

gnuplot 5.0

An Interactive Plotting Program

Thomas Williams & Colin Kelley

Version 5.0 organized by: Ethan A Merritt and many others

Major contributors (alphabetic order):

Christoph Bersch, Hans-Bernhard Bröker,

John Campbell, Robert Cunningham,

David Denholm, Gershon Elber,

Roger Fearick, Carsten Grammes,

Lucas Hart, Lars Hecking, Péter Juhász,

Thomas Koenig, David Kotz,

Ed Kubaitis, Russell Lang, Timothée Lecomte,

Alexander Lehmann, Jérôme Lodewyck,

Alexander Mai, Bastian Märkisch,

Ethan A Merritt, Petr Mikulík,

Carsten Steger, Shigeharu Takeno,

Tom Tkacik, Jos Van der Woude,

James R. Van Zandt, Alex Woo, Johannes Zellner

Copyright © 1986 - 1993, 1998, 2004 Thomas Williams, Colin Kelley

Copyright © 2004 - 2015 various authors

Mailing list for comments: gnuplot-info@lists.sourceforge.net

Mailing list for bug reports: gnuplot-bugs@lists.sourceforge.net

Web access (preferred): <http://sourceforge.net/projects/gnuplot>

This manual was originally prepared by Dick Crawford.

Version 5.0 (January 2015)

Contents

I Gnuplot	15
Copyright	15
Introduction	15
Seeking-assistance	16
New features in version 5	17
New commands	18
Changes in version 5	18
Deprecated syntax	19
Batch/Interactive Operation	19
Canvas size	20
Command-line-editing	20
Comments	21
Coordinates	21
Datastrings	21
Enhanced text mode	22
Environment	23
Expressions	24
Functions	24
Elliptic integrals	26
Random number generator	26
Value	27
Counting and extracting words	27
Operators	27
Unary	27
Binary	28
Ternary	28
Summation	29
Gnuplot-defined variables	29
User-defined variables and functions	30
Fonts	31

CONTENTS	gnuplot 5.0	3
Cairo (pdfcairo, pngcairo, epscairo, wxt terminals)		31
Gd (png, gif, jpeg terminals)		31
Postscript (also encapsulated postscript *.eps)		32
Glossary		32
Iteration		33
Linetypes, colors, and styles		33
Colorspec		33
Background color		34
Linecolor variable		34
Rgbcolor variable		35
Dashtype		35
Linestyles vs linetypes		36
Layers		36
Mouse input		36
Bind		36
Bind space		38
Mouse variables		38
Persist		38
Plotting		38
Start-up (initialization)		39
String constants and string variables		39
Substitution and Command line macros		40
Substitution of system commands in backquotes		40
Substitution of string variables as macros		40
String variables, macros, and command line substitution		41
Syntax		41
Quote Marks		42
Time/Date data		42
II Plotting styles		44
Boxerrorbars		44
Boxes		44

Boxplot	45
Boxxyerrorbars	46
Candlesticks	46
Circles	47
Ellipses	47
Dots	48
Filledcurves	48
Financebars	49
Fsteps	49
Fillsteps	50
Histeps	50
Histograms	50
Newhistogram	52
Automated iteration over multiple columns	52
Image	53
Transparency	53
Image pixels	53
Impulses	54
Labels	54
Lines	54
Linespoints	55
Parallelaxes	55
Points	55
Polar	55
Steps	56
Rgbalpha	56
Rgbimage	56
Vectors	56

CONTENTS	gnuplot 5.0	5
Xerrorbars		56
Xyerrorbars		57
Yerrorbars		57
Xerrorlines		57
Xyerrorlines		57
Yerrorlines		58
3D (surface) plots		58
2D projection (set view map)		58
III Commands		59
Cd		59
Call		59
Example		60
Old-style		60
Clear		60
Do		61
Evaluate		61
Exit		62
Fit		62
Adjustable parameters		64
Short introduction		64
Error estimates		65
Statistical overview		65
Practical guidelines		66
Control		67
Control variables		67
Environment variables		68
Multi-branch		68
Starting values		68
Tips		69
Help		69
History		70

If	70
If-old	70
For	71
Import	71
Load	72
Lower	72
Pause	72
Plot	73
Axes	74
Binary	74
General	75
Array	75
Record	75
Skip	75
Format	76
Endian	76
Filetype	76
Avs	76
Edf	76
Png	76
Keywords	77
Scan	77
Transpose	77
Dx, dy, dz	77
Flipx, flipy, flipz	77
Origin	77
Center	77
Rotate	77
Perpendicular	78
Data	78
Every	79
Example datafile	80
Index	80
Inline data	81
Skip	81
Smooth	82
Acsplines	82
Bezier	82

Csplines	82
Mcsplines	82
Sbezier	83
Unique	83
Unwrap	83
Frequency	83
Cumulative	83
Cnormal	83
Kdensity	83
Special-filenames	83
Thru	85
Using	85
Using_examples	86
Pseudocolumns	87
Xticlabels	87
X2ticlabels	88
Yticlabels	88
Y2ticlabels	88
Zticlabels	88
Volatile	88
Errorbars	88
Errorlines	89
Functions	89
Parametric	89
Ranges	90
Sampling	91
For loops in plot command	91
Title	92
With	93
Print	95
Printerr	95
Pwd	95
Quit	95
Raise	95
Refresh	96
Replot	96
Reread	96

Reset	97
Save	97
Set-show	98
Angles	98
Arrow	99
Autoscale	100
Parametric mode	101
Polar mode	102
Bars	102
Bind	102
Bmargin	102
Border	102
Boxwidth	104
Color	104
Colorsequence	104
Clabel	105
Clip	105
Cntrlabel	105
Cntrparam	106
Color box	107
Colornames	108
Contour	108
Dashtype	109
Data style	109
Datafile	109
Set datafile fortran	109
Set datafile nofpe_trap	109
Set datafile missing	109
Set datafile separator	110
Set datafile commentschars	111
Set datafile binary	111
Decimalsign	111
Dgrid3d	112
Dummy	113
Encoding	114
Fit	115
Fontpath	116
Format	116
Gprintf	117
Format specifiers	117
Time/date specifiers	118

Examples	119
Function style	119
Functions	119
Grid	119
Hidden3d	120
Historysize	122
History	122
Isosamples	122
Key	122
Key placement	124
Key samples	125
Label	125
Examples	126
Hypertext	127
Linetype	128
Link	128
Lmargin	129
Loadpath	129
Locale	129
Logscale	129
Macros	130
Mapping	130
Margin	131
Monochrome	131
Mouse	131
Doubleclick	132
Mouseformat	132
Scrolling	133
X11 mouse	133
Zoom	133
Multiplot	133
Mx2tics	135
Mxtics	135
My2tics	136
Mytics	136
Mztics	136
Object	136
Rectangle	137
Ellipse	137
Circle	138
Polygon	138
Offsets	139

Origin	139
Output	139
Parametric	140
Paxis	140
Plot	140
Pm3d	141
Algorithm	141
Position	142
Scanorder	142
Clipping	143
Color_assignment	143
Corners2color	143
Border	143
Interpolate	144
Deprecated_options	144
Palette	144
Rgbformulae	146
Defined	146
Functions	147
Gray	147
Cubehelix	148
File	148
Gamma correction	148
Postscript	149
Pointintervalbox	149
Pointsize	149
Polar	149
Print	150
Pmdir	150
Raxis	151
Rmargin	151
Rrange	151
Rtics	151
Samples	151
Size	151
Style	152
Set style arrow	153
Boxplot	154
Set style data	154
Set style fill	155
Set style fill transparent	155
Set style function	155

Set style increment	156
Set style line	156
Set style circle	157
Set style rectangle	158
Set style ellipse	158
Set style textbox	158
Surface	158
Table	159
Terminal	159
Termoption	160
Tics	160
Ticslevel	161
Ticscale	161
Timestamp	161
Timefmt	161
Title	162
Tmargin	163
Trange	163
Urange	163
Variables	163
Version	163
View	164
Equal axes	164
Vrange	164
X2data	164
X2dtics	164
X2label	164
X2mtics	165
X2range	165
X2tics	165
X2zeroaxis	165
Xdata	165
Time	165
Xdtics	166
Xlabel	166
Xmtics	167
Xrange	167
Xtics	168
Xtics timedata	171
Geographic	171
Xtics rangelimited	171
Xyplane	172

Xzeroaxis	172
Y2data	172
Y2dtics	172
Y2label	172
Y2mtics	172
Y2range	173
Y2tics	173
Y2zeroaxis	173
Ydata	173
Ydtics	173
Ylabel	173
Ymtics	173
Yrange	173
Ytics	173
Yzeroaxis	173
Zdata	173
Zdtics	174
Zzeroaxis	174
Cbdata	174
Cbdtics	174
Zero	174
Zeroaxis	174
Zlabel	175
Zmtics	175
Zrange	175
Ztics	175
Cblabel	175
Cbmtics	175
Cbrange	175
Cbtics	175
Shell	175
Splot	176
Data-file	176
Matrix	177
Uniform	177
Nonuniform	178
Examples	178
Example datafile	179
Grid data	179
Splot surfaces	180

CONTENTS	gnuplot 5.0	13
Stats (Statistical Summary)		180
System		181
Test		182
Undefine		182
Unset		182
Linetype		182
Monochrome		183
Output		183
Terminal		183
Update		183
While		183
 IV Terminal types		 184
Complete list of terminals		184
Aifm		184
Aqua		184
Be		184
Command-line_options		185
Monochrome_options		185
Color_resources		185
Grayscale_resources		186
Line_resources		186
Cairolatex		187
Canvas		189
Cgm		190
Cgm font		190
Cgm fontsize		191
Cgm linewidth		191
Cgm rotate		192
Cgm solid		192
Cgm size		192
Cgm width		192
Cgm nofontlist		192
 V Bugs		 193
Known limitations		193

External libraries	193
---------------------------	------------

VI Index	193
-----------------	------------

Part I

Gnuplot

Copyright

Copyright (C) 1986 - 1993, 1998, 2004, 2007 Thomas Williams, Colin Kelley

Permission to use, copy, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation.

Permission to modify the software is granted, but not the right to distribute the complete modified source code. Modifications are to be distributed as patches to the released version. Permission to distribute binaries produced by compiling modified sources is granted, provided you

1. distribute the corresponding source modifications from the released version in the form of a patch file along with the binaries,
2. add special version identification to distinguish your version in addition to the base release version number,
3. provide your name and address as the primary contact for the support of your modified version, and
4. retain our contact information in regard to use of the base software.

Permission to distribute the released version of the source code along with corresponding source modifications in the form of a patch file is granted with same provisions 2 through 4 for binary distributions.

This software is provided "as is" without express or implied warranty to the extent permitted by applicable law.

AUTHORS

Original Software:

Thomas Williams, Colin Kelley.

Gnuplot 2.0 additions:

Russell Lang, Dave Kotz, John Campbell.

Gnuplot 3.0 additions:

Gershon Elber and many others.

Gnuplot 4.0 and 5.0 additions:

See list of contributors at head of this document.

Introduction

Gnuplot is a portable command-line driven graphing utility for Linux, OS/2, MS Windows, OSX, VMS, and many other platforms. The source code is copyrighted but freely distributed (i.e., you don't have to pay for it). It was originally created to allow scientists and students to visualize mathematical functions and data interactively, but has grown to support many non-interactive uses such as web scripting. It is also used as a plotting engine by third-party applications like Octave. Gnuplot has been supported and under active development since 1986.

Gnuplot supports many types of plots in either 2D and 3D. It can draw using lines, points, boxes, contours, vector fields, surfaces, and various associated text. It also supports various specialized plot types.

Gnuplot supports many different types of output: interactive screen terminals (with mouse and hotkey input), direct output to pen plotters or modern printers, and output to many file formats (eps, emf, fig, jpeg, LaTeX, pdf, png, postscript, ...). Gnuplot is easily extensible to include new output modes. Recent additions include interactive terminals based on wxWidgets (usable on multiple platforms), and Qt. Mouseable plots embedded in web pages can be generated using the svg or HTML5 canvas terminal drivers.

The command language of **gnuplot** is case sensitive, i.e. commands and function names written in lowercase are not the same as those written in capitals. All command names may be abbreviated as long as the

abbreviation is not ambiguous. Any number of commands may appear on a line, separated by semicolons (;). Strings may be set off by either single or double quotes, although there are some subtle differences. See **syntax** (p. 41) and **quotes** (p. 42) for more details. Example:

```
set title "My First Plot"; plot 'data'; print "all done!"
```

Commands may extend over several input lines by ending each line but the last with a backslash (\). The backslash must be the *last* character on each line. The effect is as if the backslash and newline were not there. That is, no white space is implied, nor is a comment terminated. Therefore, commenting out a continued line comments out the entire command (see **comments** (p. 21)). But note that if an error occurs somewhere on a multi-line command, the parser may not be able to locate precisely where the error is and in that case will not necessarily point to the correct line.

In this document, curly braces ({}) denote optional arguments and a vertical bar (|) separates mutually exclusive choices. **Gnuplot** keywords or **help** topics are indicated by backquotes or **boldface** (where available). Angle brackets (<>) are used to mark replaceable tokens. In many cases, a default value of the token will be taken for optional arguments if the token is omitted, but these cases are not always denoted with braces around the angle brackets.

For built-in help on any topic, type **help** followed by the name of the topic or **help ?** to get a menu of available topics.

A large set of demo plots is available on the web page

<http://www.gnuplot.info/demo/>

When run from command line, gnuplot is invoked using the syntax

```
gnuplot {OPTIONS} file1 file2 ...
```

where file1, file2, etc. are input file as in the **load** command. On X11-based systems, you can use

```
gnuplot {X11OPTIONS} {OPTIONS} file1 file2 ...
```

see your X11 documentation and **x11** (p. ??) in this document.

Options interpreted by gnuplot may come anywhere on the line. Files are executed in the order specified, as are commands supplied by the -e option, for example

```
gnuplot file1.in -e "reset" file2.in
```

The special filename "-" is used to force reading from stdin. **Gnuplot** exits after the last file is processed. If no load files are named, **Gnuplot** takes interactive input from stdin. See help **batch/interactive** (p. 19) for more details. The options specific to gnuplot can be listed by typing

```
gnuplot --help
```

See **command line options** (p. 19) for more details.

In sessions with an interactive plot window you can hit 'h' anywhere on the plot for help about **hotkeys** and **mousing** features. Section **seeking-assistance** will help you to find further information, help and FAQ.

Seeking-assistance

The canonical gnuplot web page can be found at

<http://www.gnuplot.info>

Before seeking help, please check file FAQ.pdf or the above website for

[FAQ \(Frequently Asked Questions\) list](#).

If you need help as a gnuplot user, please use the newsgroup

```
comp.graphics.apps.gnuplot
```

Instructions for subscribing to gnuplot mailing lists may be found via the gnuplot development website on SourceForge

<http://sourceforge.net/projects/gnuplot>

Please note that before you write to any of the gnuplot mailing lists, you have to subscribe to the list first. This is necessary to keep the spam level down.

The address for mailing to list members is:

`gnuplot-info@lists.sourceforge.net`

Bug reports and code contributions should be uploaded to the trackers at

`http://sourceforge.net/projects/gnuplot/support`

Please check previous bug reports if the bug you want to report has not been already fixed in a newer version.

A mailing list for those interested in development version of gnuplot is:

`gnuplot-beta@lists.sourceforge.net`

When posting a question, please include full details of the gnuplot version, the terminal type, and the operating system you are using. A *small* script demonstrating the problem may be useful. Function plots are preferable to datafile plots.

New features in version 5

* The dot-dash pattern of a line can now be specified independent of other line properties. See **dashtype** (p. 35), **set dashtype** (p. 109), **set linetype** (p. 128).

* Text markup now supports bold and italic font settings in addition to subscript, superscript, font size and other previously available properties. Enhanced text mode is now enabled by default. See **enhanced text** (p. 22).

* Interactive terminals support hypertext labels that only appear when the mouse hovers over the label's anchor point.

* New coordinate system (Degrees, Minutes, Seconds). See **set xtics geographic** (p. 171).

* The default format for axis labels is "% h" ("%h\$" for LaTeX terminals). This format is like the C standard format %g except that the exponential term, if present, is written using a superscript. E.g. 1.2×10^5 rather than 1.2E05.

* Command scripts may place in-line data in a named data block for repeated plotting. See **inline data** (p. 81).

* Support for 32-bit Alpha channel + RGB color #AARRGGBB. See **colourspec** (p. 33).

* Support for HSV color space via a translation function `hsv2rgb(H,S,V)`.

* Secondary axes (x2, y2) may be locked to the primary axis via a mapping function. In the simplest case this guarantees that the primary and secondary axis ranges are identical. In the general case it allows you to define a non-linear axis, something that previously was possible only for log scaling. See **set link** (p. 128).

* Each function in a plot command may optionally be preceded by a sampling range. This does not affect the overall range of the plot, only the range over which this function is sampled. See **plot** (p. 73) and **piecewise.dem**.

* If the external library libcerf is available, it is used to provide complex math routines `cerf`, `cdawson`, `erfi`, `faddeeva`, and the Voigt profile `VP(x,sigma,gamma)`.

* The **import** command attaches a user-defined function name to a function provided by an external shared object (support is operating-system dependent). A template header and example source and make files for creating a suitable external shared object are provided in the demo collection.

* Previous commands in the history list of an interactive session can be reexecuted by number. For example, **history !5** will reexecute the command numbered 5 in the **history** list.

* Bit-shift operators `>>` and `<<`.

* New plot styles: **with parallelaxes**, **with table**, labeled contours.

* Shell invocation of gnuplot can pass parameters to a gnuplot script. `gnuplot -c scriptfile.gp ARG1 ARG2 ARG3 ...`

New commands

- `import f(x) from "plugin.so"` # load function from shared library
- `set history {quiet|numbers} {full|trim}` # controls output of history command
- `history !N` # re-execute prior command by number
- `plot <datafile> skip N` # skip lines at start of ascii data file
- `plot ... smooth mcsplines` # monotonic cubic spline fit to data
- `reset session` # restore initial state of current session
- `set arrow <tag> from <start> length <len> angle <ang>`
- `set colorsequence default|classic|podo` # colors used by successive plot elements
- `set monochrome` # alternative set of linetypes
- `set dashtype <tag> <dash-spec>` # user-defined dash patterns
- `set link x2 via f(x) inverse g(x)` # allows non-linear axis scaling
- `set fit quiet|results|brief|verbose` # control the amount of fit output
- `set contours; splot ... with labels` # label contour lines with numeric values
- `set style textbox` # text elements can be enclosed in a box
- `set view map {scale}` # allows resizing a 3D projection plot
- `set multiplot {next|previous}` # navigate within the auto-layout grid

Changes in version 5

These changes introduced in version 5 may cause certain scripts written for earlier versions of gnuplot to behave differently.

* Revised handling of input data containing NaN, inconsistent number of data columns, or other unexpected content. See Note under **missing** (p. 109) for examples and figures.

* Time coordinates are stored internally as the number of seconds relative to the standard unix epoch 1-Jan-1970. Earlier versions of gnuplot used a different epoch internally (1-Jan-2000). This change resolves inconsistencies introduced whenever time in seconds was generated externally. The epoch convention used by a particular gnuplot installation can be determined using the command **print strftime("%F",0)**. Time is now stored to at least millisecond precision.

* The function **timecolumn(N,"timeformat")** now has 2 parameters. Because the new second parameter is not associated with any particular data axis, this allows using the **timecolumn** function to read time data for reasons other than specifying the x or y coordinate. This functionality replaces the command sequence **set xdata time; set timefmt "timeformat"**. It allows combining time data read from multiple files with different formats within a single plot.

* The **reverse** keyword of the **set [axis]range** command affects only autoscaling. It does not invert or otherwise alter the meaning of a command such as **set xrange [0:1]**. If you want to reverse the direction of the x axis in such a case, say instead **set xrange [1:0]**.

* The **call** command is implemented by providing a set of variables ARG0, ARG1, ..., ARG9. ARG0 holds the name of the script file being executed. ARG1 to ARG9 are string variables and thus may either be referenced directly or expanded as macros, e.g. @ARG1. The older convention for referencing call parameters as tokens \$0 ... \$9 is deprecated.

* The optional bandwidth for the kernel density smoothing option is taken from a keyword rather than a data column. See **smooth kdensity** (p. 83).

* **unset xrange** (and other axis ranges) restores the original default range.

* **unset terminal** restores the original terminal of the gnuplot session.

Deprecated syntax

Gnuplot version 4 deprecated certain syntax used in earlier versions but provided a configuration option that allowed backward compatibility. Support for the old syntax has now been removed.

Deprecated in version 4 and removed in version 5:

```
set title "Old" 0,-1
set data linespoints
plot 1 2 4                # horizontal line at y=1
```

Current equivalent:

```
TITLE = "New"
set title TITLE offset char 0, char -1
set style data linespoints
plot 1 linetype 2 pointtype 4
```

Deprecated but present in version 5 if configured `--enable-backwards-compatibility`

```
if (defined(VARNAME)) ...
set style increment user
plot 'file' thru f(x)
call 'script' 1.23 ABC
    (in script: print $0, "$1", "number of args = $#")
```

Current equivalent:

```
if (exists("VARNAME")) ...
set linetype
plot 'file' using 1:(f(column(2)))
call 'script' 1.23 "ABC"
    (in script: print ARG1, ARG2, "number of args = ", ARGV)
```

Batch/Interactive Operation

Gnuplot may be executed in either batch or interactive modes, and the two may even be mixed together on many systems.

Any command-line arguments are assumed to be either program options (first character is -) or names of files containing **gnuplot** commands. The option `-e "command"` may be used to force execution of a **gnuplot** command. Each file or command string will be executed in the order specified. The special filename "-" indicates that commands are to be read from stdin. **Gnuplot** exits after the last file is processed. If no load files and no command strings are specified, **gnuplot** accepts interactive input from stdin.

Both the **exit** and **quit** commands terminate the current command file and **load** the next one, until all have been processed.

Examples:

To launch an interactive session:

```
gnuplot
```

To launch a batch session using two command files "input1" and "input2":

```
gnuplot input1 input2
```

To launch an interactive session after an initialization file "header" and followed by another command file "trailer":

```
gnuplot header - trailer
```

To give **gnuplot** commands directly in the command line, using the `"-persist"` option so that the plot remains on the screen afterwards:

```
gnuplot -persist -e "set title 'Sine curve'; plot sin(x)"
```

To set user-defined variables a and s prior to executing commands from a file:

```
gnuplot -e "a=2; s='file.png'" input.gpl
```

Canvas size

In earlier versions of gnuplot, some terminal types used the values from **set size** to control also the size of the output canvas; others did not. The use of 'set size' for this purpose was deprecated in version 4.2. Almost all terminals now behave as follows:

set term <terminal_type> size <XX>, <YY> controls the size of the output file, or "canvas". By default, the plot will fill this canvas.

set size <XX>, <YY> scales the plot itself relative to the size of the canvas. Scale values less than 1 will cause the plot to not fill the entire canvas. Scale values larger than 1 will cause only a portion of the plot to fit on the canvas. Please be aware that setting scale values larger than 1 may cause problems on some terminal types.

The major exception to this convention is the PostScript driver, which by default continues to act as it has in earlier versions. Be warned that some future version of gnuplot may change the default behaviour of the PostScript driver as well.

Example:

```
set size 0.5, 0.5
set term png size 600, 400
set output "figure.png"
plot "data" with lines
```

These commands will produce an output file "figure.png" that is 600 pixels wide and 400 pixels tall. The plot will fill the lower left quarter of this canvas. This is consistent with the way multiplot mode has always worked.

Command-line-editing

Command-line editing and command history are supported using either an external gnu readline library, an external BSD libedit library, or a built-in equivalent. This choice is a configuration option at the time gnuplot is built.

The editing commands of the built-in version are given below. Please note that the action of the DEL key is system-dependent. The gnu readline and BSD libedit libraries have their own documentation.

Command-line Editing Commands	
Character	Function
Line Editing	
^B	move back a single character.
^F	move forward a single character.
^A	move to the beginning of the line.
^E	move to the end of the line.
^H	delete the previous character.
DEL	delete the current character.
^D	delete current character. EOF if line is empty.
^K	delete from current position to the end of line.
^L, ^R	redraw line in case it gets trashed.
^U	delete the entire line.
^W	delete previous word.
History	
^P	move back through history.
^N	move forward through history.

Comments

Comments are supported as follows: a **#** may appear in most places in a line and **gnuplot** will ignore the rest of the line. It will not have this effect inside quotes, inside numbers (including complex numbers), inside command substitutions, etc. In short, it works anywhere it makes sense to work.

See also **set datafile commentschars** (p. 111) for specifying comment characters in data files. Note that if a comment line ends in `'\'` then the subsequent line is also treated as a comment.

Coordinates

The commands **set arrow**, **set key**, **set label** and **set object** allow you to draw something at an arbitrary position on the graph. This position is specified by the syntax:

```
{<system>} <x>, {<system>} <y> [{<system>} <z>]}
```

Each `<system>` can either be **first**, **second**, **graph**, **screen**, or **character**.

first places the x, y, or z coordinate in the system defined by the left and bottom axes; **second** places it in the system defined by the x2,y2 axes (top and right); **graph** specifies the area within the axes — 0,0 is bottom left and 1,1 is top right (for **splot**, 0,0,0 is bottom left of plotting area; use negative z to get to the base — see **set xyplane** (p. 172)); **screen** specifies the screen area (the entire area — not just the portion selected by **set size**), with 0,0 at bottom left and 1,1 at top right; and **character** gives the position in character widths and heights from the bottom left of the screen area (screen 0,0), **character** coordinates depend on the chosen font size.

If the coordinate system for x is not specified, **first** is used. If the system for y is not specified, the one used for x is adopted.

In some cases, the given coordinate is not an absolute position but a relative value (e.g., the second position in **set arrow ... rto**). In most cases, the given value serves as difference to the first position. If the given coordinate belongs to a log-scaled axis, a relative value is interpreted as multiplier. For example,

```
set logscale x
set arrow 100,5 rto 10,2
```

plots an arrow from position 100,5 to position 1000,7 since the x axis is logarithmic while the y axis is linear.

If one (or more) axis is timeseries, the appropriate coordinate should be given as a quoted time string according to the **timefmt** format string. See **set xdata** (p. 165) and **set timefmt** (p. 161). **Gnuplot** will also accept an integer expression, which will be interpreted as seconds relative to 1 January 1970.

Datastrings

Data files may contain string data consisting of either an arbitrary string of printable characters containing no whitespace or an arbitrary string of characters, possibly including whitespace, delimited by double quotes. The following line from a datafile is interpreted to contain four columns, with a text field in column 3:

```
1.000 2.000 "Third column is all of this text" 4.00
```

Text fields can be positioned within a 2-D or 3-D plot using the commands:

```
plot 'datafile' using 1:2:4 with labels
splot 'datafile' using 1:2:3:4 with labels
```

A column of text data can also be used to label the ticmarks along one or more of the plot axes. The example below plots a line through a series of points with (X,Y) coordinates taken from columns 3 and 4 of the input datafile. However, rather than generating regularly spaced tics along the x axis labeled numerically, **gnuplot** will position a tic mark along the x axis at the X coordinate of each point and label the tic mark with text taken from column 1 of the input datafile.

```
set xtics
plot 'datafile' using 3:4:xticlabels(1) with linespoints
```

There is also an option that will interpret the first entry in a column of input data (i.e. the column heading) as a text field, and use it as the key title for data plotted from that column. The example given below will use the first entry in column 2 to generate a title in the key box, while processing the remainder of columns 2 and 4 to draw the required line:

```
plot 'datafile' using 1:(f($2)/$4) with lines title columnhead(2)
```

Another example:

```
plot for [i=2:6] 'datafile' using i title "Results for ".columnhead(i)
```

See [labels \(p. 54\)](#), [using xticlabels \(p. 87\)](#), [plot title \(p. 92\)](#), [using \(p. 85\)](#).

Enhanced text mode

Many terminal types support an enhanced text mode in which additional formatting information is embedded in the text string. For example, "x²" will write x-squared as we are used to seeing it, with a superscript 2. This mode is selected by default when you set the terminal, but may be toggled afterward using "set termoption [no]enhanced", or by marking individual strings as in "set label 'x₂' noenhanced".

Enhanced Text Control Codes			
Control	Example	Result	Explanation
<code>^</code>	<code>a^x</code>	a^x	superscript
<code>_</code>	<code>a_x</code>	a_x	subscript
<code>@</code>	<code>a@^b_{cd}</code>	a_{cd}^b	phantom box (occupies no width)
<code>&</code>	<code>d&{space}b</code>	$d_{\text{space}}b$	inserts space of specified length
<code>~</code>	<code>~a{.8-}</code>	\tilde{a}	overprints '-' on 'a', raised by .8 times the current fontsize
	<code>{/Times abc}</code>	abc	print abc in font Times at current size
	<code>{/Times*2 abc}</code>	abc	print abc in font Times at twice current size
	<code>{/Times:Italic abc}</code>	abc	print abc in font Times with style italic
	<code>{/Arial:Bold=20 abc}</code>	abc	print abc in boldface Arial font size 20

The markup control characters act on the following single character or bracketed clause. The bracketed clause may contain a string of characters with no additional markup, e.g. $2^{\{10\}}$, or it may contain additional markup that changes font properties. This example illustrates nesting one bracketed clause inside another to produce a boldface A with an italic subscript i, all in the current font. If the clause introduced by `:Normal` were omitted the subscript would be both italic and boldface.

```
{/ :Bold A_{/ :Normal{/ :Italic i}}
```

Font specifiers MUST be preceded by a '/' character that immediately follows the opening '{'.

The phantom box is useful for `a@^b.c` to align superscripts and subscripts but does not work well for overwriting an accent on a letter. For the latter, it is much better to use an encoding (e.g. `iso_8859_1` or `utf8`) that contains a large variety of letters with accents or other diacritical marks. See [set encoding \(p. 114\)](#). Since the box is non-spacing, it is sensible to put the shorter of the subscript or superscript in the box (that is, after the `@`).

Space equal in length to a string can be inserted using the `&` character. Thus

```
'abc&{def}ghi'
```

would produce

```
'abc   ghi'.
```

The `~` character causes the next character or bracketed text to be overprinted by the following character or bracketed text. The second text will be horizontally centered on the first. Thus `~ a/` will result in an 'a' with a slash through it. You can also shift the second text vertically by preceding the second text with a number, which will define the fraction of the current fontsize by which the text will be raised or lowered. In this case the number and text must be enclosed in brackets because more than one character is necessary. If the overprinted text begins with a number, put a space between the vertical offset and the text (`~ {abc}{.5 000}`); otherwise no space is needed (`~ {abc}{.5 — }`). You can change the font for one or both strings (`~ a{.5 /*.2 o}` — an 'a' with a one-fifth-size 'o' on top — and the space between the number and the slash is necessary), but you can't change it after the beginning of the string. Neither can you use any other special syntax within either string. You can, of course, use control characters by escaping them (see below), such as `~ a{\~}`

You can specify special symbols numerically by giving a character code in octal, e.g. `{/Symbol \245}` is the symbol for infinity in the Adobe Symbol font. This does not work for multibyte encodings like UTF-8, however. In a UTF-8 environment, you should be able to enter multibyte sequences implicitly by typing or otherwise selecting the character you want.

You can escape control characters using `\`, e.g., `\\`, `\{`, and so on.

Note that strings in double-quotes are parsed differently than those enclosed in single-quotes. The major difference is that backslashes may need to be doubled when in double-quoted strings.

The file "ps_guide.ps" in the `/docs/psdoc` subdirectory of the gnuplot source distribution contains more examples of the enhanced syntax, as does the demo

[enhanced_utf8.dem](#)

Environment

A number of shell environment variables are understood by **gnuplot**. None of these are required, but may be useful.

GNUTERM, if defined, is used as the default terminal type on start-up. This can be overridden by the `~/.gnuplot` (or equivalent) start-up file (see **startup** (p. 39)) and of course by later explicit **set term** commands.

GNUHELP may be defined to be the pathname of the HELP file (gnuplot.gih).

On VMS, the logical name GNUPLOT\$HELP should be defined as the name of the help library for **gnuplot**. The **gnuplot** help can be put inside any VMS system help library.

On Unix, HOME is used as the name of a directory to search for a .gnuplot file if none is found in the current directory. On MS-DOS, Windows and OS/2, GNUPLOT is used. On Windows, the NT-specific variable USERPROFILE is also tried. VMS, SYS\$LOGIN: is used. Type **help startup**.

On Unix, PAGER is used as an output filter for help messages.

On Unix, SHELL is used for the **shell** command. On MS-DOS and OS/2, COMSPEC is used for the **shell** command.

FIT_SCRIPT may be used to specify a **gnuplot** command to be executed when a fit is interrupted — see **fit** (p. 62). **FIT_LOG** specifies the default filename of the logfile maintained by fit.

GNUPLOT_LIB may be used to define additional search directories for data and command files. The variable may contain a single directory name, or a list of directories separated by a platform-specific path separator, eg. `:'` on Unix, or `;` on DOS/Windows/OS/2 platforms. The contents of GNUPLOT_LIB are appended to the **loadpath** variable, but not saved with the **save** and **save set** commands.

Several gnuplot terminal drivers access TrueType fonts via the gd library. For these drivers the font search path is controlled by the environmental variable GDFONTPATH. Furthermore, a default font for these drivers may be set via the environmental variable GNUPLOT_DEFAULT_GDFONT.

The postscript terminal uses its own font search path. It is controlled by the environmental variable GNUPLOT_FONTPATH. The format is the same as for GNUPLOT_LIB. The contents of GNUPLOT_FONTPATH are appended to the **fontpath** variable, but not saved with the **save** and **save set**

commands.

GNUPLOT_PS_DIR is used by the postscript driver to search for external prologue files. Depending on the build process, gnuplot contains either a built-in copy of those files or a default hardcoded path. You can use this variable have the postscript terminal use custom prologue files rather than the default files. See **postscript prologue** (p. ??).

Expressions

In general, any mathematical expression accepted by C, FORTRAN, Pascal, or BASIC is valid. The precedence of these operators is determined by the specifications of the C programming language. White space (spaces and tabs) is ignored inside expressions.

Complex constants are expressed as {<real>,<imag>}, where <real> and <imag> must be numerical constants. For example, {3,2} represents $3 + 2i$; {0,1} represents 'i' itself. The curly braces are explicitly required here.

Integer constants are interpreted via the C library routine strtoll(). This means that constants beginning with "0" are interpreted as octal, and constants beginning with "0x" or "0X" are interpreted as hexadecimal.

Floating point constants are interpreted via the C library routine atof().

Note that gnuplot uses both "real" and "integer" arithmetic, like FORTRAN and C. Integers are entered as "1", "-10", etc; reals as "1.0", "-10.0", "1e1", 3.5e-1, etc. The most important difference between the two forms is in division: division of integers truncates: $5/2 = 2$; division of reals does not: $5.0/2.0 = 2.5$. In mixed expressions, integers are "promoted" to reals before evaluation: $5/2e0 = 2.5$. The result of division of a negative integer by a positive one may vary among compilers. Try a test like "print -5/2" to determine if your system chooses -2 or -3 as the answer.

The integer expression "1/0" may be used to generate an "undefined" flag, which causes a point to ignored. Or you can use the pre-defined variable NaN to achieve the same result. See **using** (p. 85) for an example.

The real and imaginary parts of complex expressions are always real, whatever the form in which they are entered: in {3,2} the "3" and "2" are reals, not integers.

Gnuplot can also perform simple operations on strings and string variables. For example, the expression ("A" . "B" eq "AB") evaluates as true, illustrating the string concatenation operator and the string equality operator.

A string which contains a numerical value is promoted to the corresponding integer or real value if used in a numerical expression. Thus ("3" + "4" == 7) and (6.78 == "6.78") both evaluate to true. An integer, but not a real or complex value, is promoted to a string if used in string concatenation. A typical case is the use of integers to construct file names or other strings; e.g. ("file" . 4 eq "file4") is true.

Substrings can be specified using a postfix range descriptor [beg:end]. For example, "ABCDEF"[3:4] == "CD" and "ABCDEF"[4:*] == "DEF". The syntax "string"[beg:end] is exactly equivalent to calling the built-in string-valued function substr("string",beg,end), except that you cannot omit either beg or end from the function call.

Functions

Arguments to math functions in **gnuplot** can be integer, real, or complex unless otherwise noted. Functions that accept or return angles (e.g. sin(x)) treat angle values as radians, but this may be changed to degrees using the command **set angles**.

Math library functions		
Function	Arguments	Returns
abs(x)	any	absolute value of x , $ x $; same type
abs(x)	complex	length of x , $\sqrt{\text{real}(x)^2 + \text{imag}(x)^2}$
acos(x)	any	$\cos^{-1} x$ (inverse cosine)

Math library functions		
Function	Arguments	Returns
acosh(x)	any	$\cosh^{-1} x$ (inverse hyperbolic cosine) in radians
airy(x)	any	Airy function $\text{Ai}(x)$
arg(x)	complex	the phase of x
asin(x)	any	$\sin^{-1} x$ (inverse sin)
asinh(x)	any	$\sinh^{-1} x$ (inverse hyperbolic sin) in radians
atan(x)	any	$\tan^{-1} x$ (inverse tangent)
atan2(y,x)	int or real	$\tan^{-1}(y/x)$ (inverse tangent)
atanh(x)	any	$\tanh^{-1} x$ (inverse hyperbolic tangent) in radians
EllipticK(k)	real $k \in (-1:1)$	$K(k)$ complete elliptic integral of the first kind
EllipticE(k)	real $k \in [-1:1]$	$E(k)$ complete elliptic integral of the second kind
EllipticPi(n,k)	real $n < 1$, real $k \in (-1:1)$	$\Pi(n, k)$ complete elliptic integral of the third kind
besj0(x)	int or real	J_0 Bessel function of x , in radians
besj1(x)	int or real	J_1 Bessel function of x , in radians
besy0(x)	int or real	Y_0 Bessel function of x , in radians
besy1(x)	int or real	Y_1 Bessel function of x , in radians
ceil(x)	any	$\lceil x \rceil$, smallest integer not less than x (real part)
cos(x)	any	$\cos x$, cosine of x
cosh(x)	any	$\cosh x$, hyperbolic cosine of x in radians
erf(x)	any	$\text{erf}(\text{real}(x))$, error function of $\text{real}(x)$
erfc(x)	any	$\text{erfc}(\text{real}(x))$, 1.0 - error function of $\text{real}(x)$
exp(x)	any	e^x , exponential function of x
expint(n,x)	int $n \geq 0$, real $x \geq 0$	$E_n(x) = \int_1^\infty t^{-n} e^{-xt} dt$, exponential integral of x
floor(x)	any	$\lfloor x \rfloor$, largest integer not greater than x (real part)
gamma(x)	any	$\text{gamma}(\text{real}(x))$, gamma function of $\text{real}(x)$
ibeta(p,q,x)	any	$\text{ibeta}(\text{real}(p, q, x))$, ibeta function of $\text{real}(p, q, x)$
inverf(x)	any	inverse error function of $\text{real}(x)$
igamma(a,x)	any	$\text{igamma}(\text{real}(a, x))$, igamma function of $\text{real}(a, x)$
imag(x)	complex	imaginary part of x as a real number
invnorm(x)	any	inverse normal distribution function of $\text{real}(x)$
int(x)	real	integer part of x , truncated toward zero
lambertw(x)	real	Lambert W function
lgamma(x)	any	$\text{lgamma}(\text{real}(x))$, lgamma function of $\text{real}(x)$
log(x)	any	$\log_e x$, natural logarithm (base e) of x
log10(x)	any	$\log_{10} x$, logarithm (base 10) of x
norm(x)	any	normal distribution (Gaussian) function of $\text{real}(x)$
rand(x)	int	pseudo random number in the interval $[0:1]$
real(x)	any	real part of x
sgn(x)	any	1 if $x > 0$, -1 if $x < 0$, 0 if $x = 0$. $\text{imag}(x)$ ignored
sin(x)	any	$\sin x$, sine of x
sinh(x)	any	$\sinh x$, hyperbolic sine of x in radians
sqrt(x)	any	\sqrt{x} , square root of x
tan(x)	any	$\tan x$, tangent of x
tanh(x)	any	$\tanh x$, hyperbolic tangent of x in radians
voigt(x,y)	real	Voigt/Faddeeva function $\frac{y}{\pi} \int \frac{\exp(-t^2)}{(x-t)^2 + y^2} dt$ Note: $\text{voigt}(x, y) = \text{real}(\text{faddeeva}(x + iy))$

Special functions from libcerf (only if available)		
Function	Arguments	Returns
cerf(z)	complex	complex error function
cdawson(z)	complex	complex extension of Dawson's integral $D(z) = \frac{\sqrt{\pi}}{2} e^{-z^2} \text{erfi}(z)$
faddeeva(z)	complex	rescaled complex error function $w(z) = e^{-z^2} \text{erfc}(-iz)$

Special functions from libcerf (only if available)		
Function	Arguments	Returns
erfi(x)	real	imaginary error function $erf(x) = -i * erf(ix)$
VP(x,σ,γ)	real	Voigt profile $VP(x, \sigma, \gamma) = \int_{-\infty}^{\infty} G(x'; \sigma) L(x - x'; \gamma) dx'$

String functions		
Function	Arguments	Returns
gprintf("format",x,...)	any	string result from applying gnuplot's format parser
sprintf("format",x,...)	multiple	string result from C-language sprintf
strlen("string")	string	int length of string in bytes
strstr("string","key")	strings	int index of first character of substring "key"
substr("string",beg,end)	multiple	string "string"[beg:end]
strftime("timeformat",t)	any	string result from applying gnuplot's time parser
strptime("timeformat",s)	string	seconds since year 1970 as given in string s
system("command")	string	string containing output stream of shell command
word("string",n)	string, int	returns the nth word in "string"
words("string")	string	returns the number of words in "string"

other gnuplot functions		
Function	Arguments	Returns
column(x)	int or string	column x during datafile manipulation.
columnhead(x)	int	string containing first entry of column x in datafile.
exists("X")	string	returns 1 if a variable named X is defined, 0 otherwise.
hsv2rgb(h,s,v)	$h,s,v \in [0:1]$	24bit RGB color value.
stringcolumn(x)	int or string	content of column x as a string.
timecolumn(N,"timeformat")	int, string	time data from column N during data input.
tm_hour(x)	int	the hour
tm_mday(x)	int	the day of the month
tm_min(x)	int	the minute
tm_mon(x)	int	the month
tm_sec(x)	int	the second
tm_wday(x)	int	the day of the week
tm_yday(x)	int	the day of the year
tm_year(x)	int	the year
time(x)	any	the current system time
valid(x)	int	test validity of column(x) during datafile manip.
value("name")	string	returns the value of the named variable.

Elliptic integrals

The **EllipticK(k)** function returns the complete elliptic integral of the first kind, i.e. the definite integral between 0 and $\pi/2$ of the function $(1-(k*\sin(p))^2)^{-0.5}$. The domain of **k** is -1 to 1 (exclusive).

The **EllipticE(k)** function returns the complete elliptic integral of the second kind, i.e. the definite integral between 0 and $\pi/2$ of the function $(1-(k*\sin(p))^2)^{0.5}$. The domain of **k** is -1 to 1 (inclusive).

The **EllipticPi(n,k)** function returns the complete elliptic integral of the third kind, i.e. the definite integral between 0 and $\pi/2$ of the function $(1-(k*\sin(p))^2)^{-0.5}/(1-n*\sin(p)^2)$. The parameter **n** must be less than 1, while **k** must lie between -1 and 1 (exclusive). Note that by definition $\text{EllipticPi}(0,k) == \text{EllipticK}(k)$ for all possible values of **k**.

Random number generator

The function **rand()** produces a sequence of pseudo-random numbers between 0 and 1 using an algorithm from P. L'Ecuyer and S. Cote, "Implementing a random number package with splitting facilities", ACM

Transactions on Mathematical Software, 17:98-111 (1991).

```

rand(0)      returns a pseudo random number in the interval [0:1]
              generated from the current value of two internal
              32-bit seeds.
rand(-1)     resets both seeds to a standard value.
rand(x)      for integer 0 < x < 2^31-1 sets both internal seeds
              to x.
rand({x,y})  for integer 0 < x,y < 2^31-1 sets seed1 to x and
              seed2 to y.

```

Value

`B = value("A")` is effectively the same as `B = A`, where `A` is the name of a user-defined variable. This is useful when the name of the variable is itself held in a string variable. See **user-defined variables (p. 30)**. It also allows you to read the name of a variable from a data file. If the argument is a numerical expression, `value()` returns the value of that expression. If the argument is a string that does not correspond to a currently defined variable, `value()` returns NaN.

Counting and extracting words

`word("string",n)` returns the *n*th word in string. For example, `word("one two three",2)` returns the string "two".

`words("string")` returns the number of words in string. For example, `words(" a b c d")` returns 4.

The **word** and **words** functions provide limited support for quoted strings, both single and double quotes can be used:

```
print words("\"double quotes\" or 'single quotes'") # 3
```

A starting quote must either be preceeded by a white space, or start the string. This means that apostrophes in the middle or at the end of words are considered as parts of the respective word:

```
print words("Alexis' phone doesn't work") # 4
```

Escaping quote characters is not supported. If you want to keep certain quotes, the respective section must be surrounded by the other kind of quotes:

```

s = "Keep \"'single quotes\" or '\"double quotes\"'"
print word(s, 2) # 'single quotes'
print word(s, 4) # "double quotes"

```

Note, that in this last example the escaped quotes are necessary only for the string definition.

Operators

The operators in **gnuplot** are the same as the corresponding operators in the C programming language, except that all operators accept integer, real, and complex arguments, unless otherwise noted. The ****** operator (exponentiation) is supported, as in FORTRAN.

Parentheses may be used to change order of evaluation.

Unary

The following is a list of all the unary operators and their usages:

Unary Operators		
Symbol	Example	Explanation
-	-a	unary minus
+	+a	unary plus (no-operation)
~	~a	* one's complement
!	!a	* logical negation
!	a!	* factorial
\$	\$3	* call arg/column during 'using' manipulation

(*) Starred explanations indicate that the operator requires an integer argument.

Operator precedence is the same as in Fortran and C. As in those languages, parentheses may be used to change the order of operation. Thus $-2**2 = -4$, but $(-2)**2 = 4$.

The factorial operator returns a real number to allow a greater range.

Binary

The following is a list of all the binary operators and their usages:

Binary Operators		
Symbol	Example	Explanation
**	a**b	exponentiation
*	a*b	multiplication
/	a/b	division
%	a%b	* modulo
+	a+b	addition
-	a-b	subtraction
==	a==b	equality
!=	a!=b	inequality
<	a<b	less than
<=	a<=b	less than or equal to
>	a>b	greater than
>=	a>=b	greater than or equal to
<<	0xff<<1	left shift unsigned
>>	0xff>>1	right shift unsigned
&	a&b	* bitwise AND
^	a^b	* bitwise exclusive OR
	a b	* bitwise inclusive OR
&&	a&&b	* logical AND
	a b	* logical OR
=	a = b	assignment
,	(a,b)	serial evaluation
.	A.B	string concatenation
eq	A eq B	string equality
ne	A ne B	string inequality

(*) Starred explanations indicate that the operator requires integer arguments. Capital letters A and B indicate that the operator requires string arguments.

Logical AND (&&) and OR (||) short-circuit the way they do in C. That is, the second && operand is not evaluated if the first is false; the second || operand is not evaluated if the first is true.

Serial evaluation occurs only in parentheses and is guaranteed to proceed in left to right order. The value of the rightmost subexpression is returned.

Ternary

There is a single ternary operator:

Ternary Operator		
Symbol	Example	Explanation
?:	a?b:c	ternary operation

The ternary operator behaves as it does in C. The first argument (a), which must be an integer, is evaluated. If it is true (non-zero), the second argument (b) is evaluated and returned; otherwise the third argument (c) is evaluated and returned.

The ternary operator is very useful both in constructing piecewise functions and in plotting points only when certain conditions are met.

Examples:

Plot a function that is to equal $\sin(x)$ for $0 \leq x < 1$, $1/x$ for $1 \leq x < 2$, and undefined elsewhere:

```
f(x) = 0<=x && x<1 ? sin(x) : 1<=x && x<2 ? 1/x : 1/0
plot f(x)
```

Note that **gnuplot** quietly ignores undefined values, so the final branch of the function ($1/0$) will produce no plottable points. Note also that $f(x)$ will be plotted as a continuous function across the discontinuity if a line style is used. To plot it discontinuously, create separate functions for the two pieces. (Parametric functions are also useful for this purpose.)

For data in a file, plot the average of the data in columns 2 and 3 against the datum in column 1, but only if the datum in column 4 is non-negative:

```
plot 'file' using 1:( $4<0 ? 1/0 : ($2+$3)/2 )
```

For an explanation of the **using** syntax, please see **plot datafile using** (p. 85).

Summation

A summation expression has the form

```
sum [<var> = <start> : <end>] <expression>
```

<var> is treated as an integer variable that takes on successive integral values from <start> to <end>. For each of these, the current value of <expression> is added to a running total whose final value becomes the value of the summation expression. Examples:

```
print sum [i=1:10] i
55.
# Equivalent to plot 'data' using 1:($2+$3+$4+$5+$6+...)
plot 'data' using 1 : (sum [col=2:MAXCOL] column(col))
```

It is not necessary that <expression> contain the variable <var>. Although <start> and <end> can be specified as variables or expressions, their value cannot be changed dynamically as a side-effect of carrying out the summation. If <end> is less than <start> then the value of the summation is zero.

Gnuplot-defined variables

Gnuplot maintains a number of read-only variables that reflect the current internal state of the program and the most recent plot. These variables begin with the prefix "GPVAL_". Examples include GPVAL_TERM, GPVAL_X_MIN, GPVAL_X_MAX, GPVAL_Y_MIN. Type **show variables all** to display the complete list and current values. Values related to axes parameters (ranges, log base) are values used during the last plot, not those currently **set**.

Example: To calculate the fractional screen coordinates of the point [X,Y]

```
GRAPH_X = (X - GPVAL_X_MIN) / (GPVAL_X_MAX - GPVAL_X_MIN)
GRAPH_Y = (Y - GPVAL_Y_MIN) / (GPVAL_Y_MAX - GPVAL_Y_MIN)
SCREEN_X = GPVAL_TERM_XMIN + GRAPH_X * (GPVAL_TERM_XMAX - GPVAL_TERM_XMIN)
SCREEN_Y = GPVAL_TERM_YMIN + GRAPH_Y * (GPVAL_TERM_YMAX - GPVAL_TERM_YMIN)
FRAC_X = SCREEN_X * GPVAL_TERM_SCALE / GPVAL_TERM_XSIZE
FRAC_Y = SCREEN_Y * GPVAL_TERM_SCALE / GPVAL_TERM_YSIZE
```

The read-only variable `GPVAL_ERRNO` is set to a non-zero value if any gnuplot command terminates early due to an error. The most recent error message is stored in the string variable `GPVAL_ERRMSG`. Both `GPVAL_ERRNO` and `GPVAL_ERRMSG` can be cleared using the command **reset errors**.

Interactive terminals with **mouse** functionality maintain read-only variables with the prefix `"MOUSE_"`. See **mouse variables** (p. 38) for details.

The **fit** mechanism uses several variables with names that begin `"FIT_"`. It is safest to avoid using such names. When using **set fit errorvariables**, the error for each fitted parameter will be stored in a variable named like the parameter, but with `"_err"` appended. See the documentation on **fit** (p. 62) and **set fit** (p. 115) for details.

See **user-defined variables** (p. 30), **reset errors** (p. 97), **mouse variables** (p. 38), and **fit** (p. 62).

User-defined variables and functions

New user-defined variables and functions of one through twelve variables may be declared and used anywhere, including on the **plot** command itself.

User-defined function syntax:

```
<func-name>( <dummy1> {,<dummy2>} ... {,<dummy12>} ) = <expression>
```

where `<expression>` is defined in terms of `<dummy1>` through `<dummy12>`.

User-defined variable syntax:

```
<variable-name> = <constant-expression>
```

Examples:

```
w = 2
q = floor(tan(pi/2 - 0.1))
f(x) = sin(w*x)
sinc(x) = sin(pi*x)/(pi*x)
delta(t) = (t == 0)
ramp(t) = (t > 0) ? t : 0
min(a,b) = (a < b) ? a : b
comb(n,k) = n!/(k!*(n-k)!)
len3d(x,y,z) = sqrt(x*x+y*y+z*z)
plot f(x) = sin(x*a), a = 0.2, f(x), a = 0.4, f(x)

file = "mydata.inp"
file(n) = sprintf("run_%d.dat",n)
```

The final two examples illustrate a user-defined string variable and a user-defined string function.

Note that the variables **pi** (3.14159...) and **NaN** (IEEE "Not a Number") are already defined. You can redefine these to something else if you really need to. The original values can be recovered by setting:

```
NaN = GPVAL_NaN
pi = GPVAL_pi
```

Other variables may be defined under various gnuplot operations like mousing in interactive terminals or fitting; see **gnuplot-defined variables** (p. 29) for details.

You can check for existence of a given variable `V` by the `exists("V")` expression. For example

```
a = 10
if (exists("a")) print "a is defined"
if (!exists("b")) print "b is not defined"
```

Valid names are the same as in most programming languages: they must begin with a letter, but subsequent characters may be letters, digits, or `"_"`.

Each function definition is made available as a special string-valued variable with the prefix `'GPFUN_'`.

Example:

```
set label GPFUN_sinc at graph .05,.95
```

See [show functions](#) (p. 119), [functions](#) (p. 89), [gnuplot-defined variables](#) (p. 29), [macros](#) (p. 40), [value](#) (p. 27).

Fonts

Gnuplot does not provide any fonts of its own. It relies on external font handling, the details of which unfortunately vary from one terminal type to another. Brief documentation of font mechanisms that apply to more than one terminal type is given here. For information on font use by other individual terminals, see the documentation for that terminal.

Although it is possible to include non-alphabetic symbols by temporarily switching to a special font, e.g. the Adobe Symbol font, the preferred method is now to specify the unicode entry point for the desired symbols using their UTF-8 encoding. See [encoding](#) (p. 114) and [locale](#) (p. 129).

Cairo (pdfcairo, pngcairo, epscairo, wxt terminals)

These terminals find and access fonts using the external fontconfig tool set. Please see the

[fontconfig user manual](#).

It is usually sufficient in gnuplot to request a font by a generic name and size, letting fontconfig substitute a similar font if necessary. The following will probably all work:

```
set term pdfcairo font "sans,12"
set term pdfcairo font "Times,12"
set term pdfcairo font "Times-New-Roman,12"
```

Gd (png, gif, jpeg terminals)

Font handling for the png, gif, and jpeg terminals is done by the external library libgd. Five basic fonts are provided directly by libgd. These are **tiny** (5x8 pixels), **small** (6x12 pixels), **medium**, (7x13 Bold), **large** (8x16) or **giant** (9x15 pixels). These fonts cannot be scaled or rotated. Use one of these keywords instead of the **font** keyword. E.g.

```
set term png tiny
```

On most systems libgd also provides access to Adobe Type 1 fonts (*.pfa) and TrueType fonts (*.ttf). You must give the name of the font file, not the name of the font inside it, in the form "<face> {,<pointsize>}". <face> is either the full pathname to the font file, or the first part of a filename in one of the directories listed in the GDFONTPATH environmental variable. That is, 'set term png font "Face"' will look for a font file named either <somedirectory>/Face.ttf or <somedirectory>/Face.pfa. For example, if GDFONTPATH contains `/usr/local/fonts/ttf:/usr/local/fonts/pfa` then the following pairs of commands are equivalent

```
set term png font "arial"
set term png font "/usr/local/fonts/ttf/arial.ttf"
set term png font "Helvetica"
set term png font "/usr/local/fonts/pfa/Helvetica.pfa"
```

To request a default font size at the same time:

```
set term png font "arial,11"
```

Both TrueType and Adobe Type 1 fonts are fully scalable and rotatable. If no specific font is requested in the "set term" command, gnuplot checks the environmental variable GNUPLOT_DEFAULT_GDFONT to see if there is a preferred default font.

Postscript (also encapsulated postscript *.eps)

PostScript font handling is done by the printer or viewing program. Gnuplot can create valid PostScript or encapsulated PostScript (*.eps) even if no fonts at all are installed on your computer. Gnuplot simply refers to the font by name in the output file, and assumes that the printer or viewing program will know how to find or approximate a font by that name.

All PostScript printers or viewers should know about the standard set of Adobe fonts **Times-Roman**, **Helvetica**, **Courier**, and **Symbol**. It is likely that many additional fonts are also available, but the specific set depends on your system or printer configuration. Gnuplot does not know or care about this; the output *.ps or *.eps files that it creates will simply refer to whatever font names you request.

Thus

```
set term postscript eps font "Times-Roman,12"
```

will produce output that is suitable for all printers and viewers.

On the other hand

```
set term postscript eps font "Garamond-Premier-Pro-Italic"
```

will produce an output file that contains valid PostScript, but since it refers to a specialized font, only some printers or viewers will be able to display the specific font that was requested. Most will substitute a different font.

However, it is possible to embed a specific font in the output file so that all printers will be able to use it. This requires that the a suitable font description file is available on your system. Note that some font files require specific licensing if they are to be embedded in this way. See **postscript fontfile (p. ??)** for more detailed description and examples.

Glossary

Throughout this document an attempt has been made to maintain consistency of nomenclature. This cannot be wholly successful because as **gnuplot** has evolved over time, certain command and keyword names have been adopted that preclude such perfection. This section contains explanations of the way some of these terms are used.

A "page" or "screen" or "canvas" is the entire area addressable by **gnuplot**. On a desktop it is a full window; on a plotter, it is a single sheet of paper; in svg mode it is the full monitor screen.

A screen may contain one or more "plots". A plot is defined by an abscissa and an ordinate, although these need not actually appear on it, as well as the margins and any text written therein.

A plot contains one "graph". A graph is defined by an abscissa and an ordinate, although these need not actually appear on it.

A graph may contain one or more "lines". A line is a single function or data set. "Line" is also a plotting style. The word will also be used in sense "a line of text". Presumably the context will remove any ambiguity.

The lines on a graph may have individual names. These may be listed together with a sample of the plotting style used to represent them in the "key", sometimes also called the "legend".

The word "title" occurs with multiple meanings in **gnuplot**. In this document, it will always be preceded by the adjective "plot", "line", or "key" to differentiate among them. A 2D graph may have up to four labeled **axes**. The names of the four axes are "x" for the axis along the bottom border of the plot, "y" for the axis along the left border, "x2" for the top border, and "y2" for the right border. See **axes (p. 74)**.

A 3D graph may have up to three labeled **axes** – "x", "y" and "z". It is not possible to say where on the graph any particular axis will fall because you can change the direction from which the graph is seen with **set view**.

When discussing data files, the term "record" will be resurrected and used to denote a single line of text in the file, that is, the characters between newline or end-of-record characters. A "point" is the datum extracted from a single record. A "datablock" is a set of points from consecutive records, delimited by blank records. A line, when referred to in the context of a data file, is a subset of a datablock. Note that the term "datablock" may also be used when referring to a named inline data block (see **datablocks (p. 81)**).

Iteration

Version 4.6 of gnuplot introduced command iteration and block-structured if/else/while/do constructs. See **if** (p. 70), **while** (p. 183), and **do** (p. 61). Simple iteration is possible inside **plot** or **set** commands. See **plot for** (p. 91). General iteration spanning multiple commands is possible using a block construct as shown below. For a related new feature, see the **summation** (p. 29) expression type. Here is an example using several of these new syntax features:

```
set multiplot layout 2,2
fourier(k, x) = sin(3./2*k)/k * 2./3*cos(k*x)
do for [power = 0:3] {
    TERMS = 10**power
    set title sprintf("%g term Fourier series",TERMS)
    plot 0.5 + sum [k=1:TERMS] fourier(k,x) notitle
}
unset multiplot
```

Linetypes, colors, and styles

In older gnuplot versions, each terminal type provided a set of distinct "linetypes" that could differ in color, in thickness, in dot/dash pattern, or in some combination of color and dot/dash. These colors and patterns were not guaranteed to be consistent across different terminal types although most used the color sequence red/green/blue/magenta/cyan/yellow. You can select this old behaviour via the command **set colorsequence classic**, but by default gnuplot version 5 uses a terminal-independent sequence of 8 colors.

You can further customize the sequence of linetype properties interactively or in an initialization file. See **set linetype** (p. 128). Several sample initialization files are provided in the distribution package.

The current linetype properties for a particular terminal can be previewed by issuing the **test** command after setting the terminal type.

Successive functions or datafiles plotted by a single command will be assigned successive linetypes in the current default sequence. You can override this for any individual function, datafile, or plot element by giving explicit line properties in the plot command.

Examples:

```
plot "foo", "bar"          # plot two files using linetypes 1, 2
plot sin(x) linetype 4     # use linetype color 4
```

In general, colors can be specified using named colors, rgb (red, green, blue) components, hsv (hue, saturation, value) components, or a coordinate along the current pm3d palette.

Examples:

```
plot sin(x) lt rgb "violet"    # one of gnuplot's named colors
plot sin(x) lt rgb "#FF00FF"   # explicit RGB triple in hexadecimal
plot sin(x) lt palette cb -45   # whatever color corresponds to -45
                                # in the current cbrange of the palette
plot sin(x) lt palette frac 0.3 # fractional value along the palette
```

See **colourspec** (p. 33), **show colornames** (p. 108), **hsv** (p. 26), **set palette** (p. 144), **cbrange** (p. 175). See also **set monochrome** (p. 131).

Linetypes also have an associated dot-dash pattern although not all terminal types are capable of using it. Gnuplot version 5 allows you to specify the dot-dash pattern independent of the line color. See **dashtype** (p. 35).

Colourspec

Many commands allow you to specify a linetype with an explicit color.

Syntax:

```
... {linecolor | lc} {"colorname" | <colourspec> | <n>}
... {textcolor | tc} {<colourspec> | {linetype | lt} <n>}
```

where <colourspec> has one of the following forms:

```
rgbcolor "colorname"      # e.g. "blue"
rgbcolor "0xRRGGBB"       # string containing hexadecimal constant
rgbcolor "0xAARRGGBB"     # string containing hexadecimal constant
rgbcolor "#RRGGBB"        # string containing hexadecimal in x11 format
rgbcolor "#AARRGGBB"      # string containing hexadecimal in x11 format
rgbcolor <integer val>     # integer value representing AARRGGBB
rgbcolor variable         # integer value is read from input file
palette frac <val>        # <val> runs from 0 to 1
palette cb <value>        # <val> lies within cbrange
palette z
variable                  # color index is read from input file
bgnd                      # background color
black
```

The "<n>" is the linetype number the color of which is used, see **test** (p. 182).

"colorname" refers to one of the color names built in to gnuplot. For a list of the available names, see **show colornames** (p. 108).

Hexadecimal constants can be given in quotes as "#RRGGBB" or "0xRRGGBB", where RRGGBB represents the red, green, and blue components of the color and must be between 00 and FF. For example, magenta = full-scale red + full-scale blue could be represented by "0xFF00FF", which is the hexadecimal representation of $(255 \ll 16) + (0 \ll 8) + (255)$.

"#AARRGGBB" represents an RGB color with an alpha channel (transparency) value in the high bits. An alpha value of 0 represents a fully opaque color; i.e., "#00RRGGBB" is the same as "#RRGGBB". An alpha value of 255 (FF) represents full transparency. **Note:** This convention for the alpha channel is backwards from that used by the "with rgbalpha" image plot mode in earlier versions of gnuplot.

The color palette is a linear gradient of colors that smoothly maps a single numerical value onto a particular color. Two such mappings are always in effect. **palette frac** maps a fractional value between 0 and 1 onto the full range of the color palette. **palette cb** maps the range of the color axis onto the same palette. See **set cbrange** (p. 175). See also **set colorbox** (p. 107). You can use either of these to select a constant color from the current palette.

"palette z" maps the z value of each plot segment or plot element into the cbrange mapping of the palette. This allows smoothly-varying color along a 3d line or surface. It also allows coloring 2D plots by palette values read from an extra column of data (not all 2D plot styles allow an extra column). There are two special color specifiers: **bgnd** for background color and **black**.

Background color

Most terminals allow you to set an explicit background color for the plot. The special linetype **bgnd** will draw in this color, and **bgnd** is also recognized as a color. Examples:

```
# This will erase a section of the canvas by writing over it in the
# background color
set term wxt background rgb "gray75"
set object 1 rectangle from x0,y0 to x1,y1 fillstyle solid fillcolor bgnd
# This will draw an "invisible" line along the x axis
plot 0 lt bgnd
```

Linecolor variable

lc variable tells the program to use the value read from one column of the input data as a linetype index, and use the color belonging to that linetype. This requires a corresponding additional column in the **using** specifier. Text colors can be set similarly using **tc variable**.

Examples:

```
# Use the third column of data to assign colors to individual points
plot 'data' using 1:2:3 with points lc variable

# A single data file may contain multiple sets of data, separated by two
# blank lines. Each data set is assigned as index value (see 'index')
# that can be retrieved via the 'using' specifier 'column(-2)'.
# See 'pseudocolumns'. This example uses to value in column -2 to
# draw each data set in a different line color.
plot 'data' using 1:2:(column(-2)) with lines lc variable
```

Rgbcolor variable

You can assign a separate color for each data point, line segment, or label in your plot. **lc rgbcolor variable** tells the program to read RGB color information for each line in the data file. This requires a corresponding additional column in the **using** specifier. The extra column is interpreted as a 24-bit packed RGB triple. If the value is provided directly in the data file it is easiest to give it as a hexadecimal value (see **rgbcolor** (p. 33)). Alternatively, the **using** specifier can contain an expression that evaluates to a 24-bit RGB color as in the example below. Text colors are similarly set using **tc rgbcolor variable**.

Example:

```
# Place colored points in 3D at the x,y,z coordinates corresponding to
# their red, green, and blue components
rgb(r,g,b) = 65536 * int(r) + 256 * int(g) + int(b)
splot "data" using 1:2:3:(rgb($1,$2,$3)) with points lc rgb variable
```

Dashtype

In gnuplot version 5 the dash pattern (**dashtype**) is a separate property associated with each line, analogous to **linecolor** or **linewidth**. It is not necessary to place the current terminal in a special mode just to draw dashed lines. I.e. the command **set term <termname> {solid|dashed}** is now ignored.

All lines have the property **dashtype solid** unless you specify otherwise. You can change the default for a particular linetype using the command **set linetype** so that it affects all subsequent commands, or you can include the desired dashtype as part of the **plot** or other command.

Syntax:

```
dashtype N          # predefined dashtype invoked by number
dashtype "pattern"  # string containing a combination of the characters
                    # dot (.) hyphen (-) underscore(_) and space.
dashtype (s1,e1,s2,e2,s3,e3,s4,e4) # dash pattern specified by 1 to 4
                    # numerical pairs <solid length>, <emptyspace length>
```

Example:

```
# Two functions using linetype 1 but distinguished by dashtype
plot f1(x) with lines lt 1 dt solid, f2(x) with lines lt 1 dt 3
```

Some terminals support user-defined dash patterns in addition to whatever set of predefined dash patterns they offer.

Examples:

```
plot f(x) dt 3          # use terminal-specific dash pattern 3
plot f(x) dt "... "     # construct a dash pattern on the spot
plot f(x) dt (2,5,2,15) # numerical representation of the same pattern
set dashtype 11 (2,4,4,7) # define new dashtype to be called by index
plot f(x) dt 11         # plot using our new dashtype
```

If you specify a dash pattern using a string the program will convert this to a sequence of <solid>,<empty> pairs. The command **show dashtype** will show both the original string and the converted numerical sequence.

Linestyles vs linetypes

A **linestyle** is a temporary association of properties `linecolor`, `linewidth`, `dashtype`, and `pointtype`. It is defined using the command **set style line**. Once you have defined a linestyle, you can use it in a `plot` command to control the appearance of one or more plot elements. In other words, it is just like a `linetype` except for its lifetime. Whereas **linetypes** are permanent (they last until you explicitly redefine them), **linestyles** last until the next reset of the graphics state.

Examples:

```
# define a new line style with terminal-independent color cyan,
# linewidth 3, and associated point type 6 (a circle with a dot in it).
set style line 5 lt rgb "cyan" lw 3 pt 6
plot sin(x) with linespoints ls 5          # user-defined line style 5
```

Layers

A gnuplot plot is built up by drawing its various components in a fixed order. This order can be modified by assigning some components to a specific layer using the keywords **behind**, **back**, or **front**. For example, to replace the background color of the plot area you could define a colored rectangle with the attribute **behind**.

```
set object 1 rectangle from graph 0,0 to graph 1,1 fc rgb "gray" behind
```

The order of drawing is

```
behind
back
the plot itself
the plot legend ('key')
front
```

Within each layer elements are drawn in the order

```
objects (rectangles, circles, ellipses, polygons) in numerical order
labels in numerical order
arrows in numerical order
```

In the case of multiple plots on a single page (multiplot mode) this order applies separately to each component plot, not to the multiplot as a whole.

Mouse input

Many terminals allow interaction with the current plot using the mouse. Some also support the definition of hotkeys to activate pre-defined functions by hitting a single key while the mouse focus is in the active plot window. It is even possible to combine mouse input with **batch** command scripts, by invoking the command **pause mouse** and then using the mouse variables returned by mouse clicking as parameters for subsequent scripted actions. See **bind** (p. 36) and **mouse variables** (p. 38). See also the command **set mouse** (p. 131).

Bind

Syntax:

```
bind {allwindows} [<key-sequence>] ["<gnuplot commands>"]
bind <key-sequence> ""
reset bind
```

The **bind** allows defining or redefining a hotkey, i.e. a sequence of gnuplot commands which will be executed when a certain key or key sequence is pressed while the driver's window has the input focus. Note that **bind** is only available if gnuplot was compiled with **mouse** support and it is used by all mouse-capable terminals. A user-specified binding supersedes any builtin bindings, except that <space> and 'q' cannot normally be rebound. For an exception, see **bind space** (p. 38).

Only mouse button 1 can be bound, and only for 2D plots.

You get the list of all hotkeys by typing **show bind** or **bind** or by typing the hotkey 'h' in the graph window.

Key bindings are restored to their default state by **reset bind**.

Note that multikey-bindings with modifiers must be given in quotes.

Normally hotkeys are only recognized when the currently active plot window has focus. **bind allwindows** <key> ... (short form: **bind all** <key> ...) causes the binding for <key> to apply to all gnuplot plot windows, active or not. In this case gnuplot variable MOUSE_KEY_WINDOW is set to the ID of the originating window, and may be used by the bound command.

Examples:

- set bindings:

```
bind a "replot"
bind "ctrl-a" "plot x*x"
bind "ctrl-alt-a" 'print "great"'
bind Home "set view 60,30; replot"
bind all Home 'print "This is window ",MOUSE_KEY_WINDOW'
```

- show bindings:

```
bind "ctrl-a"          # shows the binding for ctrl-a
bind                   # shows all bindings
show bind              # show all bindings
```

- remove bindings:

```
bind "ctrl-alt-a" ""   # removes binding for ctrl-alt-a
                        (note that builtins cannot be removed)
reset bind              # installs default (builtin) bindings
```

- bind a key to toggle something:

```
v=0
bind "ctrl-r" "v=v+1;if(v%2)set term x11 noraise; else set term x11 raise"
```

Modifiers (ctrl / alt) are case insensitive, keys not:

```
ctrl-alt-a == CtrL-altT-a
ctrl-alt-a != ctrl-alt-A
```

List of modifiers (alt == meta):

```
ctrl, alt
```

List of supported special keys:

```
"BackSpace", "Tab", "Linefeed", "Clear", "Return", "Pause", "Scroll_Lock",
"Sys_Req", "Escape", "Delete", "Home", "Left", "Up", "Right", "Down",
"PageUp", "PageDown", "End", "Begin",
"KP_Space", "KP_Tab", "KP_Enter", "KP_F1", "KP_F2", "KP_F3", "KP_F4",
"KP_Home", "KP_Left", "KP_Up", "KP_Right", "KP_Down", "KP_PageUp",
"KP_PageDown", "KP_End", "KP_Begin", "KP_Insert", "KP_Delete", "KP_Equal",
"KP_Multiply", "KP_Add", "KP_Separator", "KP_Subtract", "KP_Decimal",
"KP_Divide",
"KP_1" - "KP_9", "F1" - "F12"
```

The following are window events rather than actual keys

```
"Button1" "Close"
```

See also help for **mouse** (p. 131).

Bind space

If gnuplot was built with configuration option `-enable-raise-console`, then typing `<space>` in the plot window raises gnuplot's command window. This hotkey can be changed to `ctrl-space` by starting gnuplot as `'gnuplot -ctrlq'`, or by setting the XResource `'gnuplot*ctrlq'`. See **x11 command-line-options** (p. ??).

Mouse variables

When **mousing** is active, clicking in the active window will set several user variables that can be accessed from the gnuplot command line. The coordinates of the mouse at the time of the click are stored in `MOUSE_X` `MOUSE_Y` `MOUSE_X2` and `MOUSE_Y2`. The mouse button clicked, and any meta-keys active at that time, are stored in `MOUSE_BUTTON` `MOUSE_SHIFT` `MOUSE_ALT` and `MOUSE_CTRL`. These variables are set to undefined at the start of every plot, and only become defined in the event of a mouse click in the active plot window. To determine from a script if the mouse has been clicked in the active plot window, it is sufficient to test for any one of these variables being defined.

```
plot 'something'
pause mouse
if (exists("MOUSE_BUTTON")) call 'something_else'; \
else print "No mouse click."
```

It is also possible to track keystrokes in the plot window using the mousing code.

```
plot 'something'
pause mouse keypress
print "Keystroke ", MOUSE_KEY, " at ", MOUSE_X, " ", MOUSE_Y
```

When **pause mouse keypress** is terminated by a keypress, then `MOUSE_KEY` will contain the ascii character value of the key that was pressed. `MOUSE_CHAR` will contain the character itself as a string variable. If the pause command is terminated abnormally (e.g. by `ctrl-C` or by externally closing the plot window) then `MOUSE_KEY` will equal -1.

Note that after a zoom by mouse, you can read the new ranges as `GPVAL_X_MIN`, `GPVAL_X_MAX`, `GPVAL_Y_MIN`, and `GPVAL_Y_MAX`, see **gnuplot-defined variables** (p. 29).

Persist

Many gnuplot terminals (aqua, pm, qt, x11, windows, wxt, ...) open separate display windows on the screen into which plots are drawn. The **persist** option tells gnuplot to leave these windows open when the main program exits. It has no effect on non-interactive terminal output. For example if you issue the command

```
gnuplot -persist -e 'plot [-5:5] sinh(x)'
```

gnuplot will open a display window, draw the plot into it, and then exit, leaving the display window containing the plot on the screen. Depending on the terminal type, some mousing operations may still be possible in the persistent window. However operations like zoom/unzoom that require redrawing the plot are generally not possible because the main program has already exited.

You can also specify **persist** or **nopersist** at the time you set a new terminal type. For example

```
set term qt persist size 700,500
```

Plotting

There are four **gnuplot** commands which actually create a plot: **plot**, **splot**, **replot**, and **refresh**. **plot** generates 2D plots, **splot** generates 3D plots (actually 2D projections, of course), and **replot** appends its arguments to the previous **plot** or **splot** and executes the modified command. **refresh** reexecutes the previous **plot** or **splot** command using previously stored data rather than rereading data from a file.

Much of the general information about plotting can be found in the discussion of **plot**; information specific to 3D can be found in the **splot** section.

plot operates in either rectangular or polar coordinates – see **set polar** (p. 149). **splot** operates in Cartesian coordinates, but will accept azimuthal or cylindrical coordinates on input. See **set mapping** (p. 130). **plot** also lets you use each of the four borders – x (bottom), x2 (top), y (left) and y2 (right) – as an independent axis. The **axes** option lets you choose which pair of axes a given function or data set is plotted against. A full complement of **set** commands exists to give you complete control over the scales and labeling of each axis. Some commands have the name of an axis built into their names, such as **set xlabel**. Other commands have one or more axis names as options, such as **set logscale xy**. Commands and options controlling the z axis have no effect on 2D graphs.

splot can plot surfaces and contours in addition to points and/or lines. See **set isosamples** (p. 122) for information about defining the grid for a 3D function. See **splot datafile** (p. 176) for information about the requisite file structure for 3D data. For contours see **set contour** (p. 108), **set cntrlabel** (p. 105), and **set cntrparam** (p. 106).

In **splot**, control over the scales and labels of the axes are the same as with **plot** except that there is also a z axis and labeling the x2 and y2 axes is possible only for pseudo-2D plots created using **set view map**.

Start-up (initialization)

When gnuplot is run, it first looks for a system-wide initialization file **gnuplotrc**. The location of this file is determined when the program is built and is reported by **show loadpath**. The program then looks in the user's HOME directory for a file called **.gnuplot** on Unix-like systems or **GNUPLOT.INI** on other systems. (OS/2 will look for it in the directory named in the environment variable **GNUPLOT**; Windows will use **APPDATA**). Note: The program can be configured to look first in the current directory, but this is not recommended because it is bad security practice.

String constants and string variables

In addition to string constants, most gnuplot commands also accept a string variable, a string expression, or a function that returns a string. For example, the following four methods of creating a plot all result in the same plot title:

```
four = "4"
graph4 = "Title for plot #4"
graph(n) = sprintf("Title for plot #%d",n)

plot 'data.4' title "Title for plot #4"
plot 'data.4' title graph4
plot 'data.4' title "Title for plot #".four
plot 'data.4' title graph(4)
```

Since integers are promoted to strings when operated on by the string concatenation operator, the following method also works:

```
N = 4
plot 'data.'.N title "Title for plot #".N
```

In general, elements on the command line will only be evaluated as possible string variables if they are not otherwise recognizable as part of the normal gnuplot syntax. So the following sequence of commands is legal, although probably should be avoided so as not to cause confusion:

```
plot = "my_datafile.dat"
title = "My Title"
plot plot title title
```

Three binary operators require string operands: the string concatenation operator ".", the string equality operator "eq" and the string inequality operator "ne". The following example will print TRUE.


```
if ("A"."B" eq "AB") print "TRUE"
```

See also the two string formatting functions **gprintf** (p. 117) and **sprintf** (p. 26).

Substrings can be specified by appending a range specifier to any string, string variable, or string-valued function. The range specifier has the form [begin:end], where begin is the index of the first character of the substring and end is the index of the last character of the substring. The first character has index 1. The begin or end fields may be empty, or contain '*', to indicate the true start or end of the original string. E.g. str[:] and str[*:*] both describe the full string str.

Substitution and Command line macros

When a command line to gnuplot is first read, i.e. before it is interpreted or executed, two forms of lexical substitution are performed. These are triggered by the presence of text in backquotes (ascii character 96) or preceded by @ (ascii character 64).

Substitution of system commands in backquotes

Command-line substitution is specified by a system command enclosed in backquotes. This command is spawned and the output it produces replaces the backquoted text on the command line. Some implementations also support pipes; see **plot datafile special-filenames** (p. 83).

Command-line substitution can be used anywhere on the **gnuplot** command line, except inside strings delimited by single quotes.

Example:

This will run the program **leastsq** and replace **leastsq** (including backquotes) on the command line with its output:

```
f(x) = 'leastsq'
```

or, in VMS

```
f(x) = 'run leastsq'
```

These will generate labels with the current time and userid:

```
set label "generated on 'date +%Y-%m-%d' by 'whoami'" at 1,1
set timestamp "generated on %Y-%m-%d by 'whoami'"
```

Substitution of string variables as macros

Substitution of command line macros is disabled by default, but may be enabled using the **set macros** command. If macro substitution is enabled, the character @ is used to trigger substitution of the current value of a string variable into the command line. The text in the string variable may contain any number of lexical elements. This allows string variables to be used as command line macros. Only string constants may be expanded using this mechanism, not string-valued expressions. For example:

```
set macros
style1 = "lines lt 4 lw 2"
style2 = "points lt 3 pt 5 ps 2"
range1 = "using 1:3"
range2 = "using 1:5"
plot "foo" @range1 with @style1, "bar" @range2 with @style2
```

The line containing @ symbols is expanded on input, so that by the time it is executed the effect is identical to having typed in full

```
plot "foo" using 1:3 with lines lt 4 lw 2, \
      "bar" using 1:5 with points lt 3 pt 5 ps 2
```


The function `exists()` may be useful in connection with macro evaluation. The following example checks that `C` can safely be expanded as the name of a user-defined variable:

```
C = "pi"
if (exists(C)) print C, " = ", @C
```

Macro expansion does not occur inside either single or double quotes. However macro expansion does occur inside backquotes.

Macro expansion is handled as the very first thing the interpreter does when looking at a new line of commands and is only done once. Therefore, code like the following will execute correctly:

```
A = "c=1"
@A
```

but this line will not, since the macro is defined on the same line and will not be expanded in time

```
A = "c=1"; @A    # will not expand to c=1
```

For execution of complete commands the **evaluate** command may also be handy.

String variables, macros, and command line substitution

The interaction of string variables, backquotes and macro substitution is somewhat complicated. Backquotes do not block macro substitution, so

```
filename = "mydata.inp"
lines = `wc --lines @filename | sed "s/ .*//"`
```

results in the number of lines in `mydata.inp` being stored in the integer variable `lines`. And double quotes do not block backquote substitution, so

```
mycomputer = "`uname -n`"
```

results in the string returned by the system command **uname -n** being stored in the string variable `mycomputer`.

However, macro substitution is not performed inside double quotes, so you cannot define a system command as a macro and then use both macro and backquote substitution at the same time.

```
machine_id = "uname -n"
mycomputer = "`@machine_id`"  # doesn't work!!
```

This fails because the double quotes prevent `@machine_id` from being interpreted as a macro. To store a system command as a macro and execute it later you must instead include the backquotes as part of the macro itself. This is accomplished by defining the macro as shown below. Notice that the `sprintf` format nests all three types of quotes.

```
machine_id = sprintf('`uname -n`')
mycomputer = @machine_id
```

Syntax

Options and any accompanying parameters are separated by spaces whereas lists and coordinates are separated by commas. Ranges are separated by colons and enclosed in brackets `[]`, text and file names are enclosed in quotes, and a few miscellaneous things are enclosed in parentheses.

Commas are used to separate coordinates on the **set** commands **arrow**, **key**, and **label**; the list of variables being fitted (the list after the **via** keyword on the **fit** command); lists of discrete contours or the loop parameters which specify them on the **set cnrparam** command; the arguments of the **set** commands **dgrid3d**, **dummy**, **isosamples**, **offsets**, **origin**, **samples**, **size**, **time**, and **view**; lists of tics or the loop

parameters which specify them; the offsets for titles and axis labels; parametric functions to be used to calculate the x, y, and z coordinates on the **plot**, **replot** and **splot** commands; and the complete sets of keywords specifying individual plots (data sets or functions) on the **plot**, **replot** and **splot** commands.

Parentheses are used to delimit sets of explicit tics (as opposed to loop parameters) and to indicate computations in the **using** filter of the **fit**, **plot**, **replot** and **splot** commands.

(Parentheses and commas are also used as usual in function notation.)

Square brackets are used to delimit ranges given in **set**, **plot** or **splot** commands.

Colons are used to separate extrema in **range** specifications (whether they are given on **set**, **plot** or **splot** commands) and to separate entries in the **using** filter of the **plot**, **replot**, **splot** and **fit** commands.

Semicolons are used to separate commands given on a single command line.

Curly braces are used in the syntax for enhanced text mode and to delimit blocks in if/then/else statements. They are also used to denote complex numbers: $\{3,2\} = 3 + 2i$.

The EEPIC, Imagen, Uniplex, LaTeX, and TPIC drivers allow a newline to be specified by `\\` in a single-quoted string or `\\\\` in a double-quoted string.

Quote Marks

Gnuplot uses three forms of quote marks for delimiting text strings, double-quote (ascii 34), single-quote (ascii 39), and backquote (ascii 96).

Filenames may be entered with either single- or double-quotes. In this manual the command examples generally single-quote filenames and double-quote other string tokens for clarity.

String constants and text strings used for labels, titles, or other plot elements may be enclosed in either single quotes or double quotes. Further processing of the quoted text depends on the choice of quote marks.

Backslash processing of special characters like `\n` (newline) and `\345` (octal character code) is performed for double-quoted strings. In single-quoted strings, backslashes are just ordinary characters. To get a single-quote (ascii 39) in a single-quoted string, it has to be doubled. Thus the strings `"d\" s' b\\`" and `'d" s'' b\'` are completely equivalent.

Text justification is the same for each line of a multi-line string. Thus the center-justified string

```
"This is the first line of text.\nThis is the second line."
```

will produce

```
      This is the first line of text.
      This is the second line.
```

but

```
'This is the first line of text.\nThis is the second line.'
```

will produce

```
      This is the first line of text.\nThis is the second line.
```

Enhanced text processing is performed for both double-quoted text and single-quoted text, but only by terminals supporting this mode. See **enhanced text** (p. 22).

Back-quotes are used to enclose system commands for substitution into the command line. See **substitution** (p. 40).

Time/Date data

gnuplot supports the use of time and/or date information as input data. This feature is activated by the commands **set xdata time**, **set ydata time**, etc.

Internally all times and dates are converted to the number of seconds from the year 1970. The command **set timefmt** defines the default format for all inputs: data files, ranges, tics, label positions – anything that

accepts a time data value defaults to receiving it in this format. Only one default format can be in effect at a given time. Thus if both x and y data in a file are time/date, by default they are interpreted in the same format. However this default can be replaced when reading any particular file or column of input using the **timecolumn** function in the corresponding **using** specifier.

The conversion to and from seconds assumes Universal Time (which is the same as Greenwich Standard Time). There is no provision for changing the time zone or for daylight savings. If all your data refer to the same time zone (and are all either daylight or standard) you don't need to worry about these things. But if the absolute time is crucial for your application, you'll need to convert to UT yourself.

Commands like **show xrange** will re-interpret the integer according to **timefmt**. If you change **timefmt**, and then **show** the quantity again, it will be displayed in the new **timefmt**. For that matter, if you reset the data type flag for that axis (e.g. **set xdata**), the quantity will be shown in its numerical form.

The commands **set format** or **set tics format** define the format that will be used for tic labels, whether or not input for the specified axis is time/date.

If time/date information is to be plotted from a file, the **using** option *must* be used on the **plot** or **splot** command. These commands simply use white space to separate columns, but white space may be embedded within the time/date string. If you use tabs as a separator, some trial-and-error may be necessary to discover how your system treats them.

The **time** function can be used to get the current system time. This value can be converted to a date string with the **strftime** function, or it can be used in conjunction with **timecolumn** to generate relative time/date plots. The type of the argument determines what is returned. If the argument is an integer, **time** returns the current time as an integer, in seconds from 1 Jan 1970. If the argument is real (or complex), the result is real as well. The precision of the fractional (sub-second) part depends on your operating system. If the argument is a string, it is assumed to be a format string, and it is passed to **strftime** to provide a formatted time/date string.

The following example demonstrates time/date plotting.

Suppose the file "data" contains records like

```
03/21/95 10:00 6.02e23
```

This file can be plotted by

```
set xdata time
set timefmt "%m/%d/%y"
set xrange ["03/21/95":"03/22/95"]
set format x "%m/%d"
set timefmt "%m/%d/%y %H:%M"
plot "data" using 1:3
```

which will produce xtic labels that look like "03/21".

Gnuplot tracks time to millisecond precision. Time formats have been modified to match this. Example: print the current time to msec precision

```
print strftime("%H:%M:%.3S %d-%b-%Y",time(0.0))
18:15:04.253 16-Apr-2011
```

See **time.specifiers** (p. 118).

Part II

Plotting styles

Many plotting styles are available in gnuplot. They are listed alphabetically below. The commands **set style data** and **set style function** change the default plotting style for subsequent **plot** and **splot** commands.

You can also specify the plot style explicitly as part of the **plot** or **splot** command. If you want to mix plot styles within a single plot, you must specify the plot style for each component.

Example:

```
plot 'data' with boxes, sin(x) with lines
```

Each plot style has its own expected set of data entries in a data file. For example, by default the **lines** style expects either a single column of y values (with implicit x ordering) or a pair of columns with x in the first and y in the second. For more information on how to fine-tune how columns in a file are interpreted as plot data, see **using** (p. 85).

Boxerrorbars

The **boxerrorbars** style is only relevant to 2D data plotting. It is a combination of the **boxes** and **yerrorbars** styles. It requires 3, 4, or 5 columns of data. An additional (4th, 5th or 6th) input column may be used to provide variable (per-datapoint) color information (see **linecolor** (p. 33) and **rgbcolor variable** (p. 35)). The error bar will be drawn in the same color as the border of the box.

```
3 columns:  x  y  ydelta
4 columns:  x  y  ydelta xdelta      # boxwidth != -2
4 columns:  x  y  ylow  yhigh       # boxwidth == -2
5 columns:  x  y  ylow  yhigh  xdelta
```

The boxwidth will come from the fourth column if the y errors are given as "ydelta" and the boxwidth was not previously set to -2.0 (**set boxwidth -2.0**) or from the fifth column if the y errors are in the form of "ylow yhigh". The special case **boxwidth = -2.0** is for four-column data with y errors in the form "ylow yhigh". In this case the boxwidth will be calculated so that each box touches the adjacent boxes. The width will also be calculated in cases where three-column data are used.

The box height is determined from the y error in the same way as it is for the **yerrorbars** style — either from y-ydelta to y+ydelta or from ylow to yhigh, depending on how many data columns are provided.

Boxes

The **boxes** style is only relevant to 2D plotting. It draws a box centered about the given x coordinate that extends from the x axis (not from the graph border) to the given y coordinate. It uses 2 or 3 columns of basic data. Additional input columns may be used to provide information such as variable line or fill color (see **rgbcolor variable** (p. 35)).

```
2 columns:  x  y
3 columns:  x  y  x_width
```

The width of the box is obtained in one of three ways. If the input data has a third column, this will be used to set the width of the box. If not, if a width has been set using the **set boxwidth** command, this will be used. If neither of these is available, the width of each box will be calculated automatically so that it touches the adjacent boxes.

The interior of the boxes is drawn according to the current fillstyle. See **set style fill** (p. 155) for details. Alternatively a new fillstyle may be specified in the plot command. For fillstyle **empty** the box is not filled. For fillstyle **solid** the box is filled with a solid rectangle of the current drawing color. An optional fillstyle

parameter controls the fill density; it runs from 0 (background color) to 1 (current drawing color). For fillstyle **pattern** the box is filled in the current drawing color with a pattern.

Examples:

To plot a data file with solid filled boxes with a small vertical space separating them (bargraph):

```
set boxwidth 0.9 relative
set style fill solid 1.0
plot 'file.dat' with boxes
```

To plot a sine and a cosine curve in pattern-filled boxes style:

```
set style fill pattern
plot sin(x) with boxes, cos(x) with boxes
```

The sin plot will use pattern 0; the cos plot will use pattern 1. Any additional plots would cycle through the patterns supported by the terminal driver.

To specify explicit fillstyles for each dataset:

```
plot 'file1' with boxes fs solid 0.25, \
      'file2' with boxes fs solid 0.50, \
      'file3' with boxes fs solid 0.75, \
      'file4' with boxes fill pattern 1, \
      'file5' with boxes fill empty
```

Boxplot

Boxplots are a common way to represent a statistical distribution of values. Quartile boundaries are determined such that 1/4 of the points have a value equal or less than the first quartile boundary, 1/2 of the points have a value equal or less than the second quartile (median) value, etc. A box is drawn around the region between the first and third quartiles, with a horizontal line at the median value. Whiskers extend from the box to user-specified limits. Points that lie outside these limits are drawn individually.

Examples

```
# Place a boxplot at x coordinate 1.0 representing the y values in column 5
plot 'data' using (1.0):5
```

```
# Same plot but suppress outliers and force the width of the boxplot to 0.3
set style boxplot nooutliers
plot 'data' using (1.0):5:(0.3)
```

By default only one boxplot is produced that represents all y values from the second column of the using specification. However, an additional (fourth) column can be added to the specification. If present, the values of that column will be interpreted as the discrete levels of a factor variable. As many boxplots will be drawn as there are levels in the factor variable. The separation between these boxplots is 1.0 by default, but it can be changed by **set style boxplot separation**. By default, the value of the factor variable is shown as a tic label below (or above) each boxplot.

Example

```
# Suppose that column 2 of 'data' contains either "control" or "treatment"
# The following example produces two boxplots, one for each level of the
# factor
plot 'data' using (1.0):5:(0):2
```

The default width of the box can be set via **set boxwidth <width>** or may be specified as an optional 3rd column in the **using** clause of the plot command. The first and third columns (x coordinate and width) are normally provided as constants rather than as data columns.

By default the whiskers extend from the ends of the box to the most distant point whose y value lies within 1.5 times the interquartile range. By default outliers are drawn as circles (point type 7). The width of the bars at the end of the whiskers may be controlled using **set bars**.

These default properties may be changed using the **set style boxplot** command. See **set style boxplot** (p. 154), **bars** (p. 102), **boxwidth** (p. 104), **fillstyle** (p. 155), **candlesticks** (p. 46).

Boxxyerrorbars

The **boxxyerrorbars** style is only relevant to 2D data plotting. It is similar to the **xyerrorbars** style except that it draws rectangular areas rather than simple crosses. It uses either 4 or 6 basic columns of input data. Additional input columns may be used to provide information such as variable line or fill color (see **rgbcolor variable** (p. 35)).

```
4 columns:  x  y  xdelta ydelta
6 columns:  x  y  xlow  xhigh ylow yhigh
```

The box width and height are determined from the x and y errors in the same way as they are for the **xyerrorbars** style — either from xlow to xhigh and from ylow to yhigh, or from x-xdelta to x+xdelta and from y-ydelta to y+ydelta, depending on how many data columns are provided.

An additional (5th or 7th) input column may be used to provide variable (per-datapoint) color information (see **linecolor** (p. 33) and **rgbcolor variable** (p. 35)).

The interior of the boxes is drawn according to the current fillstyle. See **set style fill** (p. 155) and **boxes** (p. 44) for details. Alternatively a new fillstyle may be specified in the plot command.

Candlesticks

The **candlesticks** style can be used for 2D data plotting of financial data or for generating box-and-whisker plots of statistical data. The symbol is a rectangular box, centered horizontally at the x coordinate and limited vertically by the opening and closing prices. A vertical line segment at the x coordinate extends up from the top of the rectangle to the high price and another down to the low. The vertical line will be unchanged if the low and high prices are interchanged.

Five columns of basic data are required:

```
financial data:  date  open  low  high  close
whisker plot:   x    box_min whisker_min whisker_high box_high
```

The width of the rectangle can be controlled by the **set boxwidth** command. For backwards compatibility with earlier gnuplot versions, when the boxwidth parameter has not been set then the width of the candlestick rectangle is controlled by **set bars <width>**.

Alternatively, an explicit width for each box-and-whiskers grouping may be specified in an optional 6th column of data. The width must be given in the same units as the x coordinate.

An additional (6th, or 7th if the 6th column is used for width data) input column may be used to provide variable (per-datapoint) color information (see **linecolor** (p. 33) and **rgbcolor variable** (p. 35)).

By default the vertical line segments have no crossbars at the top and bottom. If you want crossbars, which are typically used for box-and-whisker plots, then add the keyword **whiskerbars** to the plot command. By default these whiskerbars extend the full horizontal width of the candlestick, but you can modify this by specifying a fraction of the full width.

The usual convention for financial data is that the rectangle is empty if (open < close) and solid fill if (close < open). This is the behavior you will get if the current fillstyle is set to "empty". See **fillstyle** (p. 155). If you set the fillstyle to solid or pattern, then this will be used for all boxes independent of open and close values. See also **set bars** (p. 102) and **financebars** (p. 49). See also the

candlestick

and

finance

demos.

Note: To place additional symbols, such as the median value, on a box-and-whisker plot requires additional plot commands as in this example:

```
# Data columns: X Min 1stQuartile Median 3rdQuartile Max
set bars 4.0
set style fill empty
plot 'stat.dat' using 1:3:2:6:5 with candlesticks title 'Quartiles', \
    '' using 1:4:4:4:4 with candlesticks lt -1 notitle

# Plot with crossbars on the whiskers, crossbars are 50% of full width
plot 'stat.dat' using 1:3:2:6:5 with candlesticks whiskerbars 0.5
```

See `set boxwidth` (p. 104), `set bars` (p. 102), `set style fill` (p. 155), and `boxplot` (p. 45).

Circles

The **circles** style plots a circle with an explicit radius at each data point. If three columns of data are present, they are interpreted as x, y, radius. The radius is always interpreted in the units of the plot's horizontal axis (x or x2). The scale on y and the aspect ratio of the plot are both ignored. If only two columns are present, the radius is taken from `set style circle`. In this case the radius may be given in graph or screen coordinates.

By default a full circle will be drawn. It is possible to plot arc segments instead of full circles by specifying a start and end angle in the 4th and 5th columns. An optional 4th or 6th column can specify per-circle color. The start and end angles of the circle segments must be specified in degrees. See `set style circle` (p. 157) and `set style fill` (p. 155).

Examples:

```
# draws circles whose area is proportional to the value in column 3
set style fill transparent solid 0.2 noborder
plot 'data' using 1:2:(sqrt($3)) with circles, \
    'data' using 1:2 with linespoints

# draws Pac-men instead of circles
plot 'data' using 1:2:(10):(40):(320) with circles

# draw a pie chart with inline data
set xrange [-15:15]
set style fill transparent solid 0.9 noborder
plot '-' using 1:2:3:4:5:6 with circles lc var
0 0 5 0 30 1
0 0 5 30 70 2
0 0 5 70 120 3
0 0 5 120 230 4
0 0 5 230 360 5
e
```

The result is similar to using a **points** plot with variable size points and pointstyle 7, except that the circles will scale with the x axis range. See also `set object circle` (p. 138) and `fillstyle` (p. 155).

Ellipses

The **ellipses** style plots an ellipse at each data point. This style is only relevant for 2D plotting. Each ellipse is described in terms of its center, major and minor diameters, and the angle between its major diameter and the x axis.

```

2 columns: x y
3 columns: x y major_diam
4 columns: x y major_diam minor_diam
5 columns: x y major_diam minor_diam angle

```

If only two input columns are present, they are taken as the coordinates of the centers, and the ellipses will be drawn with the default extent (see **set style ellipse** (p. 158)). The orientation of the ellipse, which is defined as the angle between the major diameter and the plot's x axis, is taken from the default ellipse style (see **set style ellipse** (p. 158)). If three input columns are provided, the third column is used for both diameters. The orientation angle defaults to zero. If four columns are present, they are interpreted as x, y, major diameter, minor diameter. Note that these are diameters, not radii. An optional 5th column may be used to specify the orientation angle in degrees. The ellipses will also be drawn with their default extent if either of the supplied diameters in the 3-4-5 column form is negative.

In all of the above cases, optional variable color data may be given in an additional last (3th, 4th, 5th or 6th) column. See **colorespec** (p. 33) for further information.

By default, the major diameter is interpreted in the units of the plot's horizontal axis (x or x2) while the minor diameter in that of the vertical (y or y2). This implies that if the x and y axis scales are not equal, then the major/minor diameter ratio will no longer be correct after rotation. This behavior can be changed with the **units** keyword, however.

There are three alternatives: if **units xy** is included in the plot specification, the axes will be scaled as described above. **units xx** ensures that both diameters are interpreted in units of the x axis, while **units yy** means that both diameters are interpreted in units of the y axis. In the latter two cases the ellipses will have the correct aspect ratio, even if the plot is resized.

If **units** is omitted, the default setting will be used, which is equivalent to **units xy**. This can be redefined by **set style ellipse**.

Example (draws ellipses, cycling through the available line types):

```
plot 'data' using 1:2:3:4:(0):0 with ellipses
```

See also **set object ellipse** (p. 137), **set style ellipse** (p. 158) and **fillstyle** (p. 155).

Dots

The **dots** style plots a tiny dot at each point; this is useful for scatter plots with many points. Either 1 or 2 columns of input data are required in 2D. Three columns are required in 3D.

For some terminals (post, pdf) the size of the dot can be controlled by changing the linewidth.

```

1 column      y          # x is row number
2 columns:    x  y
3 columns:    x  y  z     # 3D only (splot)

```

Filledcurves

The **filledcurves** style is only relevant to 2D plotting. Three variants are possible. The first two variants require either a function or two columns of input data, and may be further modified by the options listed below.

Syntax:

```
plot ... with filledcurves [option]
```

where the option can be one of the following

```

[closed | {above | below}
{x1 | x2 | y | r}[=<a>] | xy=<x>,<y>]

```


The first variant, **closed**, treats the curve itself as a closed polygon. This is the default if there are two columns of input data.

The second variant is to fill the area between the curve and a given axis, a horizontal or vertical line, or a point.

```
filledcurves closed    ... just filled closed curve,
filledcurves x1        ... x1 axis,
filledcurves x2        ... x2 axis, etc for y1 and y2 axes,
filledcurves y=42      ... line at y=42, i.e. parallel to x axis,
filledcurves xy=10,20  ... point 10,20 of x1,y1 axes (arc-like shape).
filledcurves above r=1.5 the area of a polar plot outside radius 1.5
```

The third variant requires three columns of input data: the x coordinate and two y coordinates corresponding to two curves sampled at the same set of x coordinates; the area between the two curves is filled. This is the default if there are three or more columns of input data.

```
3 columns:  x  y1  y2
```

Example of filling the area between two input curves.

```
fill between curves demo.
plot 'data' using 1:2:3 with filledcurves
```

The **above** and **below** options apply both to commands of the form

```
... filledcurves above {x1|x2|y1|y2|r}=<val>
```

and to commands of the form

```
... using 1:2:3 with filledcurves below
```

In either case the option limits the filled area to one side of the bounding line or curve.

Note: Not all terminal types support this plotting mode.

Zooming a filled curve drawn from a datafile may produce empty or incorrect areas because gnuplot is clipping points and lines, and not areas.

If the values of <a>, <x>, <y> are out of the drawing boundary, then they are moved to the graph boundary. Then the actually filled area in the case of option xy=<x>,<y> will depend on xrange and yrange.

Financebars

The **financebars** style is only relevant for 2D data plotting of financial data. It requires 1 x coordinate (usually a date) and 4 y values (prices).

```
5 columns:  date  open  low  high  close
```

An additional (6th) input column may be used to provide variable (per-record) color information (see **linecolor** (p. 33) and **rgbcolor variable** (p. 35)).

The symbol is a vertical line segment, located horizontally at the x coordinate and limited vertically by the high and low prices. A horizontal tic on the left marks the opening price and one on the right marks the closing price. The length of these tics may be changed by **set bars**. The symbol will be unchanged if the high and low prices are interchanged. See **set bars** (p. 102) and **candlesticks** (p. 46), and also the

```
finance demo.
```

Fsteps

The **fsteps** style is only relevant to 2D plotting. It connects consecutive points with two line segments: the first from (x1,y1) to (x1,y2) and the second from (x1,y2) to (x2,y2). The input column requires are the same

as for plot styles **lines** and **points**. The difference between **fsteps** and **steps** is that **fsteps** traces first the change in y and then the change in x. **steps** traces first the change in x and then the change in y.

See also

[steps demo.](#)

Fillsteps

The **fillsteps** style is exactly like **steps** except that the area between the curve and $y=0$ is filled in the current fill style. See **steps** (p. 56).

Histeps

The **histeps** style is only relevant to 2D plotting. It is intended for plotting histograms. Y-values are assumed to be centered at the x-values; the point at x_1 is represented as a horizontal line from $((x_0+x_1)/2, y_1)$ to $((x_1+x_2)/2, y_1)$. The lines representing the end points are extended so that the step is centered on at x. Adjacent points are connected by a vertical line at their average x, that is, from $((x_1+x_2)/2, y_1)$ to $((x_1+x_2)/2, y_2)$. The input column requires are the same as for plot styles **lines** and **points**.

If **autoscale** is in effect, it selects the xrange from the data rather than the steps, so the end points will appear only half as wide as the others. See also

[steps demo.](#)

Histograms

The **histograms** style is only relevant to 2D plotting. It produces a bar chart from a sequence of parallel data columns. Each element of the **plot** command must specify a single input data source (e.g. one column of the input file), possibly with associated tic values or key titles. Four styles of histogram layout are currently supported.

```
set style histogram clustered {gap <gapsize>}
set style histogram errorbars {gap <gapsize>} {<linewidth>}
set style histogram rowstacked
set style histogram columnstacked
set style histogram {title font "name,size" tc <colourspec>}
```

The default style corresponds to **set style histogram clustered gap 2**. In this style, each set of parallel data values is collected into a group of boxes clustered at the x-axis coordinate corresponding to their sequential position (row #) in the selected datafile columns. Thus if $\langle n \rangle$ datacolumns are selected, the first cluster is centered about $x=1$, and contains $\langle n \rangle$ boxes whose heights are taken from the first entry in the corresponding $\langle n \rangle$ data columns. This is followed by a gap and then a second cluster of boxes centered about $x=2$ corresponding to the second entry in the respective data columns, and so on. The default gap width of 2 indicates that the empty space between clusters is equivalent to the width of 2 boxes. All boxes derived from any one column are given the same fill color and/or pattern (see **set style fill** (p. 155)).

Each cluster of boxes is derived from a single row of the input data file. It is common in such input files that the first element of each row is a label. Labels from this column may be placed along the x-axis underneath the appropriate cluster of boxes with the **xticlabels** option to **using**.

The **errorbars** style is very similar to the **clustered** style, except that it requires additional columns of input for each entry. The first column holds the height (y value) of that box, exactly as for the **clustered** style.

2 columns:	y yerr	bar extends from y-yerr to y+err
3 columns:	y ymin ymax	bar extends from ymin to ymax

The appearance of the error bars is controlled by the current value of **set bars** and by the optional `<linewidth>` specification.

Two styles of stacked histogram are supported, chosen by the command **set style histogram {rowstacked|columnstacked}**. In these styles the data values from the selected columns are collected into stacks of boxes. Positive values stack upwards from $y=0$; negative values stack downwards. Mixed positive and negative values will produce both an upward stack and a downward stack. The default stacking mode is **rowstacked**.

The **rowstacked** style places a box resting on the x-axis for each data value in the first selected column; the first data value results in a box at $x=1$, the second at $x=2$, and so on. Boxes corresponding to the second and subsequent data columns are layered on top of these, resulting in a stack of boxes at $x=1$ representing the first data value from each column, a stack of boxes at $x=2$ representing the second data value from each column, and so on. All boxes derived from any one column are given the same fill color and/or pattern (see **set style fill** (p. 155)).

The **columnstacked** style is similar, except that each stack of boxes is built up from a single data column. Each data value from the first specified column yields a box in the stack at $x=1$, each data value from the second specified column yields a box in the stack at $x=2$, and so on. In this style the color of each box is taken from the row number, rather than the column number, of the corresponding data field.

Box widths may be modified using the **set boxwidth** command. Box fill styles may be set using the **set style fill** command.

Histograms always use the x1 axis, but may use either y1 or y2. If a plot contains both histograms and other plot styles, the non-histogram plot elements may use either the x1 or the x2 axis.

Examples: Suppose that the input file contains data values in columns 2, 4, 6, ... and error estimates in columns 3, 5, 7, ... This example plots the values in columns 2 and 4 as a histogram of clustered boxes (the default style). Because we use iteration in the plot command, any number of data columns can be handled in a single command. See **plot for** (p. 91).

```
set boxwidth 0.9 relative
set style data histograms
set style histogram cluster
set style fill solid 1.0 border lt -1
plot for [COL=2:4:2] 'file.dat' using COL
```

This will produce a plot with clusters of two boxes (vertical bars) centered at each integral value on the x axis. If the first column of the input file contains labels, they may be placed along the x-axis using the variant command

```
plot for [COL=2:4:2] 'file.dat' using COL:xticlabels(1)
```

If the file contains both magnitude and range information for each value, then error bars can be added to the plot. The following commands will add error bars extending from $(y-<error>)$ to $(y+<error>)$, capped by horizontal bar ends drawn the same width as the box itself. The error bars and bar ends are drawn with linewidth 2, using the border linetype from the current fill style.

```
set bars fullwidth
set style fill solid 1 border lt -1
set style histogram errorbars gap 2 lw 2
plot for [COL=2:4:2] 'file.dat' using COL:COL+1
```

This shows how to plot the same data as a rowstacked histogram. Just to be different, this example lists the separate columns explicitly rather than using iteration.

```
set style histogram rowstacked
plot 'file.dat' using 2, '' using 4:xtic(1)
```

This will produce a plot in which each vertical bar corresponds to one row of data. Each vertical bar contains a stack of two segments, corresponding in height to the values found in columns 2 and 4 of the datafile.

Finally, the commands

```
set style histogram columnstacked
plot 'file.dat' using 2, '' using 4
```

will produce two vertical stacks, one for each column of data. The stack at x=1 will contain a box for each entry in column 2 of the datafile. The stack at x=2 will contain a box for each parallel entry in column 4 of the datafile.

Because this interchanges gnuplot's usual interpretation of input rows and columns, the specification of key titles and x-axis tic labels must also be modified accordingly. See the comments given below.

```
set style histogram columnstacked
plot '' u 5:key(1)          # uses first column to generate key titles
plot '' u 5 title columnhead # uses first row to generate xtic labels
```

Note that the two examples just given present exactly the same data values, but in different formats.

Newhistogram

Syntax:

```
newhistogram {"<title>" {font "name,size"} {tc <colourspec>}}
              {lt <linetype>} {fs <fillstyle>} {at <x-coord>}}
```

More than one set of histograms can appear in a single plot. In this case you can force a gap between them, and a separate label for each set, by using the **newhistogram** command. For example

```
set style histogram cluster
plot newhistogