

Statistics and Computing

Series Editors:

J. Chambers

D. Hand

W. Härdle

Statistics and Computing

Brusco/Stahl: Branch-and-Bound Applications in Combinatorial Data Analysis.

Dalgaard: Introductory Statistics with R.

Gentle: Elements of Computational Statistics.

Gentle: Numerical Linear Algebra for Applications in Statistics.

Gentle: Random Number Generation and Monte Carlo Methods, 2nd Edition.

Härdle/Klinke/Turlach: XploRe: An Interactive Statistical Computing Environment.

Krause/Olson: The Basics of S-PLUS, 4th Edition.

Lange: Numerical Analysis for Statisticians.

Lemmon/Schafer: Developing Statistical Software in Fortran 95

Loader: Local Regression and Likelihood.

Ó Ruanaidh/Fitzgerald: Numerical Bayesian Methods Applied to Signal Processing.

Pannatier: VARIOWIN: Software for Spatial Data Analysis in 2D.

Pinheiro/Bates: Mixed-Effects Models in S and S-PLUS.

Venables/Ripley: Modern Applied Statistics with S, 4th Edition.

Venables/Ripley: S Programming.

Wilkinson: The Grammar of Graphics.

Andreas Krause
Melvin Olson

The Basics of S-PLUS

Fourth Edition



Andreas Krause
Novartis Pharma AG
Basel 4002
Switzerland

Melvin Olson
Novartis Pharma AG
Basel 4002
Switzerland

Series Editors:

J. Chambers
Bell Labs, Lucent
Technologies
600 Mountain Ave.
Murray Hill, NJ 07974
USA

D. Hand
Department of Mathematics
South Kensington Campus
Imperial College, London
London
SW7 2AZ
United Kingdom

W. Härdle
Institut für Statistik und
Ökonometrie
Humboldt-Universität zu Berlin
Spandauer Str. 1
D-10178 Berlin, Germany

Library of Congress Control Number: 2005926768

ISBN-10: 0-387-26109-5
ISBN-13: 978-0387-26109-6

e-ISBN: 0-387-22708-3

Printed on acid-free paper.

© 2005 Springer Science+Business Media, 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 Science+Business Media, Inc., 233 Spring Street, New York, NY 10013, 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 in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed in the United States of America. (SBA)

9 8 7 6 5 4 3 2 1

springeronline.com

Preface

This is now the fourth edition of “The Basics of S-Plus” since 1997. S-PLUS saw a steady growth in popularity, and it established itself in many educational and business places as a major data analysis tool. S-PLUS is valued for its modern, interactive data analysis environment, whether it is the primary system or a complement to other standards like SAS (the latter is in particular true for the industry we work in, pharmaceuticals).

We have followed the various releases with new editions of our book, introducing over time major changes like the incorporation of S Version 4 (the underlying language), Trellis graphs, a graphical user interface, in particular for the Windows operating system, and a chapter on R and its differences to S-PLUS (that are minor for the material covered in this book). This edition is an update from edition 3 to cover new functions and features of S-PLUS Version 7.0 (working from the beta release for MS Windows and Linux), adding more practical tips and examples, and correcting a few mistakes.

We are very grateful to all our readers, in particular those sending us suggestions, comments, and any other kind of feedback. You will see some of these reflected in the book.

The number of new books appearing on S-PLUS and its counterpart, R, continues to grow substantially each year. If you want to get an impression of the books available, Insightful maintains a list of S-PLUS references on their web site, <http://www.insightful.com/>. The list currently comprises about 50 books and manuscripts. You can also search for the string “Splus” or “S-Plus” (without quotes) in an online bookstore, such as <http://www.amazon.com/>.

Some of the books on S-PLUS are “Analyzing Medical Data Using S and S-PLUS,” by Everitt and Rabe-Hesketh; “Regression Modeling Strategies,” by Harrell; “Bootstrap Methods and Permutation Tests” by Hesterberg, Monaghan, et.al., “Applied Statistics in the Pharmaceutical Industry,” edited by Millard and Krause; “Mixed Effects Models in S and S-PLUS,” by Pinheiro and Bates; “Modeling Survival Data: Extending the Cox Model,” by Therneau and Grambsch; and the fourth edition of “Modern Applied Statistics with S-PLUS” as well as the second edition of “S Programming,” both by Venables and Ripley.

This book is intended to provide a head-start into the world of S-PLUS (and serves as well as an introduction to R). We are going to introduce you to S-PLUS by taking you through interactive sessions, entering commands and later analyzing data, pointing out tips and pitfalls, and complementing each session with exercises and detailed solutions.

If you have never worked with S-PLUS before and want to get a first impression, take a look at “A First Session,” “A Second Session,” and “Exploring Data” to get an idea of the philosophy of the programming language. If you intend to work with point-and-click graphical user interfaces, take a look at “Graphical User Interfaces.” We are convinced though that if you want to work permanently with S-PLUS, you will need to use the command line and appreciate it soon.

The book originates from lectures such that each chapter can serve as the basis for a 90-minute lecture. The exercises reinforce the material covered and sometimes point out a few more details. It should be interesting to work out a solution to an exercise before comparing it to the solution given here. There’s never a single solution.

