

Covers gnuplot version 5



gnuplot IN ACTION

Understanding data with graphs

SECOND EDITION

Philipp K. Janert

 MANNING

Praise for the First Edition

Knee-deep in data? This is your guidebook to exploring it with gnuplot.

—Austin King
Mozilla

Sparkles with insight about visualization, image perception, and data exploration.

—Richard B. Kreckel
GiNaC.de

Incredibly useful for beginners—indispensable for advanced users.

—Mark Pruett
Dominion

Bridges the gap between gnuplot's reference manual and real-world problems.

—Mitchell Johnson
Border Stylo

A Swiss Army knife for plotting data.

—Nishanth Sastry
University of Cambridge / IBM

Plain and simple: if you use Gnuplot and would like to understand it better, this book is for you. If you are looking for an excellent plotting tool—one that is highly configurable and can easily handle millions of data points, then download Gnuplot and get this book.

—Amazon reviewer

Gnuplot in Action

SECOND EDITION

PHILIPP K. JANERT



MANNING
SHELTER ISLAND

For online information and ordering of this and other Manning books, please visit www.manning.com. The publisher offers discounts on this book when ordered in quantity. For more information, please contact


Special Sales Department
Manning Publications Co.
20 Baldwin Road
PO Box 761
Shelter Island, NY 11964
Email: orders@manning.com

©2016 by Manning Publications Co. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by means electronic, mechanical, photocopying, or otherwise, without prior written permission of the publisher.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in the book, and Manning Publications was aware of a trademark claim, the designations have been printed in initial caps or all caps.

© Recognizing the importance of preserving what has been written, it is Manning's policy to have the books we publish printed on acid-free paper, and we exert our best efforts to that end. Recognizing also our responsibility to conserve the resources of our planet, Manning books are printed on paper that is at least 15 percent recycled and processed without the use of elemental chlorine.

 Manning Publications Co.
20 Baldwin Road
PO Box 761
Shelter Island, NY 11964

Development editors: Marina Michaels
Technical development editor: Ravishankar Rajagopalan
Copyeditor: Tiffany Taylor
Proofreader: Toma Mulligan
Technical proofreader: Clark Gaylord
Typesetter: Dottie Marsico
Cover designer: Marija Tudor

ISBN 9781633430181

Printed in the United States of America

1 2 3 4 5 6 7 8 9 10 – EBM – 21 20 19 18 17 16

The purpose of computing is insight, not numbers.

—R. W. Hamming

The purpose of computing is insight, not pictures.

—L. N. Trefethen

brief contents

PART 1	GETTING STARTED.....	1
1	■ Prelude: understanding data with gnuplot	3
2	■ Tutorial: essential gnuplot	16
3	■ The heart of the matter: the plot command	31
PART 2	CREATING GRAPHS	53
4	■ Managing data sets and files	55
5	■ Practical matters: strings, loops, and history	78
6	■ A catalog of styles	100
7	■ Decorations: labels, arrows, and explanations	125
8	■ All about axes	146
PART 3	MASTERING TECHNICALITIES.....	179
9	■ Color, style, and appearance	181
10	■ Terminals and output formats	209
11	■ Automation, scripting, and animation	236
12	■ Beyond the defaults: workflow and styles	262
PART 4	UNDERSTANDING DATA.....	287
13	■ Basic techniques of graphical analysis	289
14	■ Topics in graphical analysis	314
15	■ Coda: understanding data with graphs	344

contents

preface xvii
acknowledgments xix
about this book xx

PART 1 GETTING STARTED1

1 *Prelude: understanding data with gnuplot* 3

1.1 A busy weekend 4

Planning a marathon 4 ▪ *Determining the future* 6

1.2 What is graphical analysis? 9

Why graphical analysis? 11 ▪ *Limitations of graphical analysis* 11

1.3 What is gnuplot? 11

Gnuplot isn't GNU 12 ▪ *Why gnuplot?* 12 ▪ *Limitations* 13
Gnuplot 5: the best gnuplot there ever was! 13

1.4 Summary 15

2 *Tutorial: essential gnuplot* 16

2.1 Simple plots 17

Invoking gnuplot and first plots 17 ▪ *Plotting data from a file* 20 ▪ *Abbreviations and defaults* 24

- 2.2 Saving commands and exporting graphics 25
 - Saving and loading commands* 25 ▪ *Exporting graphs* 26
- 2.3 Managing options with set and show 29
- 2.4 Getting help 30
- 2.5 Summary 30

3 *The heart of the matter: the plot command* 31

- 3.1 Plotting functions and data 32
 - Plotting functions* 32 ▪ *Plotting data* 33
- 3.2 Math with gnuplot 34
 - Mathematical expressions* 34 ▪ *Built-in functions* 36
 - User-defined variables and functions* 37 ▪ *Mathematically undefined values and NaN (not a number)* 39
- 3.3 Data transformations 39
 - Simple data transformations* 40
- 3.4 Logarithmic plots 42
- 3.5 Smooth interpolation and approximation 45
 - Interpolation curves* 46 ▪ *Point distributions* 49
 - Deduping repeated entries* 51
- 3.6 Summary 51

PART 2 CREATING GRAPHS53

4 *Managing data sets and files* 55

- 4.1 Quickstart: the standard data-file format 56
 - Comments and header lines* 58 ▪ *Selecting columns* 58
- 4.2 Managing structured data sets 59
 - Multiple data sets per file: index* 59 ▪ *Records spanning multiple lines: the every directive* 62
- 4.3 File format options in detail 63
 - Number formats* 63 ▪ *Comments* 63 ▪ *Field separator* 64 ▪ *Missing values* 65 ▪ *Strings in data files* 66
- 4.4 Accessing columns and pseudocolumns 67
 - Accessing columns by position or name* 68 ▪ *Pseudocolumns* 69
 - Column-access functions* 69

- 4.5 Pseudofiles 70
 - Reading data from standard input* 71 ▪ *Heredocs* 72
 - Reading data from a subprocess* 74 ▪ *Writing to a pipe* 74
 - Generating data* 75
- 4.6 Metadata in data files 75
- 4.7 Other file formats 76
- 4.8 Summary 77

5 *Practical matters: strings, loops, and history* 78

- 5.1 Strings 79
 - Quotes* 80 ▪ *String operations* 80 ▪ *Worked example: plotting the Unix password file* 82
- 5.2 String expressions and string macros 85
 - String expressions in commands* 86 ▪ *Executing a string with eval* 86 ▪ *String macros inside commands* 87
- 5.3 Generating textual output 88
 - The print and set print commands* 88 ▪ *The set table command and the with table style* 89 ▪ *Reading and writing heredocs* 90
- 5.4 Simplifying work with inline loops 90
 - Loops over numbers* 90 ▪ *Loops over strings* 91
 - Summary of inline loops* 92
- 5.5 Gnuplot's internal variables 93
- 5.6 Inspecting file contents with the stats command 94
 - The stats command and internal variables* 95
 - Further options for the stats command* 96
- 5.7 Command history 96
 - Redrawing a graph* 97 ▪ *The general history feature* 97
 - Restoring session defaults* 98
- 5.8 Summary 99

6 *A catalog of styles* 100

- 6.1 Why use different plot styles? 101
- 6.2 Styles and aspects 101
 - Choosing styles inline through with* 101 ▪ *The default sequence* 102 ▪ *Customizing graph elements* 103

- 6.3 A catalog of plotting styles 105
 - Core styles: lines and points* 106 ▪ *Indicating uncertainty: styles with error bars or ranges* 108 ▪ *Styles with steps and boxes* 112
 - Filled styles* 114 ▪ *Beyond lines and points: multivariate visualization* 115
- 6.4 Putting it together 120
- 6.5 Other styles 123
- 6.6 Summary 123

7 *Decorations: labels, arrows, and explanations* 125

- 7.1 Quick start: minimal context for data 126
- 7.2 Understanding layers and locations 128
 - Locations* 128 ▪ *Layers* 129
- 7.3 Additional graph elements: decorations 130
 - Common conventions* 130 ▪ *Arrows* 131 ▪ *Text labels* 134
 - Shapes or objects* 136
- 7.4 The graph's legend or key 137
 - Turning the key on and off* 138 ▪ *Placement* 138 ▪ *Layout* 139
 - Appearance* 139 ▪ *Explanations* 140 ▪ *Default settings* 143
- 7.5 Worked example: features of a spectrum 144
- 7.6 Summary 145

8 *All about axes* 146

- 8.1 Multiple axes 147
 - Terminology* 147 ▪ *Plotting with two coordinate systems* 148
 - Linking axes* 150
- 8.2 Selecting plot ranges 153
 - What you need to know for interactive work* 154
 - What you might want to know for batch processing* 155
- 8.3 Tic marks 156
 - Overview and common conventions* 156 ▪ *Tic mark appearance and placement* 157 ▪ *Tic labels* 159 ▪ *Tic mark location and frequency* 163 ▪ *Reading tic labels from file* 165 ▪ *Grid and zero axis* 166
- 8.4 Special case: time series 168
 - Turning numbers into names: months and weekdays* 169
 - General time series: the gory details* 170 ▪ *Beyond tic labels: processing date/time information* 176
- 8.5 Summary 178

PART 3 MASTERING TECHNICALITIES179

9 *Color, style, and appearance* 181

- 9.1 Color 181
 - Explicit colors* 182 ▪ *Alpha shading and transparency* 185
 - Selecting a color through indexed lookup* 187 ▪ *Mapping a value into a continuous gradient* 188 ▪ *Using data-dependent colors* 189 ▪ *The built-in color sequences* 193 ▪ *Tips and tricks* 194
- 9.2 Lines and points 195
 - Point types and shapes* 195 ▪ *Dash pattern* 196
- 9.3 Customizing color, dash, and point sequences 198
 - Customizing line types* 199 ▪ *Special line types* 200
- 9.4 Global styles 200
 - Data and function styles* 200 ▪ *Line styles* 201 ▪ *Arrow styles* 202 ▪ *Fill styles* 202 ▪ *Other global styles* 203
- 9.5 Overall appearance: aspect ratio and borders 203
 - Size and aspect ratio* 204 ▪ *Borders* 205 ▪ *Margins* 207
 - Internal variables* 207
- 9.6 Summary 208

10 *Terminals and output formats* 209

- 10.1 The terminal abstraction 210
 - Historical digression* 210 ▪ *The terminal workflow* 210
 - Terminal capabilities and the test command* 212
- 10.2 Font selection and enhanced text mode 214
 - Font selection* 214 ▪ *Font resolution* 215 ▪ *Enhanced text mode* 217 ▪ *Worked example* 219
- 10.3 Generating PNG and PDF with cairo-based terminals 220
- 10.4 Using gnuplot with LaTeX 221
 - Including a graph in a LaTeX document* 222 ▪ *Using the cairolatex terminal* 224 ▪ *Letting LaTeX generate the graph* 228
- 10.5 Scalable graphics for the Web with SVG and HTML5 229
 - The svg terminal* 229 ▪ *The canvas terminal* 230
- 10.6 Interactive terminals 232
 - Common options* 232 ▪ *The wxt and qt terminals* 233
 - The aqua terminal* 233 ▪ *The windows terminal* 233

10.7 Other terminals 233

10.8 Summary 234

11 *Automation, scripting, and animation* 236

11.1 Loops and conditionals 237

Worked example: making graph paper 239 ▪ *Worked examples: iterating over files* 240 ▪ *Worked examples: Taylor series and Newton's method* 241

11.2 Command files 243

Scripts as subroutines 244 ▪ *Worked example: export script* 246

11.3 Batch processing 247

Using gnuplot in shell pipelines 249

11.4 Calling gnuplot from other programs 250

Worked example: calling gnuplot from Perl 250 ▪ *Worked example: calling gnuplot from Python* 251 ▪ *Helpful hints* 252

11.5 Animations 254

Introducing a delay 254 ▪ *Waiting for a user event* 255
Further examples 255

11.6 Case study: continuously monitoring a live data stream 256

Using gnuplot to monitor a file 256 ▪ *Using a driver to monitor arbitrary data sources* 259

11.7 Summary 261

12 *Beyond the defaults: workflow and styles* 262

12.1 The standard interactive workflow 263

Extracting specifics from command files 263 ▪ *Extending the command set* 264 ▪ *Session variables, loops, and macros* 266

12.2 Using external editors and viewers 267

12.3 Invoking shell commands from gnuplot 268

Worked example: plotting each file in a directory 269

12.4 Hotkeys and mousing 269

Default hotkeys 270 ▪ *Mousing* 270 ▪ *Custom hotkeys* 271
Capturing mouse events 273 ▪ *Case study: placing arrows and labels with the mouse* 274

- 12.5 Startup configurations and initialization 275
 - Startup and initialization files* 276
 - Environment variables* 277
 - Gnuplot command-line flags* 278
- 12.6 Stylesheets 279
 - Worked example: stylesheets* 279
- 12.7 Summary 285

PART 4 UNDERSTANDING DATA..... 287

13 *Basic techniques of graphical analysis* 289

- 13.1 Representing relationships 290
 - Scatter plots* 290
 - Highlighting trends* 292
- 13.2 Logarithmic plots 295
 - Large variations in data* 295
 - Power-law behavior* 298
- 13.3 Point distributions 299
 - Summary statistics and box plots* 300
 - Jitter plots and histograms* 300
 - Kernel density estimates and rug plots* 302
 - Cumulative distribution functions* 303
- 13.4 Ranked data 304
- 13.5 Pie charts 307
- 13.6 Organizational issues 308
 - The lifecycle of a graph* 308
 - Input data files* 309
 - Output files* 310
- 13.7 Presentation graphics 311
- 13.8 Summary 313

14 *Topics in graphical analysis* 314

- 14.1 Techniques for time-series plots 315
 - Plotting an Apache web server log* 315
 - Smoothing and differencing* 316
 - Monitoring and control charts* 319
 - Changing composition and stacked curves* 325
- 14.2 Graphical techniques for multivariate data sets 328
 - Introduction* 329
 - Distribution of values by attribute* 329
 - Distribution by level* 331
 - Scatter-plot matrix* 333
 - Parallel-coordinates plot* 335

- 14.3 Visual perception 338
 Banking 338 ▪ *Judging lengths and distances* 340
 Plot ranges and whether to always include zero 342
- 14.4 Summary 343

15 *Coda: understanding data with graphs* 344

- appendix A* *Obtaining, building, and installing gnuplot* 346
- appendix B* *Resources* 352
- appendix C* *Surface and contour plots* available in the e-book only
- appendix D* *Palettes and false-color plots* available in the e-book only
- appendix E* *Special plots* available in the e-book only
- appendix F* *Higher math* available in the e-book only
- index* 355

preface

On New Year's Day, 2015, the gnuplot development team released version 5.0—the first major new gnuplot release in over 10 years! I decided to take this opportunity to bring *Gnuplot in Action* up to date and to cover all the new features gnuplot has acquired since the first edition of this book was written (in 2007).

It quickly became apparent it wouldn't be sufficient to just add a couple of chapters explaining the new features. In fact, the book you're reading now has been almost entirely rewritten from scratch. Most of the material from the first edition has been retained, but it's been heavily rearranged to accommodate the addition of new topics and to reflect the changes in my own understanding and priorities.

Gnuplot 5 is largely backward compatible with previous versions, and hence most of the first edition remains valid. At the same time, new features have been added to all parts of gnuplot, either to add new functionality or to streamline and improve the existing usage. Although many of the new features are small by themselves, when taken together, their cumulative effect leads to a significantly different, more sophisticated user experience.

In the process, the book's page count has increased substantially from the first edition. To keep the physical dimensions of the printed book in check without having to sacrifice important and useful material, some topics of a more specialized nature have been relegated to the electronic (e-book) version. Access to the e-book is included in the purchase of a print copy of the book.

Today, gnuplot is still going strong. Despite increased competition, gnuplot's two most attractive features are still largely unmet by other tools:

- The ability to explore data graphically, with an absolute minimum of effort, protocol, overhead, or boilerplate

- The ability to create immaculate, very high-quality graphs, with text labels and other decorations, for presentation purposes

What's new is that gnuplot has arrived in the 21st century. Color is now the standard, font handling is up to date, and the graphing backend makes use of all contemporary technologies to create the best-looking graphs possible.

In the first edition, I wrote that gnuplot was “an indispensable part of my toolbox: one of the handful of programs I can't do without.” Several years on, this is still true.

acknowledgments

During the preparation of this book, I enjoyed conversations and correspondence with Austin King, Richard Kreckel, Ethan Merritt, Dawid Weiss, Bastian Märkisch, Daniel Sebald, Petr Mikulik, Chris Mague, Luis Moux-Dominguez, and Lee Phillips. Christoph Bersch, Zoltán Vörös, and Clark Gaylord read drafts of this book and provided many detailed suggestions; Mojca Miklavc answered several specific questions with meticulous care. Others who read the draft manuscript include Ryan Balfanz, Martin Beer, Andrew Bovill, Vitaly Bragilevsky, Anthony Cramp, Wolfgang Ecker-Lala, Wesley R. Elsberry, Nitin Gode, David Kerns, Pavol Kral, Mathew Peet, Ravishankar Rajagopalan, Karl-Friedrich Ratzsch, Jonathan Rioux, Mike Shepard, and Arthur Zubarev.

