

Cambridge University Press
0521813360 - Data Analysis and Graphics Using R: An Example-based Approach
John Maindonald and John Braun
Frontmatter
[More information](#)

Data Analysis and Graphics
Using R – an Example-based Approach

CAMBRIDGE SERIES IN STATISTICAL AND PROBABILISTIC
MATHEMATICS

Editorial Board:

- R. Gill, *Department of Mathematics, Utrecht University*
B.D. Ripley, *Department of Statistics, University of Oxford*
S. Ross, *Department of Industrial Engineering, University of California, Berkeley*
M. Stein, *Department of Statistics, University of Chicago*
D. Williams, *School of Mathematical Sciences, University of Bath*

This series of high-quality upper-division textbooks and expository monographs covers all aspects of stochastic applicable mathematics. The topics range from pure and applied statistics to probability theory, operations research, optimization, and mathematical programming. The books contain clear presentations of new developments in the field and also of the state of the art in classical methods. While emphasizing rigorous treatment of theoretical methods, the books also contain applications and discussions of new techniques made possible by advances in computational practice.

Already published

1. *Bootstrap Methods and Their Application*, A.C. Davison and D.V. Hinkley
2. *Markov Chains*, J. Norris
3. *Asymptotic Statistics*, A.W. van der Vaart
4. *Wavelet Methods for Time Series Analysis*, D.B. Percival and A.T. Walden
5. *Bayesian Methods*, T. Leonard and J.S.J. Mu
6. *Empirical Processes in M-Estimation*, S. van de Geer
7. *Numerical Methods of Statistics*, J. Monahan
8. *A User's Guide to Measure-Theoretic Probability*, D. Pollard
9. *The Estimation and Tracking of Frequency*, B.G. Quinn and E.J. Hannan

Cambridge University Press
0521813360 - Data Analysis and Graphics Using R: An Example-based Approach
John Maindonald and John Braun
Frontmatter
[More information](#)

Data Analysis and Graphics Using R – an Example-based Approach

John Maindonald

Centre for Bioinformation Science, John Curtin School of Medical Research
and Mathematical Sciences Institute, Australian National University

and

John Braun

Department of Statistical and Actuarial Science, University of Western Ontario



CAMBRIDGE
UNIVERSITY PRESS

Cambridge University Press
 0521813360 - Data Analysis and Graphics Using R: An Example-based Approach
 John Maindonald and John Braun
 Frontmatter
[More information](#)

CAMBRIDGE UNIVERSITY PRESS
 Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
 40 West 20th Street, New York, NY 10011-4211, USA
www.cambridge.org
 Information on this title: www.cambridge.org/9780521813365

© Cambridge University Press 2003

This book is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2003

Reprinted 2004, 2005

Printed in the United States of America

A catalogue record for this publication is available from the British Library.

Library of Congress Cataloguing in Publication data

Maindonald, J. H. (John Hilary), 1937–
 Data analysis and graphics using R : an example-based approach / John Maindonald and John Braun.
 p. cm. – (Cambridge series in statistical and probabilistic mathematics)
 Includes bibliographical references and index.
 ISBN 0 521 81336 0
 1. Statistical – Data processing. 2. Statistics – Graphic methods – Data processing. 3. R (Computer program language) I. Braun, John, 1963– II. Title. III. Cambridge series on statistical and probabilistic mathematics.
 QA276.4.M245 2003
 519.5'0285–dc21 2002031560

ISBN-13 978-0-521-81336-5 hardback
 ISBN-10 0-521-81336-0 hardback

Cambridge University Press has no responsibility for the persistence or accuracy of URLs for external or third-party Internet Web sites referred to in this book and does not guarantee that any content on such Web sites is, or will remain, accurate or appropriate.

It is easy to lie with statistics. It is hard to tell the truth without statistics.

[Andrejs Dunkels]

... technology tends to overwhelm common sense.

[D. A. Freedman]

Contents

Preface	<i>page xv</i>
A Chapter by Chapter Summary	xix
1 A Brief Introduction to R	1
1.1 <i>A Short R Session</i>	1
1.1.1 R must be installed!	1
1.1.2 Using the console (or command line) window	1
1.1.3 Reading data from a file	2
1.1.4 Entry of data at the command line	3
1.1.5 Online help	4
1.1.6 Quitting R	5
1.2 <i>The Uses of R</i>	5
1.3 <i>The R Language</i>	6
1.3.1 R objects	7
1.3.2 Retaining objects between sessions	7
1.4 <i>Vectors in R</i>	8
1.4.1 Concatenation – joining vector objects	8
1.4.2 Subsets of vectors	8
1.4.3 Patterned data	9
1.4.4 Missing values	9
1.4.5 Factors	10
1.5 <i>Data Frames</i>	11
1.5.1 Variable names	12
1.5.2 Applying a function to the columns of a data frame	13
1.5.3* Data frames and matrices	13
1.5.4 Identification of rows that include missing values	13
1.6 <i>R Packages</i>	14
1.6.1 Data sets that accompany R packages	14
1.7* <i>Looping</i>	14
1.8 <i>R Graphics</i>	15
1.8.1 The function <code>plot()</code> and allied functions	16
1.8.2 Identification and location on the figure region	19
1.8.3 Plotting mathematical symbols	20

1.8.4	Row by column layouts of plots	20
1.8.5	Graphs – additional notes	22
1.9	<i>Additional Points on the Use of R in This Book</i>	23
1.10	<i>Further Reading</i>	25
1.11	<i>Exercises</i>	26
2	Styles of Data Analysis	29
2.1	<i>Revealing Views of the Data</i>	29
2.1.1	Views of a single sample	30
2.1.2	Patterns in grouped data	33
2.1.3	Patterns in bivariate data – the scatterplot	34
2.1.4*	Multiple variables and times	36
2.1.5	Lattice (trellis style) graphics	37
2.1.6	What to look for in plots	41
2.2	<i>Data Summary</i>	42
2.2.1	Mean and median	43
2.2.2	Standard deviation and inter-quartile range	44
2.2.3	Correlation	46
2.3	<i>Statistical Analysis Strategies</i>	47
2.3.1	Helpful and unhelpful questions	48
2.3.2	Planning the formal analysis	48
2.3.3	Changes to the intended plan of analysis	49
2.4	<i>Recap</i>	49
2.5	<i>Further Reading</i>	50
2.6	<i>Exercises</i>	50
3	Statistical Models	52
3.1	<i>Regularities</i>	53
3.1.1	Mathematical models	53
3.1.2	Models that include a random component	54
3.1.3	Smooth and rough	55
3.1.4	The construction and use of models	56
3.1.5	Model formulae	56
3.2	<i>Distributions: Models for the Random Component</i>	57
3.2.1	Discrete distributions	57
3.2.2	Continuous distributions	58
3.3	<i>The Uses of Random Numbers</i>	60
3.3.1	Simulation	60
3.3.2	Sampling from populations	61
3.4	<i>Model Assumptions</i>	62
3.4.1	Random sampling assumptions – independence	62
3.4.2	Checks for normality	63
3.4.3	Checking other model assumptions	66
3.4.4	Are non-parametric methods the answer?	66
3.4.5	Why models matter – adding across contingency tables	67

3.5	<i>Recap</i>	68
3.6	<i>Further Reading</i>	68
3.7	<i>Exercises</i>	69
4	An Introduction to Formal Inference	71
4.1	<i>Standard Errors</i>	71
4.1.1	Population parameters and sample statistics	71
4.1.2	Assessing accuracy – the standard error	71
4.1.3	Standard errors for differences of means	72
4.1.4*	The standard error of the median	73
4.1.5*	Resampling to estimate standard errors: bootstrapping	73
4.2	<i>Calculations Involving Standard Errors: the t-Distribution</i>	74
4.3	<i>Confidence Intervals and Hypothesis Tests</i>	77
4.3.1	One- and two-sample intervals and tests for means	78
4.3.2	Confidence intervals and tests for proportions	82
4.3.3	Confidence intervals for the correlation	83
4.4	<i>Contingency Tables</i>	84
4.4.1	Rare and endangered plant species	86
4.4.2	Additional notes	87
4.5	<i>One-Way Unstructured Comparisons</i>	88
4.5.1	Displaying means for the one-way layout	90
4.5.2	Multiple comparisons	91
4.5.3	Data with a two-way structure	92
4.5.4	Presentation issues	92
4.6	<i>Response Curves</i>	93
4.7	<i>Data with a Nested Variation Structure</i>	94
4.7.1	Degrees of freedom considerations	95
4.7.2	General multi-way analysis of variance designs	96
4.8*	<i>Resampling Methods for Tests and Confidence Intervals</i>	96
4.8.1	The one-sample permutation test	96
4.8.2	The two-sample permutation test	97
4.8.3	Bootstrap estimates of confidence intervals	98
4.9	<i>Further Comments on Formal Inference</i>	100
4.9.1	Confidence intervals versus hypothesis tests	100
4.9.2	If there is strong prior information, use it!	101
4.10	<i>Recap</i>	102
4.11	<i>Further Reading</i>	102
4.12	<i>Exercises</i>	103
5	Regression with a Single Predictor	107
5.1	<i>Fitting a Line to Data</i>	107
5.1.1	Lawn roller example	108
5.1.2	Calculating fitted values and residuals	109
5.1.3	Residual plots	110
5.1.4	The analysis of variance table	113
5.2	<i>Outliers, Influence and Robust Regression</i>	114

5.3	<i>Standard Errors and Confidence Intervals</i>	116
5.3.1	Confidence intervals and tests for the slope	116
5.3.2	SEs and confidence intervals for predicted values	117
5.3.3*	Implications for design	118
5.4	<i>Regression versus Qualitative ANOVA Comparisons</i>	119
5.5	<i>Assessing Predictive Accuracy</i>	121
5.5.1	Training/test sets, and cross-validation	121
5.5.2	Cross-validation – an example	121
5.5.3*	Bootstrapping	123
5.6*	<i>A Note on Power Transformations</i>	126
5.7	<i>Size and Shape Data</i>	127
5.7.1	Allometric growth	128
5.7.2	There are two regression lines!	129
5.8	<i>The Model Matrix in Regression</i>	130
5.9	<i>Recap</i>	131
5.10	<i>Methodological References</i>	132
5.11	<i>Exercises</i>	132
6	Multiple Linear Regression	134
6.1	<i>Basic Ideas: Book Weight and Brain Weight Examples</i>	134
6.1.1	Omission of the intercept term	137
6.1.2	Diagnostic plots	138
6.1.3	Further investigation of influential points	139
6.1.4	Example: brain weight	140
6.2	<i>Multiple Regression Assumptions and Diagnostics</i>	142
6.2.1	Influential outliers and Cook's distance	143
6.2.2	Component plus residual plots	143
6.2.3*	Further types of diagnostic plot	145
6.2.4	Robust and resistant methods	145
6.3	<i>A Strategy for Fitting Multiple Regression Models</i>	145
6.3.1	Preliminaries	145
6.3.2	Model fitting	146
6.3.3	An example – the Scottish hill race data	147
6.4	<i>Measures for the Comparison of Regression Models</i>	152
6.4.1	R^2 and adjusted R^2	152
6.4.2	AIC and related statistics	152
6.4.3	How accurately does the equation predict?	153
6.4.4	An external assessment of predictive accuracy	155
6.5	<i>Interpreting Regression Coefficients – the Labor Training Data</i>	155
6.6	<i>Problems with Many Explanatory Variables</i>	161
6.6.1	Variable selection issues	162
6.6.2	Principal components summaries	163
6.7	<i>Multicollinearity</i>	164
6.7.1	A contrived example	164
6.7.2	The variance inflation factor (VIF)	167
6.7.3	Remedying multicollinearity	168

6.8	<i>Multiple Regression Models – Additional Points</i>	168
6.8.1	Confusion between explanatory and dependent variables	168
6.8.2	Missing explanatory variables	169
6.8.3*	The use of transformations	169
6.8.4*	Non-linear methods – an alternative to transformation?	170
6.9	<i>Further Reading</i>	171
6.10	<i>Exercises</i>	172
7	Exploiting the Linear Model Framework	175
7.1	<i>Levels of a Factor – Using Indicator Variables</i>	175
7.1.1	Example – sugar weight	175
7.1.2	Different choices for the model matrix when there are factors	178
7.2	<i>Polynomial Regression</i>	179
7.2.1	Issues in the choice of model	181
7.3	<i>Fitting Multiple Lines</i>	183
7.4*	<i>Methods for Passing Smooth Curves through Data</i>	187
7.4.1	Scatterplot smoothing – regression splines	188
7.4.2	Other smoothing methods	191
7.4.3	Generalized additive models	191
7.5	<i>Smoothing Terms in Multiple Linear Models</i>	192
7.6	<i>Further Reading</i>	194
7.7	<i>Exercises</i>	194
8	Logistic Regression and Other Generalized Linear Models	197
8.1	<i>Generalized Linear Models</i>	197
8.1.1	Transformation of the expected value on the left	197
8.1.2	Noise terms need not be normal	198
8.1.3	Log odds in contingency tables	198
8.1.4	Logistic regression with a continuous explanatory variable	199
8.2	<i>Logistic Multiple Regression</i>	202
8.2.1	A plot of contributions of explanatory variables	208
8.2.2	Cross-validation estimates of predictive accuracy	209
8.3	<i>Logistic Models for Categorical Data – an Example</i>	210
8.4	<i>Poisson and Quasi-Poisson Regression</i>	211
8.4.1	Data on aberrant crypt foci	211
8.4.2	Moth habitat example	213
8.4.3*	Residuals, and estimating the dispersion	215
8.5	<i>Ordinal Regression Models</i>	216
8.5.1	Exploratory analysis	217
8.5.2*	Proportional odds logistic regression	217
8.6	<i>Other Related Models</i>	220
8.6.1*	Loglinear models	220
8.6.2	Survival analysis	220
8.7	<i>Transformations for Count Data</i>	221
8.8	<i>Further Reading</i>	222
8.9	<i>Exercises</i>	223

9	Multi-level Models, Time Series and Repeated Measures	224
9.1	<i>Introduction</i>	224
9.2	<i>Example – Survey Data, with Clustering</i>	225
9.2.1	Alternative models	225
9.2.2	Instructive, though faulty, analyses	228
9.2.3	Predictive accuracy	229
9.3	<i>A Multi-level Experimental Design</i>	230
9.3.1	The ANOVA table	232
9.3.2	Expected values of mean squares	232
9.3.3*	The sums of squares breakdown	234
9.3.4	The variance components	236
9.3.5	The mixed model analysis	237
9.3.6	Predictive accuracy	239
9.3.7	Different sources of variance – complication or focus of interest?	239
9.4	<i>Within and between Subject Effects – an Example</i>	239
9.5	<i>Time Series – Some Basic Ideas</i>	242
9.5.1	Preliminary graphical explorations	242
9.5.2	The autocorrelation function	243
9.5.3	Autoregressive (AR) models	244
9.5.4*	Autoregressive moving average (ARMA) models – theory	245
9.6*	<i>Regression Modeling with Moving Average Errors – an Example</i>	246
9.7	<i>Repeated Measures in Time – Notes on the Methodology</i>	252
9.7.1	The theory of repeated measures modeling	253
9.7.2	Correlation structure	253
9.7.3	Different approaches to repeated measures analysis	254
9.8	<i>Further Notes on Multi-level Modeling</i>	255
9.8.1	An historical perspective on multi-level models	255
9.8.2	Meta-analysis	256
9.9	<i>Further Reading</i>	256
9.10	<i>Exercises</i>	257
10	Tree-based Classification and Regression	259
10.1	<i>The Uses of Tree-based Methods</i>	259
10.1.1	Problems for which tree-based regression may be used	259
10.1.2	Tree-based regression versus parametric approaches	260
10.1.3	Summary of pluses and minuses	261
10.2	<i>Detecting Email Spam – an Example</i>	261
10.2.1	Choosing the number of splits	264
10.3	<i>Terminology and Methodology</i>	264
10.3.1	Choosing the split – regression trees	265
10.3.2	Within and between sums of squares	266
10.3.3	Choosing the split – classification trees	267
10.3.4	The mechanics of tree-based regression – a trivial example	268

10.4	<i>Assessments of Predictive Accuracy</i>	270
10.4.1	Cross-validation	270
10.4.2	The training/test set methodology	271
10.4.3	Predicting the future	271
10.5	<i>A Strategy for Choosing the Optimal Tree</i>	271
10.5.1	Cost–complexity pruning	271
10.5.2	Prediction error versus tree size	272
10.6	<i>Detecting Email Spam – the Optimal Tree</i>	273
10.6.1	The one-standard-deviation rule	273
10.7	<i>Interpretation and Presentation of the rpart Output</i>	275
10.7.1	Data for female heart attack patients	275
10.7.2	Printed Information on Each Split	276
10.8	<i>Additional Notes</i>	278
10.9	<i>Further Reading</i>	279
10.10	<i>Exercises</i>	279
11	Multivariate Data Exploration and Discrimination	281
11.1	<i>Multivariate Exploratory Data Analysis</i>	282
11.1.1	Scatterplot matrices	282
11.1.2	Principal components analysis	282
11.2	<i>Discriminant Analysis</i>	285
11.2.1	Example – plant architecture	286
11.2.2	Classical Fisherian discriminant analysis	287
11.2.3	Logistic discriminant analysis	289
11.2.4	An example with more than two groups	290
11.3	<i>Principal Component Scores in Regression</i>	291
11.4*	<i>Propensity Scores in Regression Comparisons – Labor Training Data</i>	295
11.5	<i>Further Reading</i>	297
11.6	<i>Exercises</i>	298
12	The R System – Additional Topics	300
12.1	<i>Graphs in R</i>	300
12.2	<i>Functions – Some Further Details</i>	303
12.2.1	Common useful functions	303
12.2.2	User-written R functions	307
12.2.3	Functions for working with dates	310
12.3	<i>Data input and output</i>	310
12.3.1	Input	311
12.3.2	Data output	314
12.4	<i>Factors – Additional Comments</i>	315
12.5	<i>Missing Values</i>	317
12.6	<i>Lists and Data Frames</i>	320
12.6.1	Data frames as lists	320
12.6.2	Reshaping data frames; <code>reshape()</code>	320
12.6.3	Joining data frames and vectors – <code>cbind()</code>	322
12.6.4	Conversion of tables and arrays into data frames	322

12.6.5*	Merging data frames – <code>merge()</code>	322
12.6.6	The function <code>sapply()</code> and related functions	323
12.6.7	Splitting vectors and data frames into lists – <code>split()</code>	324
12.7*	<i>Matrices and Arrays</i>	324
12.7.1	Outer products	326
12.7.2	Arrays	327
12.8	<i>Classes and Methods</i>	328
12.8.1	Printing and summarizing model objects	328
12.8.2	Extracting information from model objects	329
12.9	<i>Databases and Environments</i>	330
12.9.1	Workspace management	331
12.9.2	Function environments, and lazy evaluation	331
12.10	<i>Manipulation of Language Constructs</i>	333
12.11	<i>Further Reading</i>	334
12.12	<i>Exercises</i>	334
	Epilogue – Models	338
	Appendix – S-PLUS Differences	341
	References	346
	Index of R Symbols and Functions	352
	Index of Terms	356
	Index of Names	361

Preface

This book is an exposition of statistical methodology that focuses on ideas and concepts, and makes extensive use of graphical presentation. It avoids, as much as possible, the use of mathematical symbolism. It is particularly aimed at scientists who wish to do statistical analyses on their own data, preferably with reference as necessary to professional statistical advice. It is intended to complement more mathematically oriented accounts of statistical methodology. It may be used to give students with a more specialist statistical interest exposure to practical data analysis.

The authors can claim, between them, 40 years of experience in working with researchers from many different backgrounds. Initial drafts of the monograph were constructed from notes that the first author prepared for courses for researchers, first of all at the University of Newcastle (Australia) over 1996–1997, and greatly developed and extended in the course of work in the Statistical Consulting Unit at The Australian National University over 1998–2001. We are grateful to those who have discussed their research with us, brought us their data for analysis, and allowed us to use it in the examples that appear in the present monograph. At least these data will not, as often happens once data have become the basis for a published paper, gather dust in a long-forgotten folder!

We have covered a range of topics that we consider important for many different areas of statistical application. This diversity of sources of examples has benefits, even for those whose interests are in one specific application area. Ideas and applications that are useful in one area often find use elsewhere, even to the extent of stimulating new lines of investigation. We hope that our book will stimulate such cross-fertilization. As is inevitable in a book that has this broad focus, there will be specific areas – perhaps epidemiology, or psychology, or sociology, or ecology – that will regret the omission of some methodologies that they find important.

We use the R system for the computations. The R system implements a dialect of the influential S language that is the basis for the commercial S-PLUS system. It follows S in its close linkage between data analysis and graphics. Its development is the result of a co-operative international effort, bringing together an impressive array of statistical computing expertise. It has quickly gained a wide following, among professionals and non-professionals alike. At the time of writing, R users are restricted, for the most part, to a command line interface. Various forms of graphical user interface will become available in due course.

The R system has an extensive library of packages that offer state-of-the-art-abilities. Many of the analyses that they offer were not, 10 years ago, available in any of the standard

packages. What did data analysts do before we had such packages? Basically, they adapted more simplistic (but not necessarily simpler) analyses as best they could. Those whose skills were unequal to the task did unsatisfactory analyses. Those with more adequate skills carried out analyses that, even if not elegant and insightful by current standards, were often adequate. Tools such as are available in R have reduced the need for the adaptations that were formerly necessary. We can often do analyses that better reflect the underlying science. There have been challenging and exciting changes from the methodology that was typically encountered in statistics courses 10 or 15 years ago.

The best any analysis can do is to highlight the information in the data. No amount of statistical or computing technology can be a substitute for good design of data collection, for understanding the context in which data are to be interpreted, or for skill in the use of statistical analysis methodology. Statistical software systems are one of several components of effective data analysis.

The questions that statistical analysis is designed to answer can often be stated simply. This may encourage the layperson to believe that the answers are similarly simple. Often, they are not. Be prepared for unexpected subtleties. Effective statistical analysis requires appropriate skills, beyond those gained from taking one or two undergraduate courses in statistics. There is no good substitute for professional training in modern tools for data analysis, and experience in using those tools with a wide range of data sets. No one should be embarrassed that they have difficulty with analyses that involve ideas that professional statisticians may take 7 or 8 years of professional training and experience to master.