Finally, we would like to thank all those who have directly or indirectly contributed to this edition and previous ones. For this edition, Chyi-Hung Hsu, Michael Merz and José Pinheiro gave very helpful comments on the manuscript. John Kimmel at Springer-Verlag is simply a great person to work with, dedicated and professional. David Smith, Michael O’Connell and others at Insightful, Inc. (Seattle) continue to be supportive partners by providing us with alpha and beta releases of upcoming S-PLUS versions. Following the discussion forum s-news provides an insight into frequently asked questions and problems that S-PLUS users face. And last but not least, various people from all over the world supplied comments on earlier editions of the book, pointed out possible improvements and errors and provided a good motivation to revise and update the book with the current edition 4 – which you happen to be reading right now.

We are always interested in any kind of comment that you might have about the book. You are welcome to contact us by sending an E-mail to Andreas.Krause@Novartis.com and Melvin.Olson@Novartis.com. Note though that we will not be able to provide general support, like answering programming questions. For these purposes please refer to the sections where we describe how to get help and support (pp. 413ff. and 422ff.). You

might also want to check if your institution has a contract with Insightful that covers support.

Finally, a Web page is set up to accompany this book:

<http://www.elmo.ch/doc/splus-book/>

It provides information about the book, links to relevant Internet pages, and a list of errata (which we are confident will be fairly short).

Basel, Switzerland
April 2005

Andreas Krause and Melvin Olson

Contents

Preface	v
Figures	xvii
Tables	xxi
1 Introduction	1
1.1 The History of S and S-PLUS	2
1.2 S-PLUS on Different Operating Systems	4
1.3 Notational Conventions	6
2 Graphical User Interface	9
2.1 Introduction	9
2.2 System Overview	10
2.2.1 Using a Mouse	11
2.2.2 Object Explorer	11
2.2.3 Commands Window	11
2.2.4 Toolbars	12
2.2.5 Graph Sheets	12
2.2.6 Script Window	12
2.3 Getting Started with the Interface	13
2.3.1 Importing Data	13
2.3.2 Graphs	13
2.3.3 Data and Statistics	15
2.3.4 Customizing the Toolbars	15
2.3.5 Chapters	16
2.4 Detailed Use of the GUI Interface	18

2.5	Object Explorer	18
2.6	Help	19
2.7	Data Export	21
2.8	Working Directory	23
2.9	Data Import	24
2.10	Data Summaries	27
2.11	Graphs	29
2.12	Trellis Graphs	36
2.13	Linear Regression	38
2.14	PowerPoint (Windows Only)	42
2.15	Excel (Windows Only)	44
2.16	Script Window	45
2.17	UNIX/Linux GUI	47
2.18	Summary	56
2.19	Exercises	57
2.20	Solutions	58
3	A First Session	73
3.1	General Information	73
3.1.1	Starting and Quitting	74
3.1.2	The Help System	75
3.1.3	Before Beginning	75
3.2	Simple Structures	76
3.2.1	Arithmetic Operators	76
3.2.2	Assignments	77
3.2.3	The Concatenate Command: <code>c</code>	79
3.2.4	The Sequence Command: <code>seq</code>	80
3.2.5	The Replicate Command: <code>rep</code>	81
3.3	Mathematical Operations	82
3.4	Use of Brackets	84
3.5	Logical Values	85
3.6	Review	88
3.7	Exercises	91
3.8	Solutions	92
4	A Second Session	95
4.1	Constructing and Manipulating Data	95
4.1.1	Matrices	96
4.1.2	Arrays	101
4.1.3	Data Frames	104
4.1.4	Lists	107
4.2	Introduction to Functions	108
4.3	Introduction to Missing Values	109
4.4	Merging Data	110
4.5	Putting It All Together	111

4.6	Exercises	114
4.7	Solutions	116
5	Graphics	125
5.1	Basic Graphics Commands	125
5.2	Graphics Devices	126
5.2.1	Working with Multiple Graphics Devices	128
5.3	Plotting Data	128
5.3.1	The <code>plot</code> Command	129
5.3.2	Modifying the Data Display	130
5.3.3	Modifying Figure Elements	131
5.4	Adding Elements to Existing Plots	133
5.4.1	Functions to Add Elements to Graphs	133
5.4.2	More About <code>abline</code>	135
5.4.3	More on Adding Axes	135
5.4.4	Adding Text to Graphs	137
5.5	Setting Options	138
5.6	Figure Layouts	140
5.6.1	Layouts Using Trellis Graphs	140
5.6.2	Matrices of Graphs	140
5.6.3	Multiple-Screen Graphs	141
5.6.4	Figures of Specified Size	142
5.7	Exercises	145
5.8	Solutions	146
6	Trellis Graphics	153
6.1	An Example	154
6.2	Trellis Basics	156
6.2.1	Trellis Syntax	156
6.2.2	Trellis Functions	157
6.2.3	Displaying and Storing Graphs	157
6.3	Output Devices	158
6.4	Customizing Trellis Graphs	160
6.4.1	Setting Options	160
6.4.2	Arranging the Layout of a Trellis Graph	161
6.4.3	Ordering of Graphs	163
6.4.4	Axis Customization	164
6.4.5	Modifying Panel Strips	165
6.4.6	Arranging Several Graphs on a Single Page	165
6.4.7	Updating Existing Trellis Graphs	167
6.4.8	Writing Panel Functions	168
6.5	Further Trellis Hints	171
6.5.1	Useful General Trellis Settings	172
6.5.2	Graphing Individual Profiles	173
6.5.3	Preparing Data to Use for Trellis	174