I would also like to acknowledge the tremendous impact that Wikipedia has had on the way I work. When I prepared the first edition, obtaining even basic information on topics such as color spaces, Bézier curves, and the Mandelbrot set was a real challenge—difficult, time consuming, and not always successful. For all its faults and deficiencies, Wikipedia has made it tremendously much easier to obtain at least an initial introduction (and often quite a bit more) to an incredibly wide range of topics. It is a stunning achievement.

Finally, I want to thank the people at Manning who made this book possible: publisher Marjan Bace and everyone on the editorial and production teams, including Mary Piergies, Marina Michaels, Kevin Sullivan, Tiffany Taylor, Dottie Marsico, and many others who worked behind the scenes.

about this book

This book is intended to be a comprehensive introduction to gnuplot: from the basics to the power features and beyond. In addition to providing a tutorial on gnuplot itself, it demonstrates how to apply and use gnuplot to extract insight from data.

The gnuplot program has always had complete and detailed reference documentation, but what was often missing was a continuous presentation that tied all the different bits and pieces of gnuplot together and demonstrated how to use them to achieve certain tasks. This book attempts to fill that gap. It should also serve as a handy reference for more advanced gnuplot users and as an introduction to graphical ways of knowledge discovery.

And finally, this book tries to show you how to use gnuplot to achieve some surprisingly nifty effects that will make everyone say, “How did you do that?”

Contents of this book

This book is divided into four parts. Part 1 consists of chapters 1 through 3 and is intended as a tutorial introduction to get you started with gnuplot. These three chapters cover all the truly essential material so that by the end of chapter 3, you should be able to handle most basic plotting tasks in gnuplot.

Whereas part 1 only skims the surface, part 2 goes into depth. First, chapters 4 and 5 lay more groundwork by talking about the ins and outs of file formats, string handling, and other practical matters. Then, chapters 6 through 8 discuss the various ways to change the appearance of a plot: using different plotting styles; adding labels, arrows, or other decorations; and changing the axes and their subdivisions. These chapters cover the tactical aspects of working with gnuplot in detail.

Part 3 turns its attention away from individual graphs and addresses a variety of more technical aspects. First, in chapter 9, you'll learn more about color specification, point and line types, and other relatively low-level graph elements. Chapter 10 explains how to export plots to common graphics file formats. Finally, chapters 11 and 12 address ways to improve the overall workflow through scripting and configuration changes.

In the last part, I'll mostly take gnuplot's features for granted and concentrate on the things you can do with them. Chapter 13 presents various fundamental types of graphs and explains when and how to use them. Chapter 14 is more advanced and offers solutions to some recurring topics in graphical analysis, before we end the book with a reminder of what it's all about in chapter 15.

The book has several appendixes. Appendix A explains how to obtain, build, and install gnuplot. Appendix B provides pointers to some relevant resources.

Finally, some topics of a more specialized character have been relegated to a set of supplemental appendixes: appendixes C and D discuss three-dimensional surface plots and false-color plots (heatmaps), appendix E treats some special types of graphs, and appendix F covers more mathematical topics. To reduce the physical dimensions of the print book, these four appendixes are only available in the electronic (e-book) version of this book. The purchase of a hard copy includes access to the e-book as well—you can find instructions in the front of the print book.

TIP Appendixes C through F are only available in the e-book version of this book, which is included with the purchase of the hard-copy version. Check the front of the print book for instructions on how to obtain the e-book.

How to read this book

This book was written as if readers were going to read it sequentially, cover to cover. New material is presented in order, with later chapters relying only on topics introduced earlier and avoiding forward references as much as possible. I realize that this is not a realistic picture and that the need for technical information tends to arise in a much more disjointed manner. In this spirit, I offer a few different “trail maps” to the material presented here:

- If you're new to gnuplot, begin with chapters 2 and 3 and then dive into chapters 4–8 as required to pick up the skills you need to complete whatever task you want to accomplish.
- If you're comfortable creating day-to-day graphs with gnuplot, then the material in chapters 9–12 should help you achieve greater efficiency in your work and fine-tune the results.
- If you've been using gnuplot for a long time already, then make sure you read up on the new features in gnuplot 5. Chapters 5, 9, and 10, as well as parts of chapters 11 and 12 will probably be of the most immediate interest to you.

- If you're new to graphical analysis, you may want to begin with chapter 13 to learn some of the basic methods and concepts.

Finally, keep in mind that some interesting and useful material is only available in the e-book. Three-dimensional surface and contour plots are discussed in appendix C. False-color plots (heatmaps) are treated in appendix D, together with guidelines for how to construct effective color gradients for data visualization. Appendix E explains how to combine individual graphs into composites and also discusses some other specialized types of graphs. Appendix F treats topics of a more mathematical nature.

Whom this book is for

This book is intended for anyone who wants to plot and visualize data, either to explore data sets graphically, or to create attractive, high-quality graphs for presentation and publication purposes. I had two kinds of people in mind when writing this book—those who already know gnuplot, and those who don't:

- If you already know gnuplot, I hope you'll still find it a useful reference, in particular in regard to some of the more advanced topics later in the book. I've tried to provide exactly the big-picture explanations and examples that have always been missing from the standard gnuplot reference documentation.
- If you're new to gnuplot, I think you'll find it easy enough to pick up—in fact, I can promise you that by the end of chapter 2, you'll be productive with gnuplot; and by the end of chapter 3, you'll be well equipped for most day-to-day data graphing tasks that may come your way.

