernel regression function estimates. IEEE Transactions on Information Theory, 31:91-93.
Krzyżak, A. and Pawlak, M. (1984b). Distribution-free consistency of a nonparametric kernel regression estimate and classification. IEEE Transactions on Information Theory, 30:78-81.
Kulkarni, S. (1991). Problems of computational and information complexity in machine vision and learning. PhD Thesis, Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA.
Kullback, S. (1967). A lower bound for discrimination information in terms of variation. IEEE Transactions on Information Theory, 13:126-127.
Kullback, S. and Leibler, A. (1951). On information and sufficiency. Annals of Mathematical Statistics, 22:79-86.
Kushner, H. (1984). Approximation and Weak Convergence Methods for Random Processes with Applications to Stochastic Systems Theory. MIT Press, Cambridge, MA.
Lachenbruch, P. (1967). An almost unbiased method of obtaining confidence intervals for the probability of misclassification in discriminant analysis. Biometrics, 23:639-645.
Lachenbruch, P. and Mickey, M. (1968). Estimation of error rates in discriminant analysis. Technometrics, 10:1-11.
LeCam, L. (1970). On the assumptions used to prove asymptotic normality of maximum likelihood estimates. Annals of Mathematical Statistics, 41:802-828.
LeCam, L. (1973). Convergence of estimates under dimensionality restrictions. Annals of Statistics, 1:38-53.
Li, X. and Dubes, R. (1986). Tree classifier design with a permutation statistic. Pattern Recognition, 19:229-235.
Lin, Y. and Fu, K. (1983). Automatic classification of cervical cells using a binary tree classifier. Pattern Recognition, 16:69-80.
Linde, Y., Buzo, A., and Gray, R. (1980). An algorithm for vector quantizer design. IEEE Transactions on Communications, 28:84-95.
Linder, T., Lugosi, G., and Zeger, K. (1994). Rates of convergence in the source coding theorem, empirical quantizer design, and universal lossy source coding. IEEE Transactions on Information Theory, 40:1728-1740.
Lissack, T. and Fu, K. (1976). Error estimation in pattern recognition via $l^{\alpha}$ distance between posterior density functions. IEEE Transactions on Information Theory, 22:34-45.
Ljung, L., Pflug, G., and Walk, H. (1992). Stochastic Approximation and Optimization of Random Systems. Birkhäuser, Basel, Boston, Berlin.

Lloyd, S. (1982). Least squares quantization in PCM. IEEE Transactions on Information Theory, 28:129-137.
Loftsgaarden, D. and Quesenberry, C. (1965). A nonparametric estimate of a multivariate density function. Annals of Mathematical Statistics, 36:1049-1051.
Logan, B. (1975). The uncertainty principle in reconstructing functions from projections. Duke Mathematical Journal, 42:661-706.
Logan, B. and Shepp, L. (1975). Optimal reconstruction of a function from its projections. Duke Mathematical Journal, 42:645-660.
Loh, W. and Vanichsetakul, N. (1988). Tree-structured classification via generalized discriminant analysis. Journal of the American Statistical Association, 83:715-728.
Loizou, G. and Maybank, S. (1987). The nearest neighbor and the Bayes error rates. IEEE Transactions on Pattern Analysis and Machine Intelligence, 9:254-262.
Lorentz, G. (1976). The thirteenth problem of Hilbert. In Proceedings of Symposia in Pure Mathematics, volume 28, pages 419-430. Providence, RI.
Lugosi, G. (1992). Learning with an unreliable teacher. Pattern Recognition, 25:79-87.
Lugosi, G. and Nobel, A. (1996). Consistency of data-driven histogram methods for density estimation and classification. Annals of Statistics, 24:687-706.
Lugosi, G. and Pawlak, M. (1994). On the posterior-probability estimate of the error rate of nonparametric classification rules. IEEE Transactions on Information Theory, 40:475481.

Lugosi, G. and Zeger, K. (1995). Nonparametric estimation via empirical risk minimization. IEEE Transactions on Information Theory, 41:677-678.
Lugosi, G. and Zeger, K. (1996). Concept learning using complexity regularization. IEEE Transactions on Information Theory, 42:48-54.
Lukács, E. and Laha, R. (1964). Applications of Characteristic Functions in Probability Theory. Griffin, London.
Lunts, A. and Brailovsky, V. (1967). Evaluation of attributes obtained in statistical decision rules. Engineering Cybernetics, 3:98-109.
Maass, W. (1993). Bounds for the computational power and learning complexity of analog neural nets. In Proceedings of the 25th Annual ACM Symposium on the Theory of Computing, pages 335-344. Association of Computing Machinery, New York.
Maass, W. (1994). Neural nets with superlinear vc-dimension. Neural Computation, 6:875882.

Macintyre, A. and Sontag, E. (1993). Finiteness results for sigmoidal "neural" networks. In Proceedings of the 25 th Annual ACM Symposium on the Theory of Computing, pages 325-334. Association of Computing Machinery, New York.
Mack, Y. (1981). Local properties of $k$-nearest neighbor regression estimates. SIAM Journal on Algebraic and Discrete Methods, 2:311-323.
Mahalanobis, P. (1936). On the generalized distance in statistics. Proceedings of the National Institute of Sciences of India, 2:49-55.
Mahalanobis, P. (1961). A method of fractile graphical analysis. Sankhya Series A, 23:4164.

Marron, J. (1983). Optimal rates of convergence to Bayes risk in nonparametric discrimination. Annals of Statistics, 11:1142-1155.
Massart, P. (1983). Vitesse de convergence dans le theoréme de la limite centrale pour le processus empirique. PhD Thesis, Université Paris-Sud, Orsay, France.
Massart, P. (1990). The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. Annals of Probability, 18:1269-1283.

Mathai, A. and Rathie, P. (1975). Basic Concepts in Information Theory and Statistics. Wiley Eastern Ltd., New Delhi.
Matloff, N. and Pruitt, R. (1984). The asymptotic distribution of an estimator of the bayes error rate. Pattern Recognition Letters, 2:271-274.
Matula, D. and Sokal, R. (1980). Properties of Gabriel graphs relevant to geographic variation research and the clustering of points in the plane. Geographical Analysis, 12:205222.

Matushita, K.(1956). Decision rule, based on distance, for the classification problem.Annals of the Institute of Statistical Mathematics, 8:67-77.
Matushita, K. (1973). Discrimination and the affinity of distributions. In Discriminant Analysis and Applications, Cacoullos, T., editor, pages 213-223. Academic Press, New York.
Max, J. (1960). Quantizing for minimum distortion. IEEE Transactions on Information Theory, 6:7-12.
McDiarmid, C. (1989). On the method of bounded differences. In Surveys in Combinatorics 1989, pages 148-188. Cambridge University Press, Cambridge.
McLachlan, G. (1976). The bias of the apparent error rate in discriminant analysis. Biometrika, 63:239-244.
McLachlan, G. (1992). Discriminant Analysis and Statistical Pattern Recognition. John Wiley, New York.
Meisel, W. (1969). Potential functions in mathematical pattern recognition. IEEE Transactions on Computers, 18:911-918.
Meisel, W.(1972). Computer Oriented Approaches to Pattern Recognition. Academic Press, New York.
Meisel, W. (1990). Parsimony in neural networks. In International Conference on Neural Networks, pages 443-446. Lawrence Erlbaum Associates, Hillsdale, NJ.
Meisel, W. and Michalopoulos, D. (1973). A partitioning algorithm with application in pattern classification and the optimization of decision tree. IEEE Transactions on Computers, 22:93-103.
Michel-Briand, C. and Milhaud, X. (1994). Asymptotic behavior of the AD method. Technical Report, Université Montpellier 2, Montpellier.
Mielniczuk, J. and Tyrcha, J. (1993). Consistency of multilayer perceptron regression estimators. Neural Networks, To appear.
Minsky, M. (1961). Steps towards artificial intelligence. In Proceedings of the IRE, volume 49, pages 8-30.
Minsky, M. and Papert, S. (1969). Perceptrons: An Introduction to Computational Geometry. MIT Press, Cambridge, MA.
Mitchell, A. and Krzanowski, W. (1985). The Mahalanobis distance and elliptic distributions. Biometrika, 72:464-467.
Mizoguchi, R., Kizawa, M., and Shimura, M. (1977). Piecewise linear discriminant functions in pattern recognition. Systems-Computers-Controls, 8:114-121.
Moody, J. and Darken, J. (1989). Fast learning in networks of locally-tuned processing units. Neural Computation, 1:281-294.
Moore, D., Whitsitt, S., and Landgrebe, D. (1976). Variance comparisons for unbiased estimators of probability of correct classification. IEEE Transactions on Information Theory, 22:102-105.
Morgan, J. and Sonquist, J. (1963). Problems in the analysis of survey data, and a proposal. Journal of the American Statistical Association, 58:415-434.

Mui, J. and Fu, K. (1980). Automated classification of nucleated blood cells using a binary tree classifier. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2:429443.

Myles, J. and Hand, D. (1990). The multi-class metric problem in nearest neighbour discrimination rules. Pattern Recognition, 23:1291-1297.
Nadaraya, E. (1964). On estimating regression. Theory of Probability and its Applications, 9:141-142.
Nadaraya, E. (1970). Remarks on nonparametric estimates for density functions and regression curves. Theory of Probability and its Applications, 15:134-137.
Narendra, P. and Fukunaga, K. (1977). A branch and bound algorithm for feature subset selection. IEEE Transactions on Computers, 26:917-922.
Natarajan, B. (1991). Machine Learning: A Theoretical Approach. Morgan Kaufmann, San Mateo, CA.
Nevelson, M. and Khasminskii, R. (1973). Stochastic Approximation and Recursive Estimation. Translations of Mathematical Monographs, Vol. 47. American Mathematical Society, Providence, RI.
Niemann, H. and Goppert, R. (1988). An efficient branch-and-bound nearest neighbour classifier. Pattern Recognition Letters, 7:67-72.
Nilsson, N. (1965). Learning Machines: Foundations of Trainable Pattern Classifying Systems. McGraw-Hill, New York.
Nishizeki, T. and Chiba, N. (1988). Planar Graphs: Theory and Algorithms. North Holland, Amsterdam.
Nobel, A. (1994). Histogram regression estimates using data-dependent partitions. Technical Report, Beckman Institute, University of Illinois, Urbana-Champaign.
Nobel, A. (1992). On uniform laws of averages. PhD Thesis, Department of Statistics, Stanford University, Stanford, CA.
Nolan, D. and Pollard, D. (1987). U-processes: Rates of convergence. Annals of Statistics, 15:780-799.
Okamoto, M. (1958). Some inequalities relating to the partial sum of binomial probabilities. Annals of the Institute of Statistical Mathematics, 10:29-35.
Olshen, R. (1977). Comments on a paper by C.J. Stone. Annals of Statistics, 5:632-633.
Ott, J. and Kronmal, R. (1976). Some classification procedures for multivariate binary data using orthogonal functions. Journal of the American Statistical Association, 71:391399.