6.5.4	The subset Option	175
6.5.5	Adding a Key	175
6.5.6	The subscripts Option in Panel Functions	177
6.6	Exercises	181
6.7	Solutions	183
7	Exploring Data	193
7.1	Descriptive Data Exploration	193
7.2	Graphical Exploration	204
7.2.1	Interactive Dynamic Graphics	219
7.2.2	Old-Style Graphics	219
7.3	Distributions and Related Functions	220
7.4	Confirmatory Statistics and Hypothesis Testing	225
7.5	Missing and Infinite Values	231
7.5.1	Testing for Missing Values	232
7.5.2	Supplying Data with Missing Values to Functions	232
7.5.3	Missing Values in Graphs	233
7.5.4	Infinite Values	233
7.6	Exercises	235
7.7	Solutions	238
8	Statistical Modeling	251
8.1	Introductory Examples	251
8.1.1	Regression	251
8.1.2	Regression Diagnostics	253
8.2	Statistical Models	255
8.3	Model Syntax	256
8.4	Regression	257
8.4.1	Linear Regression and Modeling Techniques	258
8.4.2	ANOVA	261
8.4.3	Logistic Regression	263
8.4.4	Survival Data Analysis	265
8.4.5	Endnote	267
8.5	Exercises	268
8.6	Solutions	271
9	Programming	285
9.1	Lists	285
9.1.1	Adding and Deleting List Elements	287
9.1.2	Naming List Elements	288
9.1.3	Applying the Same Function to List Elements	290
9.1.4	Unlisting a List	294
9.1.5	Generating a List by Using <code>split</code>	294
9.2	Writing Functions	294
9.2.1	Documenting Functions	297

9.2.2	Scope of Variables	297
9.2.3	Parameters and Defaults	298
9.2.4	Passing an Unspecified Number of Parameters to a Function	300
9.2.5	Testing for Existence of an Argument	301
9.2.6	Returning Warnings and Errors	301
9.2.7	Using Function Arguments in Graphics Labels	302
9.3	Iteration	303
9.3.1	The <code>for</code> Loop	303
9.3.2	The <code>while</code> Loop	304
9.3.3	The <code>repeat</code> Loop	305
9.3.4	Vectorizing a Loop	305
9.3.5	Large Loops	307
9.4	Debugging: Searching for Errors	308
9.4.1	Syntax Errors	309
9.4.2	Invalid Arguments	310
9.4.3	Execution or Run-Time Errors	310
9.4.4	Logical Errors	311
9.5	Output Using the <code>cat</code> Function	314
9.6	The <code>paste</code> Function	316
9.7	Exercises	318
9.8	Solutions	319

10 Object-Oriented Programming 323

10.1	Creating Classes and Objects	325
10.2	Creating Methods	328
10.3	Debugging	333
10.4	Help	334
10.5	Summary and Overview	334
10.6	Exercises	335
10.7	Solutions	336

11 Input and Output 349

11.1	Reading S-PLUS Commands from a File: The <code>source</code> Function	349
11.2	Data Import/Export: Easiest Method	350
11.3	Data Import/Export: General Method	352
11.4	Data Import/Export: Basic Method	353
11.5	Reading Data from the Terminal	354
11.6	Editing Data	355
11.7	Transferring Data	356
11.8	Recording a Session	356
11.9	Exercises	358
11.10	Solutions	359

12 Tips and Tricks	363
12.1 Useful Techniques	363
12.1.1 Housekeeping: Cleaning Up Directories	363
12.1.2 Storing and Restoring Graphical Parameters	364
12.1.3 Naming of Objects	364
12.1.4 Repeating Commands	365
12.2 Programming Environment and Techniques	366
12.2.1 The Process of Developing a Function	366
12.2.2 Setting up an Editor and Running the Code in S-PLUS	366
12.2.3 Treating Data Frames as Lists	368
12.2.4 Working with Graph Sheets	369
12.2.5 Incorporating and Accessing C and Fortran Pro- grams	371
12.2.6 Batch Jobs	374
12.2.7 Libraries	376
12.3 Factors	378
12.3.1 Creating Factors and Ordered Factors	378
12.3.2 Internal Representation of Factors	380
12.3.3 Where Levels Play a Role	381
12.3.4 Where Factors Can Lead Their Own Lives	382
12.3.5 How Factors Come Into Life	384
12.3.6 Adding and Dropping Factor Levels	385
12.4 Including Graphs in Text Processors	386
12.4.1 Generating Graphs for Windows Applications	387
12.4.2 Generating PostScript Graphs	388
12.4.3 PostScript Graphs in L ^A T _E X	389
12.4.4 If You Don't Have a PostScript Printer	390
12.4.5 Greek Letters in Graphs	390
12.5 Exercises	392
12.6 Solutions	394
13 S-Plus Internals	401
13.1 How S-PLUS Works Under UNIX	401
13.1.1 The Working Chapter	402
13.1.2 Customization on Start-Up and Exit	402
13.2 How S-PLUS Works Under Windows	404
13.2.1 Command Line Options	404
13.2.2 Start-up and Exit Functions	405
13.2.3 How the Script Window works	406
13.3 Storing Mechanism	407
13.4 Levels of Calls	408
13.5 Exercises	410
13.6 Solutions	411