This book doesn't require a strong background in mathematical methods or any in statistics, but I occasionally do expect you to have at least a fleeting familiarity with simple programming concepts. A few sections naturally require some special preliminaries (for instance, some of the discussions in chapter 10 require knowledge of LaTeX, and some sections in chapter 11 use Perl or Python code), but you can safely skip those sections if their material doesn't apply to you.

Conventions

I spell the name of the program in all lowercase (gnuplot), except at the beginning of a sentence, when I capitalize it normally. This is in accordance with the usage recommended in the gnuplot FAQ.

The gnuplot documentation is extensive, and I refer to it occasionally for additional details on topics covered only briefly or not at all here. Traditionally, the gnuplot documentation has been called the online help or online documentation, owing to the fact that it's available online during a gnuplot session. But since the advent of the internet, the word *online* seems to suggest network connectivity—falsely, in this context. To avoid confusion, I'll always refer to it as the *standard gnuplot reference documentation*.

Code examples

Gnuplot commands are shown using a monospace font, like this: `plot sin(x)`. Gnuplot commands can be entered at the gnuplot command prompt as shown in the text; the prompt itself has been suppressed to save space.

Single command lines can be long; to make them fit on a page, I occasionally had to break them across multiple lines. If so, a gray arrow (➡) has been placed at the beginning of the next line, to indicate that it is the continuation of the previous one:

```
plot "data" using 1:2 smooth csplines title "data" with lines,  
➡  sin(x) title "model"
```

The break in the original line isn't indicated separately. When using gnuplot in an interactive session, your terminal program should automatically wrap a line that's too long. Alternatively, you can break lines by escaping the newline with a backslash as usual. This is useful in command files for batch processing (and you'll see some examples in chapter 12 in the context of string macros).

Some code snippets are only intended to demonstrate the syntax and don't have a graph associated with them. In this case, I use the generic name "data" as a placeholder for the actual filename. No file named *data* exists in the downloads (in the same way that no key named *any* can be found on a computer keyboard). It's just a generic placeholder.

Occasionally, I show Unix commands that need to be entered in a Unix shell; to emphasize that these aren't gnuplot commands, I prefix them with a generic shell prompt, like this: `shell>`. Similarly, Python commands to be entered in a Python session are prefixed with `python>>>`.

Downloads

The code for all numbered listings is available for download from www.manning.com/books/gnuplot-in-action-second-edition, and so are the data sets. The only exception to this are publicly available data sets: for these, I provide the URL where they can be found.

Gnuplot searches for data files in the current directory, so the easiest way to run the supplied command files is as follows:

- 1 Change into the data directory of the downloaded bundle.
- 2 Start gnuplot.
- 3 Issue `plot` commands at the gnuplot prompt the way they're shown in the text (for example, `plot "marathon" using 1:2`), or give the full pathname to the gnuplot command file that you wish to run (for example, `load "../gnuplot/shapes.gp"`).

Command synopses

Gnuplot has a large number of options, and keeping all of them and their sub-options and optional parameters straight is a major theme running through this book. Frequently, I'll display all available options to a command in a command synopsis before discussing the options in detail. To distinguish a synopsis of available options from actual gnuplot code, a synopsis uses an italic font, like so:

```
set datafile commentschar [{"str:chars"}]
```

Within these summaries, I use a few syntactic conventions. My intent here is to stay close to the usage familiar from the standard gnuplot reference documentation, but also to follow more general conventions (such as those used for Unix man pages):

```
[ ... ]           ← Square brackets for optional parts
[ | ]           ← Vertical bars to separate alternatives
{ ... }         ← Curly braces for user-supplied input
```

For parameters supplied by the user, it's often not clear from the context what kind of information the command expects: is it a string or a number? If it's a number, is it a value selected from a fixed range of integers or a numerical factor? And so on. I've tried to clarify this situation by prefixing each user-supplied input parameter with a type indicator, terminated by a colon. I summarize the prefixes and their meanings in table 1.

Table 1 Type indicators for user-supplied parameters

Prefix	Description
<i>str:</i>	A string
<i>int:</i>	An integer number
<i>flt:</i>	A floating-point number
<i>idx:</i>	An integer number, which is interpreted as a selection from a fixed range of values
<i>clr:</i>	A color specification—for example, <code>rgbcolor "red"</code> or <code>rgb "#FFFF00"</code>
<i>pos:</i>	A pair of comma-separated coordinates, optionally containing coordinate system specifiers—for example, <code>0,0</code> or <code>first 1.1, screen 0.9</code>
<i>enum:</i>	A gnuplot keyword as unquoted string

Abbreviations

Many gnuplot commands have abbreviated forms, which I use frequently. The essential `plot` command, in particular, takes a large number of keyword directives, which I usually abbreviate to save space and keystrokes. I strongly recommend that you quickly become familiar with these shorthands and use them yourself. Table 2 lists both the abbreviated and the full forms. The `plot` command also understands a large number of appearance options (controlling aspects such as line width, style, and color), which are generally also abbreviated. A comprehensive summary of appearance options, together with their shorthands, can be found in table 6.1.

Table 2 Abbreviations for frequently used directives to the `plot` command

Abbreviation	Full
<code>i</code>	index
<code>ev</code>	every
<code>u</code>	using
<code>s</code>	smooth
<code>s acs</code>	smooth acsplines
<code>s f</code>	smooth frequency
<code>s kdens</code>	smooth kdensity
<code>t</code>	title
<code>w</code>	with
<code>w l</code>	with lines
<code>w linesp</code> or <code>w lp</code>	with linespoints
<code>w p</code>	with points
<code>w vec</code>	with vectors

Table 3 lists three frequently occurring commands that are also usually abbreviated.

Table 3 Abbreviations for frequently occurring commands

Abbreviation	Full
<code>set t</code>	set terminal
<code>set o</code>	set output
<code>set logsc</code>	set logscale

The figures in this book

The graphs in this book were generated with `gnuplot`; some special cases were handled using `pic`. All graphs were originally prepared in color, using my own set of preferred colors instead of one of `gnuplot`'s default color schemes. The color versions of the graphs are used in the electronic (e-book) version of this book. For the print book, I prepared black-and-white versions through the application of an appropriate stylesheet (see chapter 12). A handful of graphs required manual touch-ups in addition to the monochrome stylesheet to yield an optimal appearance.

You'll find the line-type definitions of both the color and the black-and-white stylesheets in table 4. The same colors and dash patterns are discussed in listings 12.7 and 12.9.

In particular in the latter part of the book, I frequently use point types (point shapes) that aren't the default, because the visual appearance of the graphs can often be improved greatly this way. If so, the point type is usually chosen explicitly in the appropriate code examples and listings.

The final version of each figure was generated using the `pdfcairo` terminal, using a (non-default) aspect ratio of $\sqrt{2}$ to 1 and Helvetica as the requested font.

Table 4 Colors and dash patterns used for the color and monochrome figures in this book

Color	Monochrome
<code>set linetype 1 lc rgb '0xee0000'</code>	<code>set linetype 1 lc black dt solid</code>
<code>set linetype 2 lc rgb '0x008b00'</code>	<code>set linetype 2 lc black dt (8,6)</code>
<code>set linetype 3 lc rgb '0x0000cd'</code>	<code>set linetype 3 lc black dt (4,3)</code>
<code>set linetype 4 lc rgb '0xff3fb3'</code>	<code>set linetype 4 lc black dt (3,6)</code>
<code>set linetype 5 lc rgb '0x00cdcd'</code>	<code>set linetype 5 lc black dt (12,5,2,5,2,5)</code>
<code>set linetype 6 lc rgb '0xcd9b1d'</code>	<code>set linetype 6 lc black dt (16,8)</code>
<code>set linetype 7 lc rgb '0x8968cd'</code>	<code>set linetype 7 lc black dt (20,6,2,6)</code>
<code>set linetype 8 lc rgb '0x8b8b83'</code>	<code>set linetype 8 lc black dt (30,10)</code>

Hardware and software requirements

This book describes gnuplot version 5.0 or higher, which was initially released in early 2015. Not all examples in this book will work with earlier gnuplot versions. If you have an earlier version of gnuplot, you should upgrade to a more current version—appendix A tells you how.

I assume you have access to a reasonably modern computer running any flavor of Unix/Linux, a recent release of MS Windows, or Mac OS X. Gnuplot has been ported to many other platforms but is actively supported primarily on the three operating systems just mentioned, and so I concentrate on them in this book.

Reference materials

Command and option references are distributed throughout the book, wherever the material is first introduced. The following pointers are intended to help you find these summaries more easily.

Graphical styles and specifications

Appearance specifiers and line options	Table 6.1	page 104
Graph locations	Figure 7.2	page 128
Graph layers	Figure 7.3	page 129
Explicit color-specification formats	Table 9.1	page 184
Point types	Figure 9.7	page 196
Dash patterns	Table 9.2	page 197

File access

Column-access methods and functions	Table 4.2	page 67
Pseudofiles	Table 4.3	page 71
Metadata in data files	Table 4.4	page 76

String handling and formatting

String functions	Table 5.1	page 81
General conversion specifiers	Table 8.2	page 160
Accuracy specifiers	Table 8.3	page 160
Time-series conversions, sorted alphabetically	Table 8.4	page 172
Time-series conversions, sorted by topic	Table 8.5	page 173
Time functions	Table 8.6	page 176
Enhanced text mode	Table 10.1	page 218

Operators and mathematical functions

Unary operators	Table 3.1	page 35
Binary operators	Table 3.2	page 35
Mathematical functions	Table F.1	page F11
Complex numbers	Table F.2	page F13

Programming constructs

Inline loops	Listing 5.3	page 92
General loops and conditionals	Table 11.1	page 238

About the author

PHILIPP K. JANERT was born and raised in Germany. He obtained a Ph.D. in theoretical physics from the University of Washington in 1997 and has been working in the tech industry ever since, including four years at Amazon.com, where he initiated and led several projects to improve Amazon's order-fulfillment process. He's the author of several books on data analysis and applied math, including the best-selling *Data Analysis with Open Source Tools* (O'Reilly, 2010). He has contributed to CPAN and is an occasional committer on the gnuplot project. Visit his company website at www.principal-value.com.

Author Online

Purchase of *Gnuplot in Action, Second Edition* includes free access to a private web forum run by Manning Publications where you can make comments about the book, ask technical questions, and receive help from the lead author and from other users. To access

the forum and subscribe to it, point your web browser to www.manning.com/books/gnuplot-in-action-second-edition. This page provides information on how to get on the forum once you are registered, what kind of help is available, and the rules of conduct on the forum.

Manning's commitment to our readers is to provide a venue where a meaningful dialog between individual readers and between readers and the author can take place. It is not a commitment to any specific amount of participation on the part of the author, whose contribution to Author Online remains voluntary (and unpaid). We suggest you try asking the author some challenging questions lest his interest stray! The Author Online forum and the archives of previous discussions will be accessible from the publisher's website as long as the book is in print.

About the cover

The figure on the cover of *Gnuplot in Action, Second Edition* is captioned "A peer of France." The title of Peer in France was held by the highest-ranking members of the French nobility. It was an extraordinary honor granted only to a few dukes, counts, and princes of the church. The illustration is taken from a 19th-century edition of Sylvain Maréchal's four-volume compendium of regional dress customs published in France. Each illustration is finely drawn and colored by hand.