Papadimitriou, C. and Bentley, J. (1980). A worst-case analysis of nearest neighbor searching by projection. In Automata, Languages and Programming 1980, pages 470-482. Lecture Notes in Computer Science \#85, Springer-Verlag, Berlin.
Park, J. and Sandberg, I. (1991). Universal approximation using radial-basis-function networks. Neural Computation, 3:246-257.
Park, J. and Sandberg, I. (1993). Approximation and radial-basis-function networks. Neural Computation, 5:305-316.
Park, Y. and Sklansky, J. (1990). Automated design of linear tree classifiers. Pattern Recognition, 23:1393-1412.
Parthasarathy, K. and Bhattacharya, P. (1961). Some limit theorems in regression theory. Sankhya Series A, 23:91-102.
Parzen, E. (1962). On the estimation of a probability density function and the mode. Annals of Mathematical Statistics, 33:1065-1076.

Patrick, E. (1966). Distribution-free minimum conditional risk learning systems. Technical Report TR-EE-66-18, Purdue University, Lafayette, IN.
Patrick, E. and Fischer, F. (1970). A generalized $k$-nearest neighbor rule. Information and Control, 16:128-152.
Patrick, E. and Fisher, F. (1967). Introduction to the performance of distribution-free conditional risk learning systems. Technical Report TR-EE-67-12, Purdue University, Lafayette, IN.
Pawlak, M. (1988). On the asymptotic properties of smoothed estimators of the classification error rate. Pattern Recognition, 21:515-524.
Payne, H. and Meisel, W. (1977). An algorithm for constructing optimal binary decision trees. IEEE Transactions on Computers, 26:905-916.
Pearl, J. (1979). Capacity and error estimates for boolean classifiers with limited complexity. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1:350-355.
Penrod, C. and Wagner, T. (1977). Another look at the edited nearest neighbor rule. IEEE Transactions on Systems, Man and Cybernetics, 7:92-94.
Peterson, D. (1970). Some convergence properties of a nearest neighbor decision rule. IEEE Transactions on Information Theory, 16:26-31.
Petrov, V. (1975). Sums of Independent Random Variables. Springer-Verlag, Berlin.
Pippenger, N. (1977). Information theory and the complexity of boolean functions. Mathematical Systems Theory, 10:124-162.
Poggio, T. and Girosi, F. (1990). A theory of networks for approximation and learning. Proceedings of the IEEE, 78:1481-1497.
Pollard, D. (1981). Strong consistency of $k$-means clustering. Annals of Statistics, 9:135140.

Pollard, D. (1982). Quantization and the method of $k$-means. IEEE Transactions on Information Theory, 28:199-205.
Pollard, D. (1984). Convergence of Stochastic Processes. Springer-Verlag, New York.
Pollard, D. (1986). Rates of uniform almost sure convergence for empirical processes indexed by unbounded classes of functions. Manuscript.
Poilard, D. (1990). Empirical Processes: Theory and Applications. NSF-CBMS Regional Conference Series in Probability and Statistics, Institute of Mathematical Statistics, Hayward, CA.
Powell, M. (1987). Radial basis functions for multivariable interpolation: a review. Algorithms for Approximation. Clarendon Press, Oxford.
Preparata, F. and Shamos, M. (1985). Computational Geometry-An Introduction. SpringerVerlag, New York.

Psaltis, D., Snapp, R., and Venkatesh, S. (1994). On the finite sample performance of the nearest neighbor classifier. IEEE Transactions on Information Theory, 40:820-837.
Qing-Yun, S. and Fu, K. (1983). A method for the design of binary tree classifiers. Pattern Recognition, 16:593-603.
Quesenberry, C. and Gessaman, M. (1968). Nonparametric discrimination using tolerance regions. Annals of Mathematical Statistics, 39:664-673.
Quinlan, J. (1993). C4.5: Programs for Machine Learning. Morgan Kaufmann, San Mateo.
Rabiner, L., Levinson, S., Rosenberg, A., and Wilson, J. (1979). Speaker-independent recognition of isolated words using clustering techniques. IEEE Transactions on Acoustics, Speech, and Signal Processing, 27:339-349.

Rao, R. (1962). Relations between weak and uniform convergence of measures with applications. Annals of Mathematical Statistics, 33:659-680.
Raudys, S. (1972). On the amount of a priori information in designing the classification algorithm. Technical Cybernetics, 4:168-174.
Raudys, S. (1976). On dimensionality, learning sample size and complexity of classification algorithms. In Proceedings of the 3rd International Conference on Pattern Recognition, pages 166-169. IEEE Computer Society, Long Beach, CA.
Raudys, S. and Pikelis, V. (1980). On dimensionality, sample size, classification error, and complexity of classification algorithm in pattern recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2:242-252.
Raudys, S. and Pikelis, V. (1982). Collective selection of the best version of a pattern recognition system. Pattern Recognition Letters, 1:7-13.
Rejtő, L. and Révész, P. (1973). Density estimation and pattern classification. Problems of Control and Information Theory, 2:67-80.
Rényi, A. (1961). On measures of entropy and information. In Proceedings of the Fourth Berkeley Symposium, pages 547-561. University of California Press, Berkeley.
Révész, P. (1973). Robbins-Monroe procedures in a Hilbert space and its application in the theory of learning processes. Studia Scientiarium Mathematicarum Hungarica, 8:391398.

Ripley, B. (1993). Statistical aspects of neural networks. In Networks and Chaos-Statistical and Probabilistic Aspects, Barndorff-Nielsen, O., Jensen, J., and Kendall, W., editors, pages 40-123. Chapman and Hall, London, U.K.
Ripley, B. (1994). Neural networks and related methods for classification. Journal of the Royal Statistical Society, 56:409-456.
Rissanen, J. (1983). A universal prior for integers and estimation by minimum description length. Annals of Statistics, 11:416-431.
Ritter, G., Woodruff, H., Lowry, S., and Isenhour, T. (1975). An algorithm for a selective nearest neighbor decision rule. IEEE Transactions on Information Theory, 21:665-669.
Robbins, H. and Monro, S. (1951). A stochastic approximation method. Annals of Mathematical Statistics, 22:400-407.
Rogers, W. and Wagner, T. (1978). A finite sample distribution-free performance bound for local discrimination rules. Annals of Statistics, 6:506-514.

Rosenblatt, F. (1962). Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms. Spartan Books, Washington, DC.
Rosenblatt, M. (1956). Remarks on some nonparametric estimates of a density function. Annals of Mathematical Statistics, 27:832-837.
Rounds, E. (1980). A combined nonparametric approach to feature selection and binary decision tree design. Pattern Recognition, 12:313-317.
Royall, R. (1966). A Class of Nonparametric Estimators of a Smooth Regression Function. PhD Thesis, Stanford University, Stanford, CA.
Rumelhart, D., Hinton, G., and Williams, R. (1986). Learning internal representations by error propagation. In Parallel Distributed Processing Vol.I, Rumelhart, D., J.L.McCelland, and the PDP Research Group, editors. MIT Press, Cambridge, MA. Reprinted in: J.A. Anderson and E. Rosenfeld, Neurocomputing-Foundations of Research, MIT Press, Cambridge, MA., pp. 673-695, 1988.
Ruppert, D. (1991). Stochastic approximation. In Handbook of Sequential Analysis, Ghosh, B. and Sen, P., editors, pages 503-529. Marcel Dekker, New York.

Samet, H. (1984). The quadtree and related hierarchical data structures. Computing Surveys, 16:187-260.
Samet, H. (1990a). Applications of Spatial Data Structures. Addison-Wesley, Reading, MA.
Samet, H. (1990b). The Design and Analysis of Spatial Data Structures. Addison-Wesley, Reading, MA.
Sansone, G. (1969). Orthogonal Functions. Interscience, New York.
Sauer, N. (1972). On the density of families of sets. Journal of Combinatorial Theory Series A, 13:145-147.

Scheffé, H. (1947). A useful convergence theorem for probability distributions. Annals of Mathematical Statistics, 18:434-458.
Schläffli, L. (1950). Gesammelte Mathematische Abhandlungen. Birkhäuser-Verlag, Basel.
Schmidt, W. (1994). Neural Pattern Classifying Systems. PhD Thesis, Technical University, Delft, The Netherlands.
Schoenberg, I. (1950). On Pólya frequency functions, II: variation-diminishing integral operators of the convolution type. Acta Scientiarium Mathematicarum Szeged, 12:97106.

Schwartz, S. (1967). Estimation of probability density by an orthogonal series. Annals of Mathematical Statistics, 38:1261-1265.
Schwemer, G. and Dunn, O. (1980). Posterior probability estimators in classification simulations. Communications in Statistics-Simulation, B9:133-140.
Sebestyen, G. (1962). Decision-Making Processes in Pattern Recognition. Macmillan, New York.

Serfling, R. (1974). Probability inequalities for the sum in sampling without replacement. Annals of Statistics, 2:39-48.
Sethi, I. (1981). A fast algorithm for recognizing nearest neighbors. IEEE Transactions on Systems, Man and Cybernetics, 11:245-248.
Sethi, I. (1990). Entropy nets: from decision trees to neural nets. Proceedings of the IEEE, 78:1605-1613.
Sethi, I. (1991). Decision tree performance enhancement using an artificial neural network interpretation. In Artificial Neural Networks and Statistical Pattern Recognition, Old and New Connections, Sethi, I. and Jain, A., editors, pages 71-88. Elsevier Science Publishers, Amsterdam.
Sethi, I. and Chatterjee, B. (1977). Efficient decision tree design for discrete variable pattern recognition problems. Pattern Recognition, 9:197-206.
Sethi, I. and Sarvarayudu, G. (1982). Hierarchical classifier design using mutual information. IEEE Transactions on Pattern Analysis and Machine Intelligence, 4:441-445.

Shannon, C. (1948). A mathematical theory of communication. Bell Systems Technical Journal, 27:379-423.
Shawe-Taylor, J. (1994). Sample sizes for sigmoidal neural networks. Technical Report, Department of Computer Science, Royal Holloway, University of London, Egham, England.

Shawe-Taylor, J., Anthony, M., and Biggs, N.L. (1993). Bounding sample size with the Vapnik-Chervonenkis dimension. Discrete Applied Mathematics, 42:65-73.
Shiryayev, A. (1984). Probability. Springer-Verlag, New York.
Short, R. and Fukunaga, K. (1981). The optimal distance measure for nearest neighbor classification. IEEE Transactions on Information Theory, 27:622-627.