14 Information Sources on and Around S-Plus	413
14.1 Insightful	413
14.2 S-News: Exchanging Information with Other Users	414
14.3 The StatLib Server	414
14.4 What Next?	415
15 R	417
15.1 Development	418
15.2 Some Similarities Between R and S	418
15.3 Some Differences Between R and S	418
15.3.1 Language	419
15.3.2 Libraries	420
15.3.3 Trellis-Type Graphs	420
15.3.4 Colors and Lines	421
15.3.5 Data Import and Export Formats	421
15.3.6 Memory Handling	421
15.3.7 Mathematical Formulae in Graphs	421
15.3.8 Graphical User Interfaces	421
15.3.9 Start-Up Mechanism	422
15.3.10 Windows Integration	422
15.3.11 Support	422
15.4 Summary	423
16 Bibliography	425
16.1 Print Bibliography	425
16.2 On-Line Bibliography	427
16.2.1 S-PLUS Related Sources	427
16.2.2 T _E X-Related Sources	429
16.2.3 Other Sources	429
Index	431

Figures

2.1	The S-PLUS screen and its components	10
2.2	The S-PLUS main toolbar	12
2.3	The 2D graph palette	14
2.4	The Object Explorer toolbar	19
2.5	The SearchPath	20
2.6	Details of the Puromycin data set	20
2.7	Data export	22
2.8	Defining a working directory	24
2.9	Data import window	25
2.10	Sorting a data frame	26
2.11	Sorted Puromycin data	26
2.12	Summary Statistics data dialog	28
2.13	Summary Statistics dialog	28
2.14	Summary Statistics Report window	29
2.15	2D Plot dialog (symbol)	31
2.16	Axis dialog box	32
2.17	Axis font dialog box	32
2.18	Annotated scatterplot	33
2.19	Saving a graph sheet	35
2.20	Trellis graphs – drag-and-drop	36
2.21	A Trellis graph conditioning on state	37
2.22	Linear Regression dialog	39
2.23	Regression Formula dialog	39
2.24	Regression Report window	40

2.25	Combining graphs onto one graph sheet	43
2.26	The Excel toolbar	45
2.27	The Script Window	46
2.28	Script Window toolbar	47
2.29	The UNIX/Linux Startup window	48
2.30	The Objects Summary Window	49
2.31	Data specification dialog	50
2.32	Data viewer	50
2.33	Trellis scatterplot data specification	51
2.34	Trellis scatterplot symbol options	52
2.35	Trellis scatterplot	52
2.36	Summary Statistics dialog	53
2.37	Summary Statistics Report window	54
2.38	Linear Regression specification dialog	54
2.39	Linear Regression Report window	55
2.40	Linear regression residual plots	55
2.41	Inserting a column of data	58
2.42	Loess smoother curves	59
2.43	Variable transformation for regression	60
2.44	Regression output using $\log(\text{conc})$	61
2.45	Residual plots using $\log(\text{conc})$	62
2.46	Regression formula using state	62
2.47	Regression output using state	63
2.48	Regression formula using interaction term	64
2.49	Regression output using interaction term	64
2.50	Object Explorer	65
2.51	Plot Symbol dialog	68
2.52	Geyser scatterplot	69
2.53	The Annotation palette	69
2.54	Plot Line dialog	70
2.55	Geyser scatterplot with annotation	71
5.1	Different options for the type parameter of the plot function	131
5.2	An example of a graph layout	132
5.3	A customized graph	133
5.4	The S-PLUS definition of figure and plot regions	138
5.5	A customized graph layout	143
5.6	A customized display of the Geyser data set	144
5.7	Trigonometric functions	146
5.8	Geometric figures (a): a rectangle.	148
5.9	Geometric figures (b): a circle	149
5.10	Geometric figures (c): a spiral	150
5.11	Lissajous figures	151

6.1	An example of a Trellis graph layout	154
6.2	Histograms of heights of singers for different voice parts .	155
6.3	Customizing a Trellis graph arrangement	162
6.4	Graph ordering in a Trellis arrangement	163
6.5	A multigraph page of Trellis displays	166
6.6	A multigraph page of Trellis displays with different sizes .	167
6.7	Trellis display with customized panel function	170
6.8	Time profiles of Theophylline concentration	174
6.9	Concentration-time profiles with dosing indicated	180
6.10	Trellis display of data spread, mean, and standard deviations	182
6.11	Trellis histograms with added key of derived statistics. . .	185
6.12	Trellis display of the Geyser data set (duration vs. waiting time) with regression lines added	187
6.13	Trellis display of the Geyser data set (waiting time vs. duration) with regression lines added	188
6.14	Trellis display of data spread, mean, and standard deviations. Group “g” consists only of missing values	191
6.15	Trellis display of data spread, mean, and standard deviations with non-default parameter settings	192
7.1	Histogram of the barley yield	205
7.2	Histogram of the barley yield by site	205
7.3	Histogram of the barley yield by site and year	206
7.4	Boxplot display of the barley yield split by site of harvest	207
7.5	Boxplot display of the barley yield for all sites by year, employing the <code>boxplot</code> function	208
7.6	Boxplot display of the barley yield for all years by site . .	209
7.7	The Old Faithful Geyser data	211
7.8	The Old Faithful Geyser data in subsets	213
7.9	A perspective plot of the Geyser data using <code>persp</code>	214
7.10	A perspective plot of the Geyser data using <code>wireframe</code> . .	215
7.11	An image plot of the Geyser data using <code>image</code>	216
7.12	A contour plot of the Geyser data using <code>contour</code>	217
7.13	A contour plot of the Geyser data using <code>contourplot</code> . .	218
7.14	A graphical comparison of three distributions	225
7.15	A graph of the binomial distribution	226
7.16	A graphical display of the <i>t</i> -test	227
7.17	Different histogram displays of normally distributed data	238
7.18	The Car data set	243
7.19	Estimating pi using a simulation technique	248
8.1	Regression diagnostic plots	254
8.2	Typical regression diagnostic plots	278
8.3	GAM regression diagnostic plots	279
8.4	Design and interaction plots for ANOVA	283