The rich variety of Maréchal's collection reminds us vividly of how culturally apart the world's towns and regions were just 200 years ago. Isolated from each other, people spoke different dialects and languages. In the streets or in the countryside, it was easy to identify where they lived and what their trade or station in life was just by their dress.

Dress codes have changed since then, and the diversity by region, so rich at the time, has faded away. It's now hard to tell apart the inhabitants of different continents, let alone different towns or regions. Perhaps we have traded cultural diversity for a more varied personal life—certainly for a more varied and fast-paced technological life.

At a time when it's hard to tell one computer book from another, Manning celebrates the inventiveness and initiative of the computer business with book covers based on the rich diversity of regional life of two centuries ago, brought back to life by Maréchal's pictures.

Part 1

Getting started

Gnuplot is a tool for visualizing data and mathematical functions. The chapters in this first part will give a first introduction to gnuplot and its most important features. Chapter 1 introduces gnuplot and describes the kinds of problems it's designed to solve. Chapter 2 provides a quick tutorial to gnuplot. By the end of this chapter, you'll be able to prepare simple plots with gnuplot and to save and export your work. Chapter 3 takes a detailed look at the all-important `plot` command, which is used to generate most graphs in gnuplot. You'll also learn about inline transformations and built-in smoothing methods.

1

Prelude: understanding data with gnuplot

This chapter covers

- Warmup examples
- What is graphical analysis?
- What is gnuplot?

NOTE TO PRINT BOOK READERS Some material of a more specialized nature is only available in the e-book version of this book. The e-book also shows all the graphs in color. To get your free e-book in PDF, ePub, or Kindle format, go to www.manning.com/books/gnuplot-in-action-second-edition to register your print book.

Gnuplot has long been one of the most popular open source programs for plotting and visualizing data. In this book, I want to show you how to use gnuplot to make plots and graphs of your data: both quick and easy graphs for your own use and highly polished graphs for presentations and publications.

But I also want to show you something else: how to solve data-analysis problems using graphical methods. The art of discovering relationships in data and extracting information from it by visual means is called *graphical analysis*, and I believe gnuplot to be an excellent tool for it.

As a teaser, let's look at some problems and how you might be able to approach them using graphical methods. The graphs here and in the rest of the book (with very few exceptions) have been, of course, generated with gnuplot.

1.1 *A busy weekend*

To get a feeling for the kinds of problems you may be dealing with and for the kinds of solutions gnuplot can help you find, let's look at two examples. Both take place during a long, busy weekend.

1.1.1 *Planning a marathon*

Imagine you're in charge of organizing the local city marathon. There will be more than 2,000 starters, traffic closed around the city, plenty of spectators—and a major Finish Line Festival to celebrate the victors. The big question is: when should the Finish Line crew be ready to deal with the majority of runners? At what point do you expect the big influx of the masses?

You have the results from last year's event. Assuming that the starters haven't improved dramatically over the last year (probably a safe assumption), you do a quick average of the completion times and find that last year's average was 282 minutes. To be on the safe side, you calculate the standard deviation as well, which comes out to about 50 minutes. So you tell your crew to be ready for the big rush starting three and a half hours (210 minutes) after the start, and you feel reasonably well prepared for the event.

So it comes as a surprise when on the big day, plenty of runners start showing up at the finish line after only two hours—a good 90 minutes earlier than the expected onset of the rush. In terms of event management, the number of runners who show up early isn't overwhelming, but it's a bit strange. The next day you wonder: what went wrong?

Let's look at the data to see what you can learn about it. So far, all you know are the mean and the standard deviation.

The mean is convenient: it's easy to calculate, and it summarizes the entire data set in a single number. But in forming the mean, you lost a lot of information. To understand the entire data set, you have to *look* at it. And because you can't understand data by looking at more than 2,000 individual finish times, this means you have to *plot* it.

It will be convenient to group the runners by completion time and to count the number of participants who finished during each five-minute interval. The resulting file might start like this:

```
# Minutes Runners
135      1
140      2
145      4
150      7
155     11
160     13
165     35
170     29
...
```




Figure 1.1 Number of finishers vs. time to complete (in minutes)

Now you plot the number of runners against the completion time (see figure 1.1). It’s immediately obvious where you went wrong: the data is *bimodal*, meaning it has two peaks. There is an early peak at around 180 minutes and a later main peak at 300 minutes.

Actually, this makes sense: a major sporting event such as a city marathon attracts two very different groups of people: athletes, who train and compete throughout the year and are in it to win, and a much larger group of amateurs, who come out once a year for a big event and are mostly there to participate. The problem is that for such data, the mean and standard deviation are obviously bad representations—so much so. that at the time when you expected the big rush (200 minutes), there’s a lull at the finish line!

The take-home message here is that it’s usually not a good idea to rely on summary statistics (such as the mean) for unknown data sets. You *always* should investigate what the data looks like. Once you’ve confirmed the basic shape, you can choose how to summarize your findings best.

And of course, there is always more to learn. In this example, for instance, you see that after about 400 minutes, almost everybody has made it, and you can start winding down the operation. The actual “tail” of the distribution is quite small—surprisingly so. (I would’ve expected to see a greater number of stragglers, but possibly many runners who are *really* slow drop out of the race when they realize they’ll place badly.)

USING GNUPLOT

Let’s look at the gnuplot command that was used to generate figure 1.1. Gnuplot is command-line oriented: after you start gnuplot, it drops you into an interactive command session, and all commands are typed at the interactive gnuplot prompt.

Gnuplot reads data from simple text files, with the data arranged in columns as shown previously. To plot a data file takes only a single command, `plot`, like this:¹

```
plot "marathon" using 1:2 with boxes
```

The `plot` command requires the name of the data file as argument in quotes. By default, gnuplot looks for the data file in the current working directory—normally the directory from which you started gnuplot. The filename provided to the `plot` command may contain path information to refer to a file that doesn't reside in the current directory.