Influences on the Modern Practice of Statistics

The development of statistics has been motivated by the demands of scientists for a methodology that will extract patterns from their data. The methodology has developed in a synergy with the relevant supporting mathematical theory and, more recently, with computing. This has led to methodologies and supporting theory that are a radical departure from the methodologies of the pre-computer era.

Statistics is a young discipline. Only in the 1920s and 1930s did the modern framework of statistical theory, including ideas of hypothesis testing and estimation, begin to take shape. Different areas of statistical application have taken these ideas up in different ways, some of them starting their own separate streams of statistical tradition. Gigerenzer et al. (1989) examine the history, commenting on the different streams of development that have influenced practice in different research areas.

Separation from the statistical mainstream, and an emphasis on “black box” approaches, have contributed to a widespread exaggerated emphasis on tests of hypotheses, to a neglect of pattern, to the policy of some journal editors of publishing only those studies that show a statistically significant effect, and to an undue focus on the individual study. Anyone who joins the R community can expect to witness, and/or engage in, lively debate that addresses these and related issues. Such debate can help ensure that the demands of scientific rationality do in due course win out over influences from accidents of historical development.

New Tools for Statistical Computing

We have drawn attention to advances in statistical computing methodology. These have led to new powerful tools for exploratory analysis of regression data, for choosing between alternative models, for diagnostic checks, for handling non-linearity, for assessing the predictive power of models, and for graphical presentation. In addition, we have new computing tools that make it straightforward to move data between different systems, to keep a record of calculations, to retrace or adapt earlier calculations, and to edit output and graphics into a form that can be incorporated into published documents.

One can think of an effective statistical analysis package as a workshop (this analogy appears in a simpler form in the JMP Start Statistics Manual (SAS Institute Inc. 1996, p. xiii).). The tools are the statistical and computing abilities that the package provides. The layout of the workshop, the arrangement both of the tools and of the working area, is important. It should be easy to find each tool as it is needed. Tools should float back of their own accord into the right place after use! In other words, we want a workshop where mending the rocking chair is a pleasure!

The workshop analogy is worth pursuing further. Different users have different requirements. A hobbyist workshop will differ from a professional workshop. The hobbyist may have less sophisticated tools, and tools that are easy to use without extensive training or experience. That limits what the hobbyist can do. The professional needs powerful and highly flexible tools, and must be willing to invest time in learning the skills needed to use them.

Good graphical abilities, and good data manipulation abilities, should be a high priority for the hobbyist statistical workshop. Other operations should be reasonably easy to implement when carried out under the instructions of a professional. Professionals also require top rate graphical abilities. The focus is more on flexibility and power, both for graphics and for computation. Ease of use is important, but not at the expense of power and flexibility.

A Note on the R System

The R system implements a dialect of the S language that was developed at AT&T Bell Laboratories by Rick Becker, John Chambers and Allan Wilks. Versions of R are available, at no cost, for 32-bit versions of Microsoft Windows, for Linux and other Unix systems, and for the Macintosh. It is available through the Comprehensive R Archive Network (CRAN). Go to <http://cran.r-project.org/>, and find the nearest mirror site.

The citation for John Chambers' 1998 Association for Computing Machinery Software award stated that S has "forever altered how people analyze, visualize and manipulate data." The R project enlarges on the ideas and insights that generated the S language. We are grateful to the R Core Development Team, and to the creators of the various R packages, for bringing into being the R system – this marvellous tool for scientific and statistical computing, and for graphical presentation.

Acknowledgements

Many different people have helped us with this project. Winfried Theis (University of Dortmund, Germany) and Detlef Steuer (University of the Federal Armed Forces, Hamburg,

Germany) helped with technical aspects of working with \LaTeX , with setting up a cvs server to manage the \LaTeX files, and with helpful comments. Lynne Billard (University of Georgia, USA), Murray Jorgensen (University of Waikato, NZ) and Berwin Turlach (University of Western Australia) gave valuable help in the identification of errors and text that required clarification. Susan Wilson (Australian National University) gave welcome encouragement. Duncan Murdoch (University of Western Ontario) helped set up the *DAAG* package. Thanks also to Cath Lawrence (Australian National University) for her Python program that allowed us to extract the R code, as and when required, from our \LaTeX files. The failings that remain are, naturally, our responsibility.

There are a large number of people who have helped with the providing of data sets. We give a list, following the list of references for the data near the end of the book. We apologize if there is anyone that we have inadvertently failed to acknowledge. Finally, thanks to David Tranah of Cambridge University Press, for his encouragement and help in bringing the writing of this monograph to fruition.

References

Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J. & Krüger, L. 1989. *The Empire of Chance*. Cambridge University Press.
SAS Institute Inc. 1996. *JMP Start Statistics*. Duxbury Press, Belmont, CA.

These (and all other) references also appear in the consolidated list of references near the end of the book.

Conventions

Text that is R code, or output from R, is printed in a verbatim text style. For example, in Chapter 1 we will enter data into an R object that we call `austpop`. We will use the `plot()` function to plot these data. The names of R packages, including our own *DAAG* package, are printed in italics.

Starred exercises and sections identify more technical items that can be skipped at a first reading.

Web sites for supplementary information

The *DAAG* package, the R scripts that we present, and other supplementary information, are available from

<http://cbis.anu.edu/DAAG>
<http://www.stats.uwo.ca/DAAG>

Solutions to exercises

Solutions to selected exercises are available from the website

<http://www.maths.anu.edu.au/~johnm/r-book.html>
See also www.cambridge.org/0521813360

A Chapter by Chapter Summary

Chapter 1: A Brief Introduction to R

This chapter aims to give enough information on the use of R to get readers started.

Note R's extensive online help facilities. Users who have a basic minimum knowledge of R can often get needed additional information from the help pages as the demand arises. A facility in using the help pages is an important basic skill for R users.

Chapter 2: Style of Data Analysis

Knowing how to explore a set of data upon encountering it for the first time is an important skill. What graphs should one draw?

Different types of graph give different views of the data. Which views are likely to be helpful?

Transformations, especially the logarithmic transformation, may be a necessary preliminary to data analysis.

There is a contrast between exploratory data analysis, where the aim is to allow the data to speak for themselves, and confirmatory analysis (which includes formal estimation and testing), where the form of the analysis should have been largely decided before the data were collected.

Statistical analysis is a form of data summary. It is important to check, as far as this is possible that summarization has captured crucial features of the data. Summary statistics, such as the mean or correlation, should always be accompanied by examination of a relevant graph. For example, the correlation is a useful summary, if at all, only if the relationship between two variables is linear. A scatterplot allows a visual check on linearity.

Chapter 3: Statistical Models

Formal data analyses assume an underlying statistical model, whether or not it is explicitly written down.

Many statistical models have two components: a *signal* (or deterministic) component; and a *noise* (or error) component.

Data from a sample (commonly assumed to be randomly selected) are used to fit the model by estimating the signal component.

The fitted model determines *fitted* or *predicted* values of the signal. The *residuals* (which estimate the noise component) are what remain after subtracting the fitted values from the observed values of the signal.

The normal distribution is widely used as a model for the noise component.

Haphazardly chosen samples should be distinguished from random samples. Inference from haphazardly chosen samples is inevitably hazardous. Self-selected samples are particularly unsatisfactory.

Chapter 4: An Introduction to Formal Inference

Formal analysis of data leads to inferences about the population(s) from which the data were sampled. Statistics that can be computed from given data are used to convey information about otherwise unknown population parameters.

The inferences that are described in this chapter require randomly selected samples from the relevant populations.

A *sampling distribution* describes the theoretical distribution of sample values of a statistic, based on multiple *independent* random samples from the population.

The standard deviation of a sampling distribution has the name *standard error*.

For sufficiently large samples, the normal distribution provides a good approximation to the true sampling distribution of the mean or a difference of means.

A *confidence interval* for a parameter, such as the mean or a difference of means, has the form

$$\text{statistic} \pm t\text{-critical-value} \times \text{standard error.}$$

Such intervals give an assessment of the level of uncertainty when using a sample statistic to estimate a population parameter.

Another viewpoint is that of *hypothesis testing*. Is there sufficient evidence to believe that there is a difference between the means of two different populations?

Checks are essential to determine whether it is plausible that confidence intervals and hypothesis tests are valid. Note however that plausibility is not proof!

Standard chi-squared tests for two-way tables assume that items enter independently into the cells of the table. Even where such a test is not valid, the standardized residuals from the “no association” model can give useful insights.

In the one-way layout, in which there are several independent sets of sample values, one for each of several groups, data structure (e.g. compare treatments with control, or focus on a small number of “interesting” contrasts) helps determine the inferences that are appropriate. In general, it is inappropriate to examine all possible comparisons.

In the one-way layout with quantitative levels, a regression approach is usually appropriate.

Chapter 5: Regression with a Single Predictor

Correlation can be a crude and unduly simplistic summary measure of dependence between two variables. Wherever possible, one should use the richer regression framework to gain deeper insights into relationships between variables.

The line or curve for the regression of a response variable y on a predictor x is different from the line or curve for the regression of x on y . Be aware that the inferred relationship is conditional on the values of the predictor variable.

The model matrix, together with estimated coefficients, allows for calculation of predicted or fitted values and residuals.

Following the calculations, it is good practice to assess the fitted model using standard forms of graphical diagnostics.

Simple alternatives to straight line regression using the data in their raw form are

- transforming x and/or y ,
- using polynomial regression,
- fitting a smooth curve.

For size and shape data the allometric model is a good starting point. This model assumes that regression relationships among the logarithms of the size variables are linear.

Chapter 6: Multiple Linear Regression

Scatterplot matrices may provide useful insight, prior to fitting a regression model.

Following the fitting of a regression, one should examine relevant diagnostic plots.

Each regression coefficient estimates the effect of changes in the corresponding explanatory variable when other explanatory variables are held constant.

The use of a different set of explanatory variables may lead to large changes in the coefficients for those variables that are in both models.

Selective influences in the data collection can have a large effect on the fitted regression relationship.

For comparing alternative models, the AIC or equivalent statistic (including Mallows C_p) can be useful. The R^2 statistic has limited usefulness.

If the effect of variable selection is ignored, the estimate of predictive power can be grossly inflated.

When regression models are fitted to observational data, and especially if there are a number of explanatory variables, estimated regression coefficients can give misleading indications of the effects of those individual variables.

The most useful test of predictive power comes from determining the predictive accuracy that can be expected from a new data set.

Cross-validation is a powerful and widely applicable method that can be used for assessing the expected predictive accuracy in a new sample.

Chapter 7: Exploiting the Linear Model Framework

In the study of regression relationships, there are many more possibilities than regression lines! If a line is adequate, use that. But one is not limited to lines!

A common way to handle qualitative factors in linear models is to make the initial level the baseline, with estimates for other levels estimated as offsets from this baseline.

Polynomials of degree n can be handled by introducing into the model matrix, in addition to a column of values of x , columns corresponding to x^2, x^3, \dots, x^n . Typically, $n = 2, 3$ or 4 .

Multiple lines are fitted as an interaction between the variable and a factor with as many levels as there are different lines.

Scatterplot smoothing, and smoothing terms in multiple linear models, can also be handled within the linear model framework.

Chapter 8: Logistic Regression and Other Generalized Linear Models

Generalized linear models (GLMs) are an extension of linear models, in which a function of the expectation of the response variable y is expressed as a linear model. A further generalization is that y may have a binomial or Poisson or other non-normal distribution.

Common important GLMs are the logistic model and the Poisson regression model.

Survival analysis may be seen as a further specific extension of the GLM framework.

Chapter 9: Multi-level Models, Time Series and Repeated Measures

In a multi-level model, the random component possesses structure; it is a sum of distinct error terms.

Multi-level models that exhibit suitable balance have traditionally been analyzed within an analysis of variance framework. Unbalanced multi-level designs require the more general multi-level modeling methodology.

Observations taken over time often exhibit time-based dependence. Observations that are close together in time may be more highly correlated than those that are widely separated. The autocorrelation function can be used to assess levels of serial correlation in time series.

Repeated measures models have measurements on the same individuals at multiple points in time and/or space. They typically require the modeling of a correlation structure similar to that employed in analyzing time series.

Chapter 10: Tree-based Classification and Regression

Tree-based models make very weak assumptions about the form of the classification or regression model. They make limited use of the ordering properties of continuous or ordinal explanatory variables. They are unsuitable for use with small data sets.

Tree-based models can be an effective tool for analyzing data that are non-linear and/or involve complex interactions.

The decision trees that tree-based analyses generate may be complex, giving limited insight into model predictions.

Cross-validation, and the use of training and test sets, are essential tools both for choosing the size of the tree and for assessing expected accuracy on a new data set.

Chapter 11: Multivariate Data Exploration and Discrimination

Principal components analysis is an important multivariate exploratory data analysis tool.

Examples are presented of the use of two alternative discrimination methods – logistic regression including multivariate logistic regression, and linear discriminant analysis.

Both principal components analysis, and discriminant analysis, allow the calculation of scores, which are values of the principal components or discriminant functions, calculated observation by observation. The scores may themselves be used as variables in, e.g., a regression analysis.

Chapter 12: The R System – Additional Topics

This final chapter gives pointers to some of the further capabilities of R. It hints at the marvellous power and flexibility that are available to those who extend their skills in the use of R beyond the basic topics that we have treated. The information in this chapter is intended, also, for use as a reference in connection with the computations of earlier chapters.