9.1	Two-dimensional functions drawn with <code>persp</code>	320
10.1	Country data graph for a single country.	331
10.2	Country data graph for two countries	332
10.3	Country data graph for four countries	336
11.1	Capital development with a given interest rate	362
12.1	The “play” button to run S-PLUS scripts	367
12.2	Renaming graph sheet pages	370
12.3	Barley site factor levels reordered by maximum yield . . .	395

Tables

1.1	S-PLUS system requirements	5
1.2	Notational conventions	7
2.1	File types available with the Import From File facility	13
3.1	Arithmetic operations on vectors	83
3.2	Different brackets in S-PLUS	84
3.3	Symbols for logical operators	86
3.4	Truncation functions	94
3.5	Truncation results	94
5.1	The most common output devices	127
5.2	Commands for graphics devices	129
5.3	Options for <code>type</code> in the <code>plot</code> command	130
5.4	Options for plotting functions	132
5.5	Commands that add to existing graphs	134
5.6	Some examples of uses for <code>abline</code>	135
5.7	Optional parameters for plotting axes	136
5.8	Optional parameters for plotting axes using the <code>axis</code> function	136
5.9	Optional parameters for adding text	137
5.10	Layout parameters	139
6.1	Trellis functions	157

6.2	Trellis graphics devices	160
6.3	Trellis plot options	161
6.4	Trellis panel functions	168
7.1	Descriptive statistics functions	200
7.2	Tabulation and split functions	200
7.3	Multivariate dynamic graphics functions	219
7.4	Old-style graphics functions	219
7.5	Old-style multivariate graphics functions	220
7.6	Categorization of distribution-related functions in S-PLUS	221
7.7	Distribution-related functions in S-PLUS:	
	Discrete distributions	222
7.8	Distribution-related functions in S-PLUS:	
	Continuous distributions	222
7.9	Statistical tests in S-PLUS	228
7.10	Estimating pi using a simulation technique	247
8.1	Swiss fertility data set	252
8.2	Selected statistical models in S-PLUS	256
8.3	Model syntax	257
8.4	Kyphosis data frame	263
8.5	Leukemia data frame	266
9.1	Some commands in browser , debugger , and recover . .	315
9.2	Control sequences in S-PLUS	315
11.1	File types and extensions with the import/export facility in UNIX	351
11.2	File types and extensions with the import/export facility in Windows	351
12.1	S-PLUS types and C/Fortran equivalents	373

1

Introduction

Over the years, the S language and S-PLUS have undergone many changes. Since its development in the mid-1970s, the three main authors of S, Rick Becker, John Chambers, and Allan Wilks, have enhanced the entire language considerably. All their work was done at Bell Labs with the original goal of defining a language to make it easier to do repetitive tasks in data analysis, such as calculating a linear model.

In the following years, many people contributed to the S project in one form or another. People outside Bell Labs also became aware of the interesting development and took part in it, and this is to a great extent the way S and S-PLUS are still developed today. A very lively user community works with and on S/S-PLUS, and they especially appreciate the S style of working in an integrated environment. Special strengths are the modern and flexible language and high quality graphics generation.

It is noteworthy that the authors do not consider S as a primarily statistical system. The system was developed to be flexible and interactive, especially designed for doing data analysis of an exploratory nature, which began to boom in the late 1970s after the release of Tukey's book (1977) on the subject. Most of the statistical functionality was added later, and many statistics routines, such as estimation, regression, and testing, were incorporated by the S-PLUS team.

This chapter describes the development of S and S-PLUS over the years, clarifies the differences between the two, and points to some further references.

1.1 The History of S and S-PLUS

The S Language, and some years later the NEW S, were developed at AT&T Bell Labs in the late 1970s and early 1980s, mainly by Rick Becker and John Chambers. Some years later, Allan Wilks joined the core team. Several other people have been involved in the project. Becker (1994) describes in great detail the foundation and development of S until 1994 (S Version 3). S Version 4 was developed by John Chambers at Bell Laboratories, Lucent Technologies.

The year 1976 can perhaps be viewed as the founding year of S, the year in which the first concepts were developed and implemented. At first, the “system” consisted of a library of routines together with an interface, such that the kernel code itself could be kept unmodified. In 1981, the S team decided to rewrite the system in C and port it to the UNIX operating system. From 1981 to 1993, the source code was available to interested people outside Bell Labs.

The next years revealed a strongly increasing interest among statisticians and data analysts in using the system, which was called S by then. It is remarkable that the important steps in the development of S were all marked by books, such that S users today talk about the days of the *Brown Book*, the *Blue Book*, and the *White Book*. The *Green Book* marks the most recent milestone.

In 1984, as the interest in S began to rise, a real manual was necessary. The first book, today referred to as the *Brown Book*, was written by Becker and Chambers (1984). This version of S is now referred to as “Old S,” as no version numbers existed at the time.

The QPE (Quantitative Programming Environment) developed by John Chambers set a milestone in the development of S. In 1988, it introduced the function concept (replacing the former macros), and new programming concepts were added. This work is described in the *Blue Book* (Becker, Chambers, and Wilks, 1988).

During all these years, the user community added substantial functionality to S, and many sophisticated techniques such as tree regression, nonparametric smoothing, survival analysis, object-oriented programming, and new model formulation became a part of S. This step in the development was manifested and accompanied by the *White Book* (Chambers and Hastie, 1992).

Version 4 of S came out in 1998 and is described in full detail in Chambers (1998), the *Green Book*. The object-oriented paradigm forms the basis of the entire language, documentation is integrated into an object, and the general paradigm “everything is an object” is followed throughout the language.

In 1987, Douglas Martin at the University of Washington, Seattle, founded a small company to make S more popular. He realized that the major drawback of S was the need of professional support for end-users. Hence,

he started the company Statistical Sciences, Inc. (StatSci). StatSci became the Data Analysis Products Division of MathSoft in 1994. When the company became independent again in 2001, it received a new name: Insightful Corporation.

The company added more functionality to S, ported it to other hardware platforms, and provided the necessary support for technical and statistical problems. The enhanced version of S received a new name: S-PLUS.

S-PLUS helped popularize S among nontechnical people. StatSci ported S-PLUS to the only non-UNIX platform, releasing S-PLUS for DOS in 1989 and S-PLUS for Windows in 1993. Up to version 3, S-PLUS for DOS/Windows and UNIX provided essentially the same functionality.

S-PLUS Version 4 was only released for the Windows platform. It was enhanced by a new graphical user interface (GUI) and a graphics system based on the Axum package. It adopts the Windows standard and gives a menu-based interface to S-PLUS. Windows-specific functionality such as creation of PowerPoint slides or direct data exchange with other programs via DLLs was added. Much of the S-PLUS functionality is available via menus and buttons, and the graphics are shown in an editable graph sheet. A major strength is that all functionality that is accessible via the menus can also be called from the command line. S-PLUS comes with a script window where the corresponding command line input is shown.

In fall 1998, S-PLUS for Windows was released in version 4.5, which is split into a “Professional Edition” that comprises the full functionality and a menu-only version called “S-PLUS Standard.”

Version 5 of S-PLUS was only released on UNIX systems in late 1998 and 1999. It is based on S Version 4 as described in Chambers (1998). For the first time, the Linux system for Intel PCs is supported.

S-PLUS 2000 for Windows systems (analogous to Office 2000 by Microsoft) was released in mid-1999, still based on S Version 3.

Insightful released S-PLUS Version 6 for UNIX in 2000 and for Windows in 2001. For the first time, the two systems are built from the same code base. The UNIX system has a graphical user interface based on Java and supports some Windows-specific elements such as reading Excel data files and writing graphs in WMF (Windows MetaFile) format.

Version 6.2 introduced an important feature for automated batch runs, the so-called verbose logging. By calling S-PLUS with the switch /SBATCH an exhaustive log file is created that facilitates scanning for possible problems in batch runs.

Version 7.0 of S-PLUS is about to be released by the time of writing this edition of the book (early 2005). We have an alpha and a beta version of S-PLUS 7.0 for Windows and a beta version for Linux available. The new major version includes a library for handling big data sets without keeping it entirely in memory. The advertised limits are “no limits in the number of rows” and “tens of thousands of columns.” The library comprises functions for extracting simple statistics from columns and for calculating statistical

models such as linear regression, generalized linear models, principal components, and clustering on big data sets. The GUI is revised and now based on a new toolkit, Eclipse. It comes with a workbench to develop S-PLUS programs, track changes, and source control facilities. These new functionalities are only available for the newly introduced “Enterprise Developer” edition though and not part of the “Professional Developer” edition.

S is still the heart of the system, and the core S team continues to work on the S system. The whole S functionality is incorporated in S-PLUS and enhanced, and today the S system is no longer publicly available. In the remainder of the book, we will use S-PLUS as the standard reference.

In addition to S-PLUS, the R project started in 1997. R is a software that is “not unlike S” and freely available. It has become increasingly popular in recent years since a core developer team continuously improved and developed the system further. A fairly large group of supporters around the world contributes libraries, in very much the way S started out a long time before. This group of people includes many of the contributors of S-PLUS libraries. We will point out some of the similarities and differences between the two systems in a separate chapter.

1.2 S-PLUS on Different Operating Systems

Sometimes it is important to know about differences in software on various hardware systems or under different operating systems. This can be the case if you work on more than one computer system with S-PLUS (and therefore need to exchange data files) or if you want to be informed about the differences before deciding in favor of a specific system. In this section, we discuss some details about different systems supported by S-PLUS.

In addition, the chapter provides some basic information about the general setup of files and structures in S-PLUS. More information on the internal workings of S-PLUS can be found in Sections 13.1 and 13.2.

At present, S-PLUS supports two major operating systems: UNIX (with most of its variants) and Windows (with most of its variants). Table 1.1 summarizes the currently supported hardware and operating systems.

As the S source code is no longer available, machines not binary compatible to the ones supported are not able to run S or S-PLUS. In those cases, the R system offers an alternative. We will get to the details in Chapter 15 (page 417).

S-PLUS has minimum requirement specifications regarding main memory and hard disk size. The requirements depend on the particular version and operating system. For a full installation, 300MB will be sufficient (including PDF manuals, etc.). For reasonable performance, 128MB of main memory (RAM) is a lower limit (per simultaneous user). Large data sets or intensive computations will require more memory.

Table 1.1. S-PLUS system requirements.

Operating systems (hardware)	System requirements (recommended minimum)
Windows 2000, Windows XP Home Edition, Windows XP Professional Edition, Windows 2003 Server (on Intel platforms)	Pentium III, 512MB RAM, 450MB disk space
Solaris 2.8, 2.9 (SPARC processors)	160 MB RAM, 340 MB hard disk space
Red Hat Enterprise Linux WS 3.0, kernel 2.4.21, gcc version 3.2.3, glibc version 2.3.2 (Intel/AMD x86, 32 and 64 bit)	160 MB RAM, 340 MB hard disk space
AIX 5.1, 5.2 (POWER processors, IBM)	160 MB RAM, 340 MB hard disk space
HP-UX 11.0, 11i (PA-RISC processor, HP)	360 MB RAM, 340 MB hard disk space

The systems listed here are the officially supported systems for version 7.0. S-PLUS may work other systems as well (for example, other Linux distributions). See the Insightful web page, <http://www.insightful.com/>, for further information.

The minimum RAM requirements for Solaris, Linux, AIX, and HP-UX can be reduced to about 60 MB if S-PLUS is used from the command line and the Java GUI is not used.

As a side note, S-PLUS consumes and releases memory dynamically during a session, depending on the needs. Therefore, it does not run out of memory until there is no more main memory and swap space available. If S-PLUS runs out of main memory (RAM), the operating system assigns virtual memory (i.e., hard disk space) as a substitute. As this slows down the execution time dramatically, the machine should be equipped with a reasonable amount of memory. For improving performance, main memory is the first speed-up factor. If you are not satisfied with the performance, watch for permanent hard disk access while executing commands, or use a monitoring tool (like “top” under UNIX or the “Task Manager” under Windows) to track swapping activity.

Differences Between Versions

S-PLUS has some differences in its implementation between the UNIX and the Windows versions. The most visible difference is the user interface. S-PLUS for Windows has a typical Windows-like menu-based user interface. Many graphical functions can be started using the toolbar interface. Graphics are editable using point-and-click. Data and graphics transfer from and to other Windows applications (such as Word and Powerpoint) is smoothly

integrated. Graphs can be put into Word or Powerpoint documents by using copy and paste, and if a large number of graphs is generated, the Powerpoint Wizard creates a Powerpoint presentation of all S-PLUS graphs with a few mouse clicks.

S-PLUS for UNIX has a graphical user interface, too, but it differs slightly from the look and feel under Windows (due to its implementation in Java). It also lacks the toolbars, and graphs cannot be edited.

In essence, it is straightforward to generate S-PLUS graphs on Windows and include them in Windows document, and the same is true for S-PLUS graphs on Unix and inclusion in documents that reside on the Unix machine. It requires a little extra effort if S-PLUS runs on one system and the documents are edited on the other system. Unless you create a larger number of graphs it should not be too much extra work though. We will cover graph generation and inclusion in documents later on in more detail.

Under S-PLUS for Windows, the standard graphics device is a so-called graph sheet. The graph sheet can hold multiple pages. If some S-PLUS code generates multiple graphs, they are shown on different pages and the user can switch between them by mouse clicks. S-PLUS for UNIX does not have the graph sheet device. If a new graph page is created on the motif device, the page is cleared and the graph lost. The newly introduced Java device has that capability of holding multiple pages, too.

S-PLUS for UNIX allows easy integration of C and Fortran code by basically automatic compilation of a “shared object file” that is linked to the system at runtime. S-PLUS for Windows allows one to dynamically link libraries (DLLs) with some more manual effort.

S-PLUS for Windows comes in two editions: “Professional” and “Standard.” The standard version does not provide command line access.

S-PLUS for UNIX has a special release called “Application Server.” This new implementation allows access to S-PLUS via the Java interface (including the command line accessible from a Java interface that looks similar to the original command line). The application server allows users to access S-PLUS using a Web browser.

There are more differences, but most of them are not directly visible to the user. With very few exceptions, S-PLUS code that runs under S-PLUS for Windows will also run under S-PLUS for UNIX, and vice versa.

1.3 Notational Conventions

Before starting let us introduce a few notational conventions used throughout the book. To begin with, you must be aware that when running S-PLUS, you will be asked for a new command with the greater than sign: `>`. If a single command extends over one line of input, S-PLUS will determine that the command is incomplete and ask for completion. The prompt changes

to the plus sign (+) and you can enter the remainder of the command. A preview of this is shown below.

```
> 3+(4*2
+ )
```

We mostly omit the continuation prompt, as the code is easier to read and so that the + sign is not accidentally typed in as a “plus” operator. For longer program codes, we omit the prompt, too, as we do not assume interactive input but the use of an editor instead. S-PLUS objects or expressions are written in a special font as with the example `print`.

There are a few examples of commands to either the UNIX or DOS shells. For these examples, no prompt is used.

When presenting commands, we sometimes include descriptive text. The descriptive text is written in S-PLUS syntax for comments. Anything after the number sign # until the end of the line is treated as a comment and not interpreted as a command.

A summary of these conventions is found in Table 1.2.

Table 1.2. Notational conventions.

Convention	Explanation
>	S-PLUS prompt
+	Command has continued onto next line (displayed by S-PLUS but omitted in this monograph)
<code>Commands</code>	S-PLUS commands are shown in typewriter font
No prompt	For calls to the UNIX or DOS shell
#	Comment symbol indicating start of a comment
<i>placeholders</i>	Need to be replaced by an appropriate expression, such as <i>filename</i> , which needs to be replaced by a valid file name
CTRL-C	Press and hold down the CTRL key and press C
Menu	Menu entries and buttons are referred to in this font
Note	Notes point out something important, such as an example, an application, or an exception. The end of a note is marked by the symbol ◁

Finally, if S-PLUS prints an object with dimensions – a vector or a matrix, for example – each row is preceded by the index of the following element, as in the following example where [1] indicates that the first element, 6, has the index 1.

```
> 6:9
[1] 6 7 8 9
```

We mostly omit these indices to facilitate reading the output.

2

Graphical User Interface

2.1 Introduction

S-PLUS has an excellent graphical user interface (GUI) that allows the user to access its functionalities by pointing and clicking with the mouse. This convenient way of working with S-PLUS was previously only available for the Windows platform but has now been added for the UNIX/Linux platform as well.

The advantage of the GUI for the novice of S-PLUS is that you don't have to know the syntax of S-PLUS to get started. All you need is a little familiarity with typical Windows software and a data set in some sort of standard format.

The descriptions in this chapter are, for the most part, based on S-PLUS for Windows. Users of UNIX/Linux will benefit from this chapter, despite it being aimed at the Windows user of S-PLUS, as we will show typical applications and examples and point out major differences between the two operating systems. The chapter will end with a section dedicated solely to the GUI available with UNIX/Linux.

The approach we will take in this chapter is to quickly introduce the S-PLUS system design under Windows, show the briefest of explanations of how it functions, and finish by describing in detail the various tasks that will be needed to complete a data analysis, from data input to printing and saving the results. It is by no means intended to be an extensive or exhaustive exploration of the GUI but merely a way of familiarizing you

with its structure, where to find things, and, most importantly, where and what to try for more detailed options.

2.2 System Overview

When you open S-PLUS under Windows by double-clicking on the S-PLUS icon, you are greeted by a screen layout as shown in Figure 2.1. This may vary slightly according to the version of S-PLUS you are using. The main elements that are visible include the Object Explorer, the Commands window, the menus, and the toolbar. Optionally, a graphics window can be opened. A short description of each of these components is given in the next several subsections.¹




Figure 2.1. The S-PLUS screen and its components: the Object Explorer, the Commands window, and the toolbars.

The general layout of the S-PLUS system is similar to that of many popular Windows systems in that it has pull-down menus at the top and toolbars just below the menus. To use such a system it is useful to be a little familiar with basic point-and-click operations and how to use a mouse.

¹According to the version of S-PLUS being used, some components might not be present or look slightly different.

For those who are not that comfortable with window- and icon-based software, the following subsections provide a crash introduction to the essentials. The pull-down menus across the top are used to group categories of commands or options to be set. Under the **File** menu, we find actions relating to files (**Open, Close, Import Data, Save**) as well as to exiting the system. The toolbars below the menus contain buttons that are convenient shortcuts to commands found through the layers in the menus. Some of the toolbar buttons (e.g., **2D Plots**) open a palette containing a myriad of options to complete your task.

An online help facility is available through the main menu. The typical approach to handling the help system (i.e., search by index or topic) has been used. The big advantage of the help system in S-PLUS is that it includes online manuals. We encourage liberal use of the help system.

2.2.1 Using a Mouse

Using a mouse efficiently is important to get the most out of the system. Clicking once on the left mouse button is usually used to highlight an item in a list (e.g., a file out of a list of files) or to select a menu heading or button from a toolbar. You will not always be able to guess at the function of a toolbar button merely by looking at the icon, but if you are at a loss, simply position the mouse over the button in question and a short text description will appear below it. Double-clicking on the left mouse button is used to select and execute. Examples include double-clicking on a file name to select it and start the function, or on a part of a graph to select and edit it. Clicking once with the right mouse button opens a context menu that changes depending on the item selected.

2.2.2 Object Explorer

The Object Explorer is used to get an overview of what is available on the system, including data, functions, and graphs. It operates in much the same way as the Windows Explorer in that there is a tree-like structure in the left pane and the details are displayed in the right pane. The Object Explorer can be opened either from the menu or from a toolbar button. Not only can it show what objects exist, but with data, for example, the Object Explorer is used to view (browse) it and even edit it. Double-click on the data you want to view or edit, and a spreadsheet will open containing the selected data. Once the spreadsheet is open, the data can be edited.

2.2.3 Commands Window

The Commands window is actually the heart of the S-PLUS system. Every command that is performed via menus and buttons can be issued as a