Simon, H. (1991). The Vapnik-Chervonenkis dimension of decision trees with bounded rank. Information Processing Letters, 39:137-141.
Simon, H. (1993). General lower bounds on the number of examples needed for learning probabilistic concepts. In Proceedings of the Sixth Annual ACM Conference on Computational Learning Theory, pages 402-412. Association for Computing Machinery, New York.
Sklansky, J. and Michelotti (1980). Locally trained piecewise linear classifiers. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2:101-111.
Sklansky, J. and Wassel, G. (1979). Pattern Classifiers and Trainable Machines. SpringerVerlag, New York.
Slud, E. (1977). Distribution inequalities for the binomial law. Annals of Probability, 5:404412.

Specht, D. (1967). Generation of polynomial discriminant functions for pattern classification. IEEE Transactions on Electronic Computers, 15:308-319.
Specht, D. (1971). Series estimation of a probability density function. Technometrics, 13:409-424.
Specht, D. (1990). Probabilistic neural networks and the polynomial Adaline as complementary techniques for classification. IEEE Transactions on Neural Networks, 1:111-121.
Spencer, J. (1987). Ten Lectures on the Probabilistic Method. SIAM, Philadelphia, PA.
Sprecher, D. (1965). On the structure of continuous functions of several variables. Transactions of the American Mathematical Society, 115:340-355.
Steele, J. (1975). Combinatorial entropy and uniform limit laws. PhD Thesis, Stanford University, Stanford, CA.
Steele, J. (1986). An Efron-Stein inequality for nonsymmetric statistics. Annals of Statistics, 14:753-758.
Stengle, G. and Yukich, J. (1989). Some new Vapnik-Chervonenkis classes. Annals of Statistics, 17:1441-1446.
Stoffel, J. (1974). A classifier design technique for discrete variable pattern recognition problems. IEEE Transactions on Computers, 23:428-441.
Stoller, D. (1954). Univariate two-population distribution-free discrimination. Journal of the American Statistical Association, 49:770-777.
Stone, C. (1977). Consistent nonparametric regression. Annals of Statistics, 5:595-645.
Stone, C. (1982). Optimal global rates of convergence for nonparametric regression. Annals of Statistics, 10:1040-1053.
Stone, C. (1985). Additive regression and other nonparametric models. Annals of Statistics, 13:689-705.
Stone, M. (1974). Cross-validatory choice and assessment of statistical predictions. Journal of the Royal Statistical Society, 36:111-147.
Stute, W. (1984). Asymptotic normality of nearest neighbor regression function estimates. Annals of Statistics, 12:917-926.
Sung, K. and Shirley, P. (1992). Ray tracing with the BSP tree. In Graphics Gems III, Kirk, D., editor, pages 271-274. Academic Press, Boston, MA.

Swonger, C. (1972). Sample set condensation for a condensed nearest neighbor decision rule for pattern recognition. In Frontiers of Pattern Recognition, Watanabe, S., editor, pages 511-519. Academic Press, New York.
Szarek, S. (1976). On the best constants in the Khintchine inequality. Studia Mathematica, 63:197-208.

Szegő, G. (1959). Orthogonal Polynomials, volume 32. American Mathematical Society, Providence, RI.
Talagrand, M. (1987). The Glivenko-Cantelli problem. Annals of Probability, 15:837-870.
Talagrand, M. (1994). Sharper bounds for Gaussian and empirical processes. Annals of Probability, 22:28-76.
Talmon, J. (1986). A multiclass nonparametric partitioning algorithm. In Pattern Recognition in Practice II, Gelsema, E. and Kanal, L., editors. Elsevier Science Publishers, Amsterdam.
Taneja, I. (1983). On characterization of J-divergence and its generalizations. Journal of Combinatorics, Information and System Sciences, 8:206-212.
Taneja, I. (1987). Statistical aspects of divergence measures. Journal of Statistical Planning and Inference, 16:137-145.
Tarter, M. and Kronmal, R. (1970). On multivariate density estimates based on orthogonal expansions. Annals of Mathematical Statistics, 41:718-722.
Tomek, I. (1976a). A generalization of the $k$-nN rule. IEEE Transactions on Systems, Man and Cybernetics, 6:121-126.
Tomek, I. (1976b). Two modifications of cnn. IEEE Transactions on Systems, Man and Cybernetics, 6:769-772.
Toussaint, G. (1971). Note on optimal selection of independent binary-valued features for pattern recognition. IEEE Transactions on Information Theory, 17:618.
Toussaint, G. (1974a). Bibliography on estimation of misclassification. IEEE Transactions on Information Theory, 20:472-479.
Toussaint, G. (1974b). On the divergence between two distributions and the probability of misclassification of several decision rules. In Proceedings of the Second International Joint Conference on Pattern Recognition, pages 27-34. Copenhagen.
Toussaint, G. and Donaldson, R. (1970). Algorithms for recognizing contour-traced handprinted characters. IEEE Transactions on Computers, 19:541-546.
Tsypkin, Y. (1971). Adaptation and Learning in Automatic Systems. Academic Press, New York.
Tutz, G. (1985). Smoothed additive estimators for non-error rates in multiple discriminant analysis. Pattern Recognition, 18:151-159.
Tutz, G. (1986). An alternative choice of smoothing for kernel-based density estimates in discrete discriminant analysis. Biometrika, 73:405-411.
Tutz, G. (1988). Smoothing for discrete kernels in discrimination. Biometrics Journal, 6:729-739.
Tutz, G. (1989). On cross-validation for discrete kernel estimates in discrimination. Communications in Statistics-Theory and Methods, 18:4145-4162.
Ullmann, J. (1974). Automatic selection of reference data for use in a nearest-neighbor method of pattern classification. IEEE Transactions on Information Theory, 20:541543.

Vajda, I. (1968). The estimation of minimal error probability for testing finite or countable number of hypotheses. Problemy Peredaci Informacii, 4:6-14.
Vajda, I. (1989). Theory of Statistical Inference and Information. Kluwer Academic Publishers, Dordrecht.
Valiant, L. (1984). A theory of the learnable. Communications of the ACM, 27:1134-1142.
Van Campenhout, J. (1980). The arbitrary relation between probability of error and measurement subset. Journal of the American Statistical Association, 75:104-109.

Van Ryzin, J. (1966). Bayes risk consistency of classification procedures using density estimation. Sankhya Series A, 28:161-170.
Vapnik, V. (1982). Estimation of Dependencies Based on Empirical Data. Springer-Verlag, New York.
Vapnik, V. and Chervonenkis, A. (1971). On the uniform convergence of relative frequencies of events to their probabilities. Theory of Probability and its Applications, 16:264-280.
Vapnik, V. and Chervonenkis, A. (1974a). Ordered risk minimization. I. Automation and Remote Control, 35:1226-1235.
Vapnik, V. and Chervonenkis, A. (1974b). Ordered risk minimization. II. Automation and Remote Control, 35:1403-1412.
Vapnik, V. and Chervonenkis, A. (1974c). Theory of Pattern Recognition. Nauka, Moscow. (in Russian); German translation: Theorie der Zeichenerkennung, Akademie Verlag, Berlin, 1979.

Vapnik, V. and Chervonenkis, A. (1981). Necessary and sufficient conditions for the uniform convergence of means to their expectations. Theory of Probability and its Applications, 26:821-832.

Vidal, E. (1986). An algorithm for finding nearest neighbors in (approximately) constant average time. Pattern Recognition Letters, 4:145-157.

Vilmansen, T. (1973). Feature evaluation with measures of probabilistic dependence. IEEE Transactions on Computers, 22:381-388.
Vitushkin, A. (1961). The absolute $\epsilon$-entropy of metric spaces. Translations of the American Mathematical Society, 17:365-367.
Wagner, T. (1971). Convergence of the nearest neighbor rule. IEEE Transactions on Information Theory, 17:566-571.
Wagner, T. (1973). Convergence of the edited nearest neighbor. IEEE Transactions on Information Theory, 19:696-699.
Wang, Q. and Suen, C. (1984). Analysis and design of decision tree based on entropy reduction and its application to large character set recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence, 6:406-417.
Warner, H., Toronto, A., Veasey, L., and Stephenson, R. (1961). A mathematical approach to medical diagnosis. Journal of the American Medical Association, 177:177-183.
Wassel, G. and Sklansky, J. (1972). Training a one-dimensional classifier to minimize the probability of error. IEEE Transactions on Systems, Man, and Cybernetics, 2:533-541.
Watson, G. (1964). Smooth regression analysis. Sankhya Series A, 26:359-372.
Weiss, S. and Kulikowski, C. (1991). Computer Systems that Learn. Morgan Kaufmann, San Mateo, CA.
Wenocur, R. and Dudley, R. (1981). Some special Vapnik-Chervonenkis classes. Discrete Mathematics, 33:313-318.
Wheeden, R. and Zygmund, A. (1977). Measure and Integral. Marcel Dekker, New York.
White, H. (1990). Connectionist nonparametric regression: multilayer feedforward networks can learn arbitrary mappings. Neural Networks, 3:535-549.
White, H. (1991). Nonparametric estimation of conditional quantiles using neural networks. In Proceedings of the 23 rd Symposium of the Interface: Computing Science and Statistics, pages 190-199. American Statistical Association, Alexandria, VA.
Widrow, B. (1959). Adaptive sampled-data systems-a statistical theory of adaptation. In IRE WESCON Convention Record, volume part 4, pages 74-85.

Widrow, B. and Hoff, M. (1960). Adaptive switching circuits. In IRE WESCON Convention Record, volume part 4, pages 96-104. Reprinted in J.A. Anderson and E. Rosenfeld, Neurocomputing: Foundations of Research, MIT Press, Cambridge, MA., 1988.
Wilson, D. (1972). Asymptotic properties of nearest neighbor rules using edited data. IEEE Transactions on Systems, Man and Cybernetics, 2:408-421.
Winder, R. (1963). Threshold logic in artificial intelligence. In Artificial Intelligence, pages 107-128. IEEE Special Publication S-142.
Wolverton, C. and Wagner, T. (1969a). Asymptotically optimal discriminant functions for pattern classification. IEEE Transactions on Systems, Science and Cybernetics, 15:258265.

Wolverton, C. and Wagner, T. (1969b). Recursive estimates of probability densities. IEEE Transactions on Systems, Science and Cybernetics, 5:307.
Wong, W. and Shen, X. (1992). Probability inequalities for likelihood ratios and convergence rates of sieve MLE's. Technical Report 346, Department of Statistics, University of Chicago, Chicago, IL.
Xu, L., Krzyżak, A., and Oja, E. (1993). Rival penalized competitive learning for clustering analysis, RBF net and curve detection. IEEE Transactions on Neural Networks, 4:636649.

Xu, L., Krzyżak, A., and Yuille, A. (1994). On radial basis function nets and kernel regression: Approximation ability, convergence rate and receptive field size. Neural Networks. To appear.
Yatracos, Y. (1985). Rates of convergence of minimum distance estimators and kolmogorov's entropy. Annals of Statistics, 13:768-774.
Yau, S. and Lin, T. (1968). On the upper bound of the probability of error of a linear pattern classifier for probabilistic pattern classes. Proceedings of the IEEE, 56:321-322.
You, K. and Fu, K. (1976). An approach to the design of a linear binary tree classifier. In Proceedings of the Symposium of Machine Processing of Remotely Sensed Data, pages 3A-10. Purdue University, Lafayette, IN.
Yukich, J. (1985). Laws of large numbers for classes of functions. Journal of Multivariate Analysis, 17:245-260.
Yunck, T. (1976). A technique to identify nearest neighbors. IEEE Transactions on Systems, Man, and Cybernetics, 6:678-683.
Zhao, L. (1987). Exponential bounds of mean error for the nearest neighbor estimates of regression functions. Journal of Multivariate Analysis, 21:168-178.
Zhao, L. (1989). Exponential bounds of mean error for the kernel estimates of regression functions. Journal of Multivariate Analysis, 29:260-273.
Zhao, L., Krishnaiah, P., and Chen, X. (1990). Almost sure $L_{r}$-norm convergence for databased histogram estimates. Theory of Probability and its Applications, 35:396-403.
Zygmund, A. (1959). Trigonometric Series I. University Press, Cambridge.

## Author Index

Abou-Jaoude, S., 143, 144, 364
Abram, G.D., 342
Aitchison, J., 474, 476
Aitken, C.G.G., 474, 476
Aizerman, M.A., 159
Akaike, H., 149, 292, 295
Alexander, K., 207
Anderson, A.C., 340
Anderson, J.A., 259
Anderson, M.W., 375
Anderson, T.W., 49, 322, 374, 375
Angluin, D., 130
Anthony, M., 493, 534
Antos, A., vi
Argentiero, P., 340, 342
Arkadjew, A.G., 149
Ash, R.B., 575
Assouad, P., 231, 232, 243
Azuma, K., 133, 135

Bahadur, R.R., 472
Bailey, T., 73
Barron, A.R., vi, 144, 208, 292, 295, 364, 494, 509, 518, 520, 527, 531, 533
Barron, R.L., vi, 509, 532, 533
Bartlett, P., 522, 544
Bashkirov, O., 159, 540

Baskett, F., 62
Baum, E.B., 521, 522
Beakley, G.W., 375
Beaudet, P., 340, 342
Beck, J., 169
Becker, P., 42, 56
Beirlant, J., vi
Ben-Bassat, M., 562
Benedek, G.M., 211, 297, 301, 485
Bennett, G., 124
Benning, R.D., 375
Bentley, J.L., 62, 333
Beran, R.H., 23
Berlinet, A., vi
Bernstein, S.N., 124
Bhattacharya, P.K., 169, 320
Bhattacharyya, A., 23
Bickel, P.J., 169
Birgé, L., 118, 243
Blumer, A., 202, 233, 236, 245
Bosq, D., vi
Brailovsky, V.L., 125
Braverman, E.M., 149, 159, 540
Breiman, L., 169, 185, 318, 336, 337, 339, $342,347,360,364,376,383,447$
Brent, R.P., 542
Broder, A.J., 62

Broeckaert, I., 149
Broniatowski, M., vi
Broomhead, D.S., 540
Buescher, K.L., 211, 212, 297, 485
Burbea, J., 28
Burman, P., vi
Burshtein, D., 340
Buzo, A., 380
Cacoullos, T., 149
Cao, R., vi
Carnal, H., 229
Casey, R.G., 340
Cencov, N.N., 279
Chang, C.L., 310
Chang, C.Y., 563
Chatterjee, B., 340
Chen, C., 125, 557, 558
Chen, H., 518
Chen, T., 518
Chen, X.R., 364
Chen, Y.C., 541, 542
Chen, Z., 128
Chernoff, H., 24, 41, 51, 56, 122, 123, 129
Chervonenkis, A.Y., v, 5, 50, 126, 187, 189, 190, 197, 198, 199, 203, 216, 235, 240, 291, 491, 492
Chin, R., 340, 342
Chou, P.A., 339
Chou, W.S., 541, 542
Chow, C.K., 18, 466
Chow, Y.S., 575
Ciampi, A., 339
Collomb, G., 169
Conway, J.H., 475
Coomans, D., 149
Cormen, T.H., 316, 326
Cover, T.M., v, vi, 5, 22, 25, 26, 62, 70, 80 , 81, 87, 114, 125, 169, 198, 219, 222, 223, 230, 292, 295, 395, 562, 564, 579
Cramér, H., 45
Csibi, S., v, 16, 506
Csiszár, I., 31, 37, 144, 219
Csűrös, M., vi, 301
Cuevas, A., vi
Cybenko, G., 518

Dasarathy, B.V., 62, 309
Das Gupta, S., 84, 375
Day, N.E., 260
de Guzmán, M., 580
Deheuvels, P., vi, 433
Della Pietra, V.D., 340
Delp, E.J., 339
Devijver, P.A., vi, 22, 30, 45, 62, 74, 77, 82, 88, 304, 309, 313, 455, 554, 562, 563
Devroye, L., 16, 33, 62, 70, 76, 77, 86, 87, $88,90,112,114,126,134,137$, $138,143,149,150,160,161,162$, $164,169,170,171,176,177,178$, 182, 183, 198, 202, 206, 208, 212, 223, 229, 236, 240, 243, 278, 304, 333, 334, 388, 389, 398, 403, 411, $414,415,416,421,430,484,588$
Diaconis, P., 538, 546
Dillon, W.R., 463
Donahue, M., 547
Donaldson, R.W., 556
Do-Tu, H., 56
Dubes, R.C., 125, 339, 557, 558
Duda, R.O., 49, 225, 270, 472, 501, 566, 572
Dudani, S.A., 88
Dudley, R.M., 193, 198, 211, 221, 223, 231, 365, 485, 496, 498
Duin, R.P.W., 540, 542
Dunn, O.J., 554
Durrett, R., 575
Dvoretzky, A., 207, 505
Edelsbrunner, H., 339, 380, 512
Efron, B., 125, 137, 556
Ehrenfeucht, A., 202, 233, 236, 245
Elashoff, J.D., 564
Elashoff, R.M., 564
Fabian, V., 505
Fano, R.M., 26
Faragó, A., vi, 62, 231, 523, 525
Faragó, T., 567
Feinholz, L., 198, 394
Feller, W., 587
Finkel, R.A., 62, 333
Fisher, F.P., 85, 184, 374
Fisher, R.A., 46

Fitzmaurice, G.M., 554
Fix, E., 61, 149, 169
Flick, T.E., 455, 538
Forney, G.D., 18
Fraiman, R., vi
Friedman, J.H., 62, 318, 336, 337, 339, $342,347,360,364,375,376,383$, 535, 538
Fritz, J., v, 56, 58
Fu, K.S., 36, 128, 340, 342, 343, 375, 562
Fuchs, H., 341, 342
Fukunaga, K., 62, 128, 307, 415, 455, 554, 555, 563
Funahashi, K., 518

Gabor, D., 532
Gabriel, K.R., 90
Gaenssler, P., 198, 207
Gallant, A.R., 527
Ganesalingam, S., 554
Garnett, J.M., 128
Gates, G.W., 307, 455
Gelfand, S.B., 339
Geman, S., 527, 540
Gerberich, C.L., 340
Gessaman, M.P., 374, 375, 393
Gessaman, P.H., 374
Geva, S., 310
Gijbels, I., vi
Giné, E., 198
Girosi, F., 540, 547
Glick, N., vi, 16, 96, 125, 225, 463, 474, 550
Goldberg, P., 524
Goldmann, G.E., 564
Goldstein, M., 463
Golea, M., 542
Gonzalez-Manteiga, W., vi
Goodman, R.M., 339
Goppert, R., 62
Gordon, L., 96, 339, 347, 364, 376
Gowda, K.C., 307
Grant, E.D., 342
Gray, R.M., 380
Greblicki, W., 149, 161, 165, 166, 279
Grenander, U., 527, 540
Grimmett, G.R., 575
Guo, H., 339
Gurvits, L., 547

Györfi, L., 16, 18, 33, 56, 57, 58, 76, 86, $138,144,150,161,162,169,170$, $171,176,208,430,505,567,588$
Györfi, Z., 18, 76, 169
Haagerup, U., 588
Habbema, J.D.F., 444
Hagerup, T., 130
Hall, P., vi, 474, 538
Hand, D.J., 149, 455, 554, 557, 558
Härdle, W., 445
Hart, P.E., v, 5, 22, 49, 62, 70, 81, 87, 225, 270, 304, 455, 472, 501, 566, 572, 579
Hartigan, J.A., 377
Hartmann, C.R.P., 340
Hashlamoun, W.A., 29, 36
Hastie, T., 356, 537, 538
Haussler, D., 202, 210, 233, 235, 236, 245, 493, 496, 498, 506, 522, 526, 527
Hecht-Nielsen, R., 534
Hellman, M.E., 81
Henrichon, E.G., 340, 342, 343, 375
Herman, C., 538
Hermans, J., 444
Hertz, J., 56, 509
Hills, M., 474
Hinton, G.E., 525
Hjort, N., 270
Hodges, J.L., 61, 149, 169
Hoeffding, W., 122, 133, 135, 589
Hoff, M.E., 54, 531
Holden, S.B., 534
Holmström, L., 150, 430
Hora, S.C., 554
Horibe, Y., 24
Horne, B., 59
Hornik, K., 518, 543
Hostetler, L.D., 455
Huber, P.J., 538
Hudimoto, H., 24
Hummels, D.M., 128
Hush, D., 59
Hwang, C.R., 527, 540
Installé, M., 56
Isenhour, T.L., 307, 455
Isogai, E., vi
Itai, A., 211, 297, 485

Ito, T., 24, 472
Ivakhnenko, A.G., 532, 533
Jain, A.K., 73, 125, 557, 558
Jeffreys, H., 27
Jerrum, M., 524
Johnson, D.S., 54
Jones, L.K., 538
Kailath, T., 24
Kanal, L.N., 125, 562
Kanevsky, D., 340
Kaplan, M.R., 342
Karlin, A.S., 434
Karp, R.M., 130
Karpinski, M., 525
Kazmierczak, H., 476
Kearns, M., 233, 236, 245
Kedem, Z.M., 342
Kégl, B., vi
Kemp, R., 469, 476
Kemperman, J.H.B., 37
Kerridge, N.E., 260
Kessel, D.L., 128, 415, 554, 555
Khasminskii, R.Z., 505
Kiefer, J., 207, 505
Kim, B.S., 62
Kittler, J., 22, 30, 45, 77, 88, 304, 309, 313, 455, 554, 562, 563
Kizawa, M., 342
Klemelä, J., 150, 430
Knoke, J.D., 550
Kohonen, T., 310
Kolmogorov, A.N., 22, 479, 481, 482, 486, 487, 510, 534, 536, 537
Konovalenko, V.V., 532
Körner, J., 219
Koutsougeras, C., 542
Kraaijveld, M.A., 540, 542
Krasits'kyy, M.S., 532
Krishna, G., 307
Krishnaiah, P.R., 364
Krogh, A., 56, 509
Kronmal, R.A., 279, 473
Krzanowski, W.J., 31, 463
Krzyżak, A., vi, 149, 150, 161, 163, 164, 165, 166, 169, 170, 176, 447, 540, 541, 547
Kulikowski, C.A., 509

Kulkarni, S.R., 211, 485
Kullback, S., 27, 37
Kumar, P.R., 211, 212, 297
Kůrková, V., 537
Kushner, H.J., 505
Lachenbruch, P.A., 415
Laforest, L., 334
Laha, R.G., 45
Landgrebe, D.A., 554
LeCam, L., 23, 33
Leibler, A., 27
Leiserson, C.E., 316, 326
Levinson, S.E., 85
Li, T.L., 562
Li, X., 339
Lin, H.E., 84
Lin, T.T., 45
Lin, Y.K., 343
Linde, Y., 380
Linder, T., vi, 62, 506, 540, 541, 547
Lissack, T., 36
Littlestone, N., 210, 233, 235
Liu, R., 518
Ljung, L., 505
Lloyd, S.P., 380
Loftsgaarden, D.O., 184
Logan, B.F., 539
Loh, W.Y., 342, 343
Loizou, G., 84
Lorentz, G.G., 510, 534
Lowe, D., 540
Lowry, S.L., 307, 455
Lugosi, G., 62, 89, 109, 169, 170, 176, 210, 236, 240, 243, 292, 301, 364, 365, $369,378,383,506,523,525,527$, $528,540,541,547,551,554,555$, 559
Lukács, E., 45
Lunts, A.L., 125
Maass, W., 523, 524
Machell, F., 160
Macintyre, A., 524, 525
Mack, Y.P., vi, 169, 185
Mahalanobis, P.C., 30, 320
Mantock, J.M., 307
Marchand, M., 542
Marron, J.S., 114, 445

Massart, P., 44, 198, 207
Mathai, A.M., 22
Matloff, N., 554
Matula, D.W., 90
Matushita, K., 23, 31
Max, J., 380
Maybank, S.J., 84
McDiarmid, C., 3, 133, 134, 136
McLachlan, G.J., 49, 125, 251, 259, 270, $399,415,468,472,554,557,558$
Mehrotra, K.G., 340
Meisel, W., 149, 185, 339, 359, 365, 447, 542, 562
Michalopoulos, W.S., 359, 365
Michel-Briand, C., 339
Michelotti, 342
Mickey, P.A., 415
Mielniczuk, J., 527
Milhaud, X., 339
Min, P.J., 562
Minsky, M.L., 466, 509
Mitchell, A.F.S., 31
Mizoguchi, R., 342
Monro, S., 505
Moody, J., 540
Moore, D.S., 554
Morgan, J.N., 339
Muchnik, I.E., 159, 540
Mui, J.K., 343
Myles, J.P., 455

Nadaraya, E.A., 149
Nádas, A., vi, 340
Nagy, G., 340
Narendra, P.M., 62, 563
Natarajan, B.K., 202, 468
Naylor, B., 342
Nevelson, M.B., 505
Niemann, H., 62
Nilsson, N.J., 44, 509, 525
Nobel, A.B., vi, 364, 365, 369, 378, 383, 499
Nolan, D., 498, 499

Oja, E., 540
Okamoto, M., 51, 75, 122, 131
Olshen, R.A., 96, 178, 318, 336, 337, 339, $342,347,360,364,376,383$
Ott, J., 472

Páli, I., vi
Palmer, R.G., 56, 509
Papachristou, C.A., 542
Papadimitriou, C.H., 62
Papert, S., 509
Park, J., 547
Park, S.B., 62
Park, Y., 342
Parthasarathy, K.R., 320
Parzen, E., 149, 161
Patrick, E.A., 85, 184, 374
Pawlak, M., vi, 149, 161, 165, 166, 279, 551, 554, 555, 559
Payne, H.J., 339
Pearl, J., 470
Penrod, C.S., 309, 313
Peterson, D.W., 87
Petrache, G., 532
Petrov, V.V., 416
Petry, H., vi
Pflug, G., vi, 505
Pikelis, V., 49, 567
Pinsker, M., 37
Pintér, M., vi
Pippinger, N., 470
Poggio, T., 540, 547
Pollard, D., 193, 198, 365, 492, 493, 496, $498,499,500,506$
Powell, M.J.D., 540
Preparata, F.P., 54, 177
Priest, R.G., 538
Pruitt, R., 554
Psaltis, D., 87
Purcell, E., 185, 447
Pyatniskii, E.S., 159

Qing-Yun, S., 340, 342
Quesenberry, C.P., 184, 374
Quinlan, J.R., 339

Rabiner, L.R., 85
Rao, C.R., 28
Rao, R.R., 228
Rathie, P.N., 22
Raudys, S., 49, 567
Ravishankar, C.S., 339
Rejtő, L., 149
Rényi, A., 28
Révész, P., v, 149, 161

Richardson, T., 211, 485
Ripley, B.D., 509
Rissanen, J., 292, 295
Ritter, G.L., 307, 455
Rivest, R.L., 316, 326
Robbins, H., 505
Rogers, W.H., 4, 411, 413
Rosenberg, A.E., 85
Rosenblatt, F., 5, 14, 39, 44, 509
Rosenblatt, M., 149, 161
Rounds, E.M., 339
Roussas, G., vi
Royall, R.M., 71, 73, 179, 185, 453
Rozonoer, L.I., 159
Rüb, C., 130
Rumelhart, D.E., 525
Ruppert, D., 505
Samarasooriya, V.N.S., 29, 36
Samet, H., 326, 333, 342
Sandberg, I.W., 547
Sansone, G., 281
Sarvarayudu, G.P.R., 339
Sauer, N., 216
Scheffé, H., 32, 208, 211, 383
Schläffli, L., 395
Schmidt, W.F., 525
Schoenberg, I.J., 434
Schroeder, A., 538
Schwartz, S.C., 279
Schwemer, G.T., 554
Sebestyen, G., 149, 340, 432, 562
Serfling, R.J., 589
Sethi, I.K., 62, 339, 340, 542
Shahshahani, M., 538, 546
Shamos, M.I., 177
Shannon, C.E., 25
Shawe-Taylor, J., 493, 525
Shen, X., 254, 527
Shepp, L., 539
Shimura, M., 342
Shirley, P., 342
Shiryayev, A.N., 575
Short, R.D., 455
Shustek, L.J., 62
Silverman, B.W., 536, 538
Simar, L., vi
Simon, H.U., 198, 240, 477
Sitte, J., 310

Sklansky, J., 56, 342
Sloane, N.J.A., 475
Slud, E.V., 87, 588
Smyth, P., 339
Snapp, R.R., 87
Sokal, R.R., 90
Sonquist, J.A., 339
Sontag, E., 524, 547
Specht, D.F., 160, 279, 531, 532
Spencer, J., 496
Sprecher, D.A., 534
Steele, J.M., 137, 221, 229
Stein, C., 137
Steinbuch, K., 476
Stengle, G., 198
Stephenson, R., 468
Stinchcombe, M., 518
Stirzaker, D.R., 575
Stoffel, J.C., 340
Stoller, D.S., 40, 43
Stone, C.J., v, 3, 6, 65, 70, 97, 98, 100, 101, $114,169,175,179,181,183,318$, $336,337,339,342,347,360,364$, 376, 383, 538
Stone, M., 125
Stuetzle, W., 538
Stute, W., vi, 169, 207
Suen, C.Y., 339
Sung, K., 342
Swonger, C.W., 307
Szabados, T., vi
Szabó, R., 301
Szarek, S.J., 588
Szegő, G., 289
Talagrand, M., 199, 207
Talmon, J.L., 339
Taneja, I.J., 28
Tarter, M.E., 279
Thomas, J.A., 25, 26, 219
Tibshirani, R.J., 536, 537, 538
Tikhomirov, V.M., 479, 481, 482, 486, 487
Tomek, I., 307, 309, 455
Toronto, A.F., 468
Toussaint, G.T., vi, 28, 35, 125, 556, 564, 565
Tsypkin, Y.Z., 505
Tukey, J., 538
Tulupchuk, Y.M., 532

Tuteur, F.B., 375
Tutz, G.E., 441, 474, 477, 550
Tymchenko, I.K., 532
Tyrcha, J., 527

Ullmann, J.R., 307, 455
Vajda, Igor, vi, 22, 32
Vajda, István, 18, 56, 57
Valiant, L.G., 130, 202, 233, 236, 245
Van Campenhout, J.M., 562, 564
van den Broek, K., 444
van der Meulen, E.C., vi, 144, 208
Vanichsetakul, N., 342, 343
Van Ryzin, J., 16, 149
Vapnik, V.N., v, 5, 50, 126, 187, 189, 190, 197, 198, 199, 202, 203, 206, 207, 211, 216, 235, 240, 291, 485, 491, 492, 493, 501
Varshney, P.K., 29, 36, 340
Veasey, L.G., 468
Venkatesh, S.S., 87
Vidal, E., 62
Vilmansen, T.E., 563
Vitushkin, A.G., 482

Wagner, T.J., v, vi, 4, 16, 20, 62, 84, 125, 149, 161, 198, 202, 206, 304, 309, $313,389,398,403,411,413,414$, 415, 416, 421, 455
Walk, H., vi, 505
Wand, M.P., 474
Wang, Q.R., 339
Warmuth, M.K., 202, 210, 233, 235, 236, 245
Warner, H.R., 468

Wassel, G.N., 56
Watson, G.S., 149
Weiss, S.M., 509
Wenocur, R.S., 498
Wheeden, R.L., 580
White, H., 518, 527
Whitsitt, S.J., 554
Widrow, B., 54, 531
Wilcox, J.B., 554
Williams, R.J., 525
Wilson, D.L., 309, 313, 455
Wilson, J.G., 85
Winder, R.O., 466
Wise, G.L., 177, 182
Wold, H., 45
Wolfowitz, J., 207, 505
Wolverton, C.T., 16, 20, 161
Wong, W.H., 254, 527
Woodruff, H.B., 307, 455
Wouters, W., v
Xu, L., 540

Yakowitz, S., vi
Yatracos, Y.G., vi, 278
Yau, S.S., 45, 128
You, K.C., 340
Yuille, A.L., 540
Yukich, J.E., 198
Yunck, T.P., 62

Zeitouni, O., 211, 485
Zhao, L.C., 150, 169, 170, 364
Zeger, K., vi, 292, 301, 506, 527, 528
Zinn, J., 198
Zygmund, A., 281, 575, 580

## Subject Index

absolute error, 11
accuracy, 201, 202, 206, 246
Adaline, 531
additive model, 536, 538, 546
admissible transformation, 570, 571, 572, 573
agreement, 163
AID criterion, 339
Alexander's inequality, 207, 210, 211, 278, 295, 300
ancestral rule, 180, 181, 185, 420
a posteriori probability, $10,11,16,92,97$, $102,108,113,301,479,483,493$, 549, 551, 552, 554, 555, 563, 568, 570, 572
apparent error rate, see error estimation, resubstitution
approximation error, $188,284,290,295$, 297, 395, 430, 431, 454, 487, 517, 523, 529, 547
arrangement, $511,512,513,514,515,516$, 517, 542
arrangement classifier, 511,512,514,515, 516, 517
association inequality, 22, 583
asymptotic optimality, $390,393,431,445$
back propagation, 525
Bahadur-Lazarsfeld expansion, 472
bandwidth, see smoothing factor
Barron network, 545, 546
basis function, $225,266,279,285,290$, 398, 501, 506
complete orthonormal, 280, 281, 282, 287
Haar, 281, 288
Hermite, 279, 281, 288
Laguerre, 281, 288
Legendre, 281, 288
Rademacher, 282
Rademacher-Walsh, 470, 472
trigonometric, 279, 287, 288
Walsh, 282
Bayes classifier (or rule), 2, 10, 11, 12, 13, $14,15,16,17,19,20,47,48,57$, $58,59,61,69,202,226,239,249$, 264, 265, 266, 271, 272, 274, 275, 276, 277, 280, 294, 295, 300, 301, 307, 357, 387, 464, 466, 467, 471, $472,473,474,487,527,569,570$
Bayes decision, see Bayes classifier
Bayes error, 2, 11, 12, 13, 14, 15, 17, 18, $19,21,23,28,29,30,31,35,36$, $39,57,79,82,90,91,111,112$,
$113,128,129,176,188,236,237$, 241, 276, 285, 289, 295, 297, 298, 307, 308, 313, 357, 389, 401, 451, $533,545,561,562,563,565,567$, $568,569,570,571,572,573$
estimation of, 4, 128,566
Bayes problem, 2, 9
Bayes risk, see Bayes error
Bennett's inequality, 124, 130
Beppo-Levy theorem, 577
Bernstein perceptron, 544
Bernstein polynomial, 544
Bernstein's inequality, 124, 130, 210, 494
beta distribution, 77, 88, 266, 276
Bhattacharyya affinity, 23, 24, 343
bias, 121, 125, 126, 144, 151, 395, 555
binary space partition tree, $316,317,318$, $338,340,341,342,343,344,347$, 362, 383, 542
balanced, 362
raw, 342, 362
binomial distribution, $51,71,72,75,76$, $77,78,79,81,82,83,84,85,86$, $89,94,106,109,122,124,127$, $129,130,131,152,184,230,238$, $242,276,306,309,321,332,345$, $410,444,449,462,586,587,588$
binomial theorem, 218
boolean classification, 461, 468
Borel-Cantelli lemma, 142, 193, 227, 284, $293,357,367,381,425,452,505$, 585
bracketing metric entropy, 253, 254, 262
branch-and-bound method, 563
BSP tree, see binary space partition tree
CART, 336, 338, 339
Catalan number, 469, 476
Cauchy-Schwarz inequality, 20, 33, 52, 76, $93,95,99,103,104,141,155,158$, $173,256,345,401,402,414,463$, 582, 583
central limit theorem, 85, 270
channel, 109
characteristic function, 45
Chebyshev-Cantelli inequality, 42, 237, 463, 582
Chebyshev's inequality, 45, 56, 97, 194, 210, 269, 276, 322, 332, 414,

494, 582
Chernoff's affinity, 24, 35
Chernoff's bounding method, 117, 123, $129,130,131,135,230,261$
$\chi^{2}$-divergence, 34,36
class-conditional density, $14,15,17,19$, $31,148,149,150,166,184,260$, 263, 264, 266, 268, 269, 271, 272, 276, 277, 280, 282, 285
class-conditional distribution, 56, 118, 263, $264,265,271,274,275,276$
classifier selection, 125, 127, 132, 187, 188, 189, 193, 199, 201, 233, 234, $245,250,404,549,554,555,559$
lower bounds for, 206, 233, 234, 239, 245, 246, 275
class probability, $15,17,18,19,32,34,58$, 150
clustering, $343,363,377,378,379,380$, 384, 392, 506, 542
committee machine, 525, 526, 543
complexity penalty, 291, 292, 297, 300, 301, 395
complexity regularization, 266, 289, 290, 291, 292, 293, 294, 295, 296, 298, 300, 301, 395, 524, 525, 527, 531
concept learning, see learning
confidence, 201, 202, 246
consistency, 92, 93, 104, 106, 109, 128, $132,184,188,232,268,296,389$, 390, 468, 513, 555
definition of, 3, 91
of Fourier series rule, 282, 284
of generalized linear rules, 286
of kernel rules, 160, 161, 163, 171, 439, 441, 442, 446, 448, 474
of maximum likelihood, 249, 250, 252, $253,254,257,259,262,268,473$
of nearest neighbor rules, $80,174,176$, $307,308,309,310,311,312,313$, 457
of neural network classifiers, 514, 523, 525, 527, 531, 534, 542, 544, 545
of partitioning rules, $94,96,319,320$, $322,323,324,332,333,335,338$, $339,340,341,342,343,344,346$, $348,349,350,357,358,359,361$, $362,364,368,370,374,376,377$, $381,384,385,392,394,395$
of skeleton estimates, 483,487
of squared error minimization, 490, 493
strong, see strong consistency
strong universal, see strong universal consistency
universal, see universal consistency
weak, $91,100,103,133,139,149,162$, $169,184,425,430,452$
consistent rule, see consistency
convex hull, 229, 232
Cover-Hart inequality, $5,22,24,27,35,62$, 70, 448
covering
of a class of classifiers, 211, 212
of a class of functions, 479, 480, 483, $484,485,486,487,489,499,500$
empirical, 211, 298, 485
simple empirical, 289, 296, 297, 298, 300, 301
covering lemma, $67,153,157,158,164$, $165,171,176,181,182$
covering number, 479, 483, 492, 493, 494, $496,499,504,526,530,547$
cross-validation, 125, 445
curse of dimensionality, $469,486,561$
decision with rejection, 18
denseness
in $L_{1}, 286,287,517,520,542$
in $L_{2}, 502,503$
in $L_{p}, 543,579$
with respect to the supremum norm, 287, 517, 518, 533, 539
density estimation, $107,118,143,149$, $150,151,161,162,169,184,208$, $266,278,279,364,433,445,447$, 484, 527
density of a class of sets, 231
diamond, 231
dimension, 17, 49, 52, 222, 363, 482, 567
reduction of, 561, 567
distance
data-based, 169, 391, 451, 456, 459
empirical, 178, 182
Euclidean, 63, 101, 149, 169, 368, 370, 378, 456
generalized, 182
$L_{p}, 182$
rotation-invariant, 179,180
tie breaking, $4,61,69,101,170,174$, $175,176,177,178,180,181,182$, $304,305,377,413,414,451,456$, 458, 553
distribution
beta, $77,88,266,276$
binomial, 51, 71, 72, 75, 76, 77, 78, 79, $81,82,83,84,85,86,89,94,106$, $109,122,124,127,129,130,131$, 152, 184, 230, 238, 242, 276, 306, $309,321,332,345,410,444,449$, 462, 586, 587, 588
class-conditional, 56, 118, 263, 264, $265,271,274,275,276$
exponential, 361, 590
gamma, 266, 268, 276, 573, 590
geometric, 80, 153, 161, 167
hypergeometric, 589
Maxwell, 266
multinomial, 227, 321, 589
normal, 31, 35, 47, 48, 49, 56, 57, 59, 76, 85, 86, 249, 264, 265, 266, 276, $343,390,399,415,564,566,567$, 572, 588, 590
Rayleigh, 266, 277
strictly separable, 104, 162
distribution function, 142
class-conditional, 40, 41, 45
empirical, 142, 193, 325, 375
divergence, see Kullback-Leibler divergence
dominated convergence theorem, 63,156 , 267, 329, 514, 577
Dvoretzky-Kiefer-Wolfowitz-Massart inequality, 44, 57, 207
elementary cut, 350
embedding, 64
empirical classifier selection, see classifier selection
empirical covering, see covering
empirical error, 43, 49, 54, 56, 125, 187, $188,191,200,310,335,351,362$, 482
empirically optimal classifier, 187, 188, $200,206,225,275$
empirical measure, 43, 142, 190, 193, 203, 208, 210, 211, 213, 265, 270, 271,
$272,278,304,365,366,367,398$, 414, 443, 553, 556
empirical risk, see empirical error empirical risk minimization, $5,6,49,50$, $53,54,56,125,126,131,187,191$, 192, 193, 199, 201, 202, 209, 210, $211,212,225,227,232,234,235$, $239,246,249,251,257,258,265$, $266,268,274,275,277,285,286$, $287,289,290,291,292,294,306$, $310,311,334,339,346,347,354$, $375,388,404,420,428,441,444$, 454, 455, 457, 458, 467, 469, 476, $477,479,483,487,489,491,493$, $501,502,506,512,514,521,523$, $525,526,527,534,537,539,540$, $541,542,546,549,551$
empirical squared Euclidean distance error, $377,378,380,384$
entropy, 25, 26, 29, 35, 218, 219, 252, 286, 337, 339, 373
$\epsilon$-covering number, see covering number
$\epsilon$-entropy, see metric entropy
$\epsilon_{m}$-optimality, $390,393,454$
$\epsilon$-packing number, see packing number
error estimation, $4,121,124,125,127$, $132,343,397,408,409,549$
0.632 estimator, 557
bootstrap, $125,549,556,557,558$
complexity penalized, 291, 297, 396
cross validated, see leave-one-out estimate
deleted, see leave-one-out estimate
double bootstrap, 557
E0 estimator, 557
error counting, $121,123,126,131,549$, 550
holdout, $125,347,388,389,391,393$, $395,397,428,430,453,549,556$
leave-one-out, $4,125,309,397,407$, $408,410,413,414,415,417,419$, 420, 441, 444, 446, 458, 459, 474, $549,550,554,555,556,557$
variance of, $415,419,552,556$
posterior probability, $549,554,555,559$
randomized bootstrap, 557
resubstitution, $125,304,343,347,382$, $397,398,399,400,402,405,407$, $415,419,439,458,474,476,549$,

550, 557
bias of, $343,397,398,399,400,405$, 419, 550
rotation, 125,556
smoothed, $125,549,550,551,552,554$, 555
$U$-method, see leave-one-out estimate
estimation error, $188,284,290,295,395$, $431,454,455,471,487,523,529$
Euclidean norm, 99, 370, 567
Euler's theorem, 380
exponential distribution, 361,590
exponential family, $266,275,276$

Fano's inequality, 26, 35
Fatou's lemma, 577
$f$-divergence, $31,32,34,36$
feature extraction, $14,21,128,561,562$, 567, 572, 573
$F$-error, 29, 32, 34, 36, 573
fingering, $191,192,277$
fingering dimension, 192, 209
Fisher linear discriminant, $47,49,58$
flexible grid, 367, 368, 369
Fourier series classifier, 279, 280, 282, 285, 471
Fourier series density estimation, 279
Fubini's theorem, 139, 157, 371, 579, 582, 583
fundamental rule, $440,461,462,463,471$, 474
fundamental theorem of mathematical statistics, see Glivenko-Cantelli theorem

Gabriel graph, 90
Gabriel neighbor, 90
gamma distribution, $266,268,276,573$, 590
gaussian distribution, see normal distribution
gaussian noise, 57, 285
geometric distribution, $80,153,161,167$
Gessaman's rule, $375,376,384,393$
ghost sample, 193, 203, 210, 213
Gini index, 338, 339, 340, 343, 361, 362
Glivenko-Cantelli theorem, 192, 193, 321, 374
gradient optimization, $54,56,251,525$

Grenander's method of sieves, 527, 540
grid complexity, 357
group method of data handling, 532
Hamming distance, 474, 477
ham-sandwich theorem, 339
Hart's rule, 304, 305, 306
Hellinger distance, 23, 32, 243, 474
Hellinger integral, 32, 36
histogram density estimation, 107, 143, 208
histogram rule, 94, 95, 96, 98, 106, 109, $133,138,139,143,144,145,147$, $152,153,155,159,170,253,261$, 287, 288, 290, 301, 318, 342, 347, $351,363,368,370,378,384,394$, 399, 400, 401, 404, 405, 408, 409, 417, 418, 419, 420, 462, 463, 506, $518,549,552,553,555,559$
cubic, $95,96,107,133,138,144,381$, 382, 384, 465, 475, 512
data-dependent, $317,365,368,369,373$, 382, 385, 391, 403, 404
lazy, 143, 152, 408, 409, 414
randomized, 475
Hoeffding's inequality, $51,52,75,122$, $124,127,133,134,135,137,190$, 193, 198, 212, 213, 284, 299, 306, $311,366,428,440,445,448,483$, 487, 490, 495
Hölder's inequality, 108, 582
hypergeometric distribution, 589
imperfect training, 89, 109
impurity function, $336,337,338,340,343$, 361, 362
inequality
Alexander's, 207, 210, 211, 278, 295, 300
association, 22, 583
Bennett's, 124, 130
Bernstein's, 124, 130, 210, 494
Cauchy-Schwarz, 20, 33, 52, 76, 93, 95 , $99,103,104,141,155,158,173$, $256,345,401,402,414,463,582$, 583
Chebyshev-Cantelli, 42, 237, 463, 582
Chebyshev's, 45, 56, 97, 194, 210, 269, $276,322,332,414,494,582$

Cover-Hart, 5, 22, 24, 27, 35, 62, 70, 448
Dvoretzky-Kiefer-Wolfowitz-Massart, 44, 57, 207
Fano's, 26, 35
Hoeffding's, 51, 52, 75, 122, 124, 127, 133, 134, 135, 137, 190, 193, 198, 212, 213, 284, 299, 306, 311, 366, $428,440,445,448,483,487,490$, 495
Hölder's, 108, 582
Jensen's, 23, 26, 29, 32, 76, 99, 141, 242, 252, 288, 401, 402, 418, 463, 558, 568, 583
Khinchine's, 410, 588
large deviation, 136, 190, 364, 365, 379
LeCam's, 33, 244
Markov's, 76, 117, 123, 256, 309, 379, 582
McDiarmid's, 3, 136, 142, 158, 162, 172, 174, 201, 211, 346, 414, 448, 553, 558
Pinsker's, 37
Rogers-Wagner, 4, 411, 413, 417, 418
Vapnik-Chervonenkis, 5, 138, 192, 196, 197, 198, 201, 202, 203, 206, 225, 226, 229, 233, 271, 273, 275, 278, $295,304,312,356,357,364,365$, 366, 388, 399, 429, 436, 467, 476, 492, 494
information divergence, see KullbackLeibler divergence

Jeffreys' divergence, 27, 28, 29
Jensen's inequality, 23, 26, 29, 32, 76, 99, 141, 242, 252, 288, 401, 402, 418, 463, 558, 568, 583
Jessen-Marcinkiewitz-Zygmund theorem, 329, 580, 581

Karhunen-Loève expansion, 568
k-d tree, 321, 326, 327, 328, 333, 334, 358, 359, 375
chronological, 327, 328, 343, 344, 358
deep, 327, 333, 358
permutation-optimized chronological, 344
well-populated, 327
kernel, 147, 148, 149, 151, 153, 159, 160, $161,162,163,165,166,185,416$,

420, 421, 423, 428, 430, 432, 438, $439,440,441,446,447,448,474$, $476,518,540,541,542,547,549$
Cauchy, 148, 149, 165
Deheuvels', 433
de la Vallée-Poussin, 446
devilish, 162
Epanechnikov, 148, 149, 162
exponential, 433
gaussian, 105, 109, 148, 149, 165, 416, $421,426,433,448,531$
Hermite, 151
multiparameter, 435, 447
naïve, $148,149,150,153,156,157,158$, $163,167,416,420,426,442,445$
negative valued, 151,153
polynomial, 434
product, 435, 447
regular, $149,150,156,158,164,165$, $166,185,415,424,425,426,433$, 446
star-shaped, 432
uniform, see naïve kernel
window, see naïve kernel
kernel complexity, 429, 430, 431, 432, 433, $434,435,436,446,447$
kernel density estimation, $149,150,151$, 161, 162, 208, 433, 447
kernel rule, $98,105,109,147,148,149$, $150,151,153,159,160,161,162$, $163,166,167,171,185,285,290$, $390,408,409,415,416,417,423$, $428,429,430,439,440,442,444$, 445, 446, 447, 462, 474, 476, 477, 531, 532, 540, 547, 552
automatic, $150,423,424,425,426$
variable, 447, 448
Khinchine's inequality, 410, 588
$k$-local rule, 64, 69, 414
$k$-means clustering, $380,381,384,542$
Kolmogorov-Lorentz network, 535
Kolmogorov-Lorentz representation, 511, 534, 535, 538, 546
Kolmogorov-Smirnov distance, 41, 271, 272, 274
generalized, 271, 272, 273, 274, 277, 278
Kolmogorov-Smirnov statistic, 142, 207
Kolmogorov variational distance, 22, 36
$k$-spacing rule, $320,363,373,374,376$, 420
Kullback-Leibler divergence, 27, 34, 36, 252, 258, 261
$\mathbf{L}_{1}$ consistency, 109, 143, 161, 162, 184
$L_{1}$ convergence, 93,270
$L_{1}$ distance, $15,16,492$
$L_{1}$ error, 20, 107, 108, 150, 162, 208, 525, 526, 527, 528, 542
$L_{2}$ consistency, 102, 284, 559
$L_{2}$ error, 104, 283, 284, 445, 526, 542
large deviation inequality, 136, 190, 364, 365, 379
learning, 2, 201, 202, 211, 233
algorithm, 201
supervised, 2
with a teacher, 2
learning vector quantization, 310
Lebesgue's density theorem, 580
LeCam's inequality, 33, 244
likelihood product, 250, 258, 259, 260, 268
linear classifier, $14,19,39,40,44,47,48$, $49,50,51,53,54,55,57,58,59$, $191,224,225,260,311,342,343$, $390,399,466,467,468,476,490$, 501, 507, 510, 540, 567
generalized, $39,160,224,225,266,279$, 285, 287, 289, 290, 398, 399, 501, 505, 506, 532, 549
linear discriminant, see linear classifier
linear independence, 222
linear ordering by inclusion, 231
Lloyd-Max algorithm, 380
local average estimator, 98
logistic discrimination, 259
log-linear model, 472
$L_{p}$ error, 525, 526, 527
lying teacher, see imperfect training
Mahalanobis distance, 30, 36, 47, 48, 57, 566, 572
majority vote, $61,84,86,94,96,98,101$, 106, 109, 147, 152, 166, 170, 176, $177,178,183,317,319,322,327$, $328,333,334,339,341,342,343$, $346,347,349,356,361,363,368$, $376,377,383,384,392,420,435$,
$463,474,475,512,513,514,516$, 525
Markov's inequality, 76, 117, 123, 256, 309, 379, 582
martingale, 134, 135, 142
martingale difference sequence, 133,134 , 135,136
Matushita error, 23, 29, 34, 36, 109
maximum likelihood, 249, 250, 251, 252, 253, 254, 257, 259, 260, 261, 262, 265, 266, 268, 269, 270, 276, 277, 444, 467, 472, 473, 475, 527, 540, 551
distribution format, 250, 260
regression format, 249, 250, 260
set format, 249
Maxwell distribution, 266
McDiarmid's inequality, $3,136,142,158$, $162,172,174,201,211,346,414$, 448, 553, 558
median tree, $322,323,324,325,334,358$, 362, 375
theoretical, 323, 324, 358
method of bounded differences, 133, 135
metric, see distance
metric entropy, $480,481,484,486,492$, 526
minimax lower bound, 234, 235, 236
minimum description length principle, 292, 295
minimum distance estimation, 265,266 , $270,271,272,274,277,278$
moment generating function, 56,589
monomial, 468, 469, 476, 531, 533, 534
monotone layer, 183, 226, 227, 229, 232, 487
moving window rule, $147,148,150,552$, 553, 555, 559
multinomial discrimination, 461
multinomial distribution, 227, 321, 589
multivariate normal distribution, see normal distribution
natural classifier, $317,318,320,323,327$, $342,347,358,513,514,542$
nearest neighbor, $62,63,69,80,84,85,89$, $90,101,166,169,170,171,175$, $177,179,180,181,182,184,185$,

307, 310, 392, 413, 414, 441, 448, 453, 456, 553
nearest neighbor clustering, 377,379
nearest neighbor density estimation, 184
nearest neighbor error, $22,23,29,31,57$, $63,69,132,307,308,361,551,559$
nearest neighbor rule
1-NN, 5, 70, 80, 84, 87, 88, 106, 107, 152, 169, 180, 181, 185, 285, 303, 304, 305, 306, 307, 308, 309, 310, $311,313,359,360,387,390,391$, 395, 398, 405, 407, 414, 454, 458, 475, 550
$3-\mathrm{NN}, 14,80,84,86,405,565,566$
admissibility of the 1 -NN rule, $5,81,87$, 174
asymptotic error probability of the $1-\mathrm{NN}$ rule, see nearest neighbor error
asymptotic error probability of the $k$-NN rule, $70,71,73,75,79,82,169$
automatic, 169, 387, 451
based on reference data, see prototype nN rule
condensed, 303, 304, 307, 309, 455
data-based, see automatic nearest neighbor rule
edited, $62,307,309,313,455$
$k$-NN, 3, 4, 5, 6, 61, 62, 64, 69, 70, 73, 74, $80,81,82,83,85,86,87,98,100$, $101,169,170,175,176,177,179$, 180, 181, 182, 185, 290, 309, 313, $343,389,397,408,413,414,415$, $420,421,451,452,453,455,456$, 457, 458, 459, 549, 552, 553, 555
( $k, l$ )-NN, 81, 82, 88
layered, 183
prototype, 307, 309, 310, 311, 377, 455
recursive, 177, 179, 182
relabeling, $169,180,181,185,309,420$
selective, 455
variable metric, 455
weighted, $61,71,73,74,85,87,88,90$, $179,183,451,453$
neural network, $444,507,510,514,520$, 527, 528, 530, 531
classifier, 5, 39, 290, 516, 517, 543, 549
multilayer, 509, 523, 545
with one hidden layer, $507,508,509$, $510,517,518,519,520,521,525$,

## 538, 543, 544

with two hidden layers, 509, 510, 514, 515, 516, 517, 522, 537, 542, 546
Neyman-Pearson lemma, 18
normal distribution, $31,35,47,48,49,56$, 57, 59, 76, 85, 86, 249, 264, 265, $266,276,343,390,399,415,564$, 566, 567, 572, 588, 590
order statistics, $88,178,257,320,328,358$, 363, 372, 382
outer layer, see monotone layer
overfitting, 188, 286, 291, 297, 348, 395, 454, 527, 532, 544
packing number, 481, 496
Padaline, 531, 532, 545
parameter estimation, 250, 268, 270, 277, 399
parametric classification, 24, 250, 263, $270,275,390,415,472$
Parseval's identity, 283
partition, $94,106,109,140,143,144,147$, $152,287,301,307,310,312,315$, $317,318,319,320,321,327,328$, $333,343,346,347,349,350,351$, $352,353,354,356,357,360,363$, $364,365,366,368,372,373,374$, $378,381,382,383,392,394,395$, $399,400,402,403,404,417,418$, 419, 420, 454, 477, 506, 511, 513, 553
complexity of a family of, $365,367,369$, 403, 404
cubic, $95,96,107,133,138,143,382$, 475
data-dependent, $94,363,364,370,404$, 420
$\epsilon$-effective cardinality of, 144
family of, $365,367,368,369,403,404$, 405
$\mu$-approximating, 144
partitioning rule, see histogram rule
perceptron, $6,14,39,40,44,507,509,510$, $514,515,516,532,544$
perceptron criterion, 59
pigeonhole principle, 238,476
Pinsker's inequality, 37
planar graph, 312,380
plug-in decision, $15,16,17,19,20,24,93$, 265, 266, 267, 270, 271, 276, 277, 467, 468, 471, 472, 475
polynomial discriminant function, 160 , 279, 287, 531
polynomial network, 532, 533, 534
potential function rule, $159,160,279,540$
principal component analysis, 455
probabilistic method, 112, 234, 496
projection pursuit, 538, 539, 546
pseudo dimension, 498, 506
pseudo probability of error, 82
quadratic discrimination rule, 225, 265, $276,343,390,468,476$
quadtree, $321,333,334,348,361$
chronological, 333, 361
deep, 333, 361
quantization, 380, 461, 463, 567
radial basis function, 540, 541, 542, 547
Radon-Nikodym derivative, 272, 313
Radon-Nikodym theorem, 578,586
rate of convergence, $4,102,104,108,111$, $113,114,118,128,239,262,264$, 267, 269, 275, 276, 284, 294, 295, 296, 297, 300, 357, 390, 483, 485, 486, 494, 527
Rayleigh distribution, 266, 277
record, 330, 331, 358
recursive kernel rule, 160, 161, 166
regression function, $9,11,93,254,261$, $445,454,474,475,479,483,484$, 487, 489
estimation, $92,93,102,103,104,108$, $114,149,169,172,185,405,445$, 527
relative stability, 162
Rogers-Wagner inequality, 4, 411, 413, 417, 418
rotation invariance, see transformation invariance
Royall's rule, 185
sample complexity, 201, 202, 205, 206, 211, 234, 245
sample scatter, 46, 47
sampling without replacement, 213
scale invariance, see transformation invariance
scaling, $144,166,177,179,381,385,387$, 391
scatter matrix, 46
Scheffe's theorem, 32, 36, 208, 211, 278, 383
search tree, 62,326
balanced, 322
shatter coefficient, 196, 197, 198, 200, 212, $215,216,219,220,222,223,226$, 271, 273, 312, 347, 365, 389, 390, $404,429,453,492,493,496,497$, 512, 522, 541
shattering, 196, 198, 199, 217, 218, 219, 220, 221, 222, 223, 226, 229, 235, 239, 277, 436, 437, 438, 476, 477, 498, 499, 521, 541
sieve method, 444
sigmoid, 56, 59, 507, 510, 515, 517, 519, $520,521,524,526,528,542,543$, 544, 545
arctan, 508, 524
gaussian, 508, 524
logistic, see standard sigmoid
piecewise polynomial, 524
standard, 508, 524, 525
threshold, 507, 508, 516, 517, 519, $521,522,523,524,525,526,537$, 543, 544
signed measure, 272, 273
skeleton estimate, 292, 480, 482, 483, 484, 486, 487, 489, 493
smart rule, 106, 109
smoothing factor, $148,149,153,160,162$, $163,165,166,185,390,416,423$, 424, 428, 429, 435, 436, 439, 441, $442,444,446,462,474,531$
data-dependent, $423,424,425,426,474$
spacing, $328,358,372$
splitting criterion, $334,335,336,337,338$, 339, 348
splitting function, 334
splitting the data, $132,185,211,298,309$, $347,387,388,389,390,391,392$, 397, 428, 436, 452, 453, 454, 541
squared error, $11,54,55,56,58,59,445$, $448,489,490,491,492,493,501$, $502,505,506,525,526,527,551$,

559
standard empirical measure, see empirical measure
statistically equivalent blocks, 178, 363, 372, 374, 375, 384, 393
Stirling's formula, 83, 238, 415, 449, 587
stochastic approximation, 161, 505
stochastic process
expansion of, 573
sampling of, 573
Stoller split, 40, 334, 335, 336, 338, 350, 361, 375
empirical, 43, 339, 343
theoretical, 40, 41, 42, 43, 56
Stoller's rule, 43, 44
Stone's lemma, 65, 66, 69, 83
Stone's theorem, 97, 100, 101, 107, 181, 182, 183, 185, 456
Stone-Weierstrass theorem, 281, 287, 519, 538, 543, 580
strictly separable distribution, 104, 162
strong consistency, 92, 93, 128
definition of, 3, 91
of Fourier series rule, 282, 284
of kernel rules, 149, 161, 162, 431
of nearest neighbor rules, $100,169,174$, 451
of neural network classifiers, 528
of partitioning rules, $143,155,341,347$, $368,369,370,376,378,383$
strong universal consistency, 395
of complexity regularization, 290, 292, 298
definition of, 92
of generalized linear rules, 287,502, 505
of kernel rules, $149,150,166,424,426$, $429,430,435$
of nearest neighbor rules, 170, 176, 453, 458, 459
of neural network classifiers, 523,528, 545
of partitioning rules, $96,133,138,142$, $143,144,170,288,369,376,381$, 382, 384, 462
structural risk minimization, see complexity regularization
sufficient statistics, $268,270,571,572$, 573, 574
superclassifier, 5, 118
support, $63,80,84,579$
symmetric rule, $109,411,413,414,419$
symmetrization, $193,194,203,278$
tennis rule, 82,88
testing sequence, $121,124,125,126,127$, $131,388,390,392,393,397,428$, 453, 455, 457
total boundedness, $480,483,486,487,493$, 494
total positivity, 434
total variation, 22, 32, 36, 208
transformation invariance, $6,62,84,177$, $178,179,182,183,285,318,319$, $320,328,334,337,348,373,375$, 385, 512, 531
tree
$\alpha$-balanced, 358
binary, $315,316,364$
BSP, see binary space partition tree
classifier, $6,39,316,319,322,341,343$, $347,364,383,390,394,477,516$, 542
depth of, $315,333,358$
expression, 469
greedy, $348,349,350$
height of, $315,334,474$
Horton-Strahler number of, 477
hyperplane, see binary space partition tree
k-d, see k-d tree
median, see median tree
ordered, 316
ordinary, $316,317,318$
perpendicular splitting, $334,336,348$
pruning of, see trimming of a tree
sphere, 316,317
trimming of, $327,336,338,339$
unbiased estimator, 121, 285
uniform deviations, see uniform laws of large numbers
uniform laws of large numbers, 190,196 , $198,199,207,228,366,490,529$
uniform marginal transformation, 319
universal consistency, $97,98,106,109$, $113,118,226,285,387,510$
of complexity regularization, 293
definition of, 3,92
of generalized linear rules, 286, 501, 502
of kernel rules, $149,153,163,164,165$, $166,167,423,425,441,442$
of nearest neighbor rules, $61,100,169$, $170,175,177,178,180,181,182$, $183,184,185,453,454,455$
of neural network classifiers, 512,524 , $526,528,532,539,540,547$
of partitioning rules, $95,96,107,133$, $253,344,346,364,382,393$

Vapnik-Chervonenkis dimension, see vC dimension
Vapnik-Chervonenkis inequality, 5, 138, $192,196,197,198,201,202,203$, $206,225,226,229,233,271,273$, $275,278,295,304,312,356,357$, $364,365,366,388,399,429,436$, $467,476,492,494$
variation, 144
vc class, 197, 210, 251, 290, 297
vc dimension, 5, 192, 196, 197, 200, 201, $202,206,210,215,216,218,219$, 220, 221, 223, 224, 226, 229, 230, $231,232,233,235,236,238,239$, $240,243,245,247,257,272,273$, $275,277,278,286,289,290,291$, $292,294,295,297,298,300,301$, $346,347,398,436,439,444,447$, $467,471,476,477,486,496,521$, $522,523,524,526,530,533,534$, $537,538,541,542,544,546,547$
Voronoi cell, $62,304,312,378$
Voronoi partition, 62, 304, 312, 377, 378, 380
weighted average estimator, 98
X-property, $319,320,322,333,334,339$, $342,343,357,358,362,513$

Pattern recognition presents one of the most significant challenges for scientists and engineers, and many different approaches have been proposed. The aim of this book is to provide a self-contained account of probabilistic analysis of these approaches. The book includes a discussion of distance measures, nonparametric methods based on kernels or nearest neighbors, Vapnik-Chervonenkis theory, epsilon entropy, parametric classification, error estimation, tree classifiers, and neural networks.

Wherever possible, distribution-free properties and inequalities are derived. A substantial portion of the results or the analysis is new. Over 430 problems and exercises complement the material.