The rest of the command line specifies which columns to use for the plot and in which way to represent the data. The `using 1:2` declaration tells gnuplot to use the first and second columns in the file called `marathon`. The final part of the command, `with boxes`, selects a box style, which is often suitable to display counts of events.

Gnuplot handles most everything else by itself: it sizes the graph and selects the most interesting plot range, it draws the border, and it draws the tic marks and their labels. All these details can be customized, but gnuplot typically does a good job at anticipating what the user wants.

NOTE The little markers along the edge that define the *scale* of the corresponding axis are called *tick marks* (or *tic marks*). The gnuplot standard reference documentation uses the spelling *tic* mark; the relevant commands are called `set xtics`, `set ytics`, and so on. In order to avoid confusion, I use the same spelling (*tic*) throughout this book.

1.1.2 *Determining the future*

The same weekend when 2,000 runners are running through the city, a diligent graduate student is working on his research topic. He studies diffusion limited aggregation (DLA), a process wherein a particle performs a random walk until it comes into contact with a growing cluster of particles. At the moment of contact, the particle sticks to the cluster at the location where the contact occurred and becomes part of the cluster. Then a new random walker is released to perform a random walk, until *it* sticks to the cluster. And so on. Clusters grown through this process have a remarkably open, tenuous structure (as shown in figure 1.2): they're *fractals*.²

The DLA process is simple, so it seems straightforward to write a program to grow such clusters in a computer, and this is what the busy graduate student has done. Initially, all seems well; but as the simulation progresses, the cluster appears to grow more and more slowly—excruciatingly slowly, in fact. The goal was to grow a DLA cluster in excess of 100,000 particles. Will the program ever finish?

¹ Depending on your gnuplot setup and initialization, your graphs may look slightly different from the figures shown in this chapter. We'll discuss user-defined appearance options starting with chapter 6.

² The original paper on DLA was "Diffusion Limited Aggregation, A Kinetic Critical Phenomenon" by T. A. Witten and L. M. Sander, *Physical Review Letters* 41 (1981): 1400. It's one of the most-quoted papers from that journal of all time. If you want to learn more about DLA and similar processes, check out *Fractals, Scaling, and Growth Far From Equilibrium* by Paul Meakin (Cambridge University Press, 1998).




Figure 1.2 A DLA cluster of $N=50,000$ particles, drawn with gnuplot

Luckily, the simulation program periodically writes information about its progress to a log file: for each new particle added to the cluster, the time (in seconds) since the start of the simulation is recorded. The grad student should be able to predict the completion time from this data, but an initial plot (figure 1.3) isn't helpful; there are too many ways this curve can be extrapolated to larger cluster sizes.




Figure 1.3 Time required to grow a DLA cluster

The time consumed by many computer algorithms grows as a simple power of the size of the problem. In this case, this would be the number N of particles in the cluster $T \sim N^k$, for some value of k . The research student therefore plots the running time of his simulation program on a double-logarithmic plot versus the cluster size (see figure 1.4). The data points fall on a straight line, indicating a power law. (I'll explain later how and why this works.) Through a little trial and error, he also finds an equation that approximates the data quite well. The equation can be extended to any cluster size desired and will give the time required. For $N=100,000$ (which was the original goal), he can read off almost $T=100,000$ seconds (or more), corresponding to more than 24 hours, so there is no point in your friend spending the weekend in the lab—he should go out (maybe run a marathon) and come back on Monday, or perhaps work on a better algorithm. (For simulations of DLA cluster growth, dramatic speedups over the naive implementation are possible. Try it if you like.)




Figure 1.4 Time required to grow a DLA cluster in a double-logarithmic plot, together with an approximate mathematical model

USING GNUPLOT

Again, let's see how the graphs in this section were created. The easiest to understand is figure 1.3. Given a file containing two columns, one listing the cluster size and the other listing the completion time, the command is just

```
plot "runtime" using 1:2 with lines
```

The only difference compared to figure 1.1 is the style: rather than boxes, I use line segments to connect consecutive data points: with lines.

Did you notice that figure 1.3 and figure 1.4 contain more than just data? Both axes are now labelled! Details such as labels and other helpful decorations often make the difference between a mediocre and a high-quality graph, because they provide the observer with the necessary context to fully understand the graph.

In gnuplot, all details of a graph's appearance are handled by setting the appropriate options. To place the labels on the x and y axes in figure 1.3, I used

```
set xlabel "Cluster size"
set ylabel "Run time [sec]"
```

Figure 1.4 is drawn using double-logarithmic axes. This is another option, which is set as follows:

```
set logscale
```

Figure 1.4 shows two curves: the data together with a best “fit.” Plotting several data sets or mathematical functions together in one plot is easy—you list them one after another on the command line for the `plot` command:

```
plot "runtime" using 1:2 title "Data" with lines,
➤ (x/2500)**3 title "Model"
```

This command introduces a further gnuplot feature: the `title` directive. It takes a string as argument, which is displayed together with a line sample in the plot's key or legend (visible at upper left in figure 1.4).

Finally, we come to figure 1.2. It's a somewhat different beast. Notice that the border and the tic marks are missing. The aspect ratio (the ratio of the graph's width to its height) has been constrained to 1, and a single dot has been placed at the position of each particle in the cluster. Here are the most important commands that I used:

```
unset border
unset xtics
unset ytics

set size square

plot "cluster" using 1:2 with dots
```

You can see that gnuplot is simple to use. In the next section, I talk more about using graphical methods to understand a data set, before coming back to gnuplot and discussing why it's my favorite tool for this kind of activity.

1.2 What is graphical analysis?

The previous two examples should have given you an idea of what graphical analysis is and how it works. The basic steps are always the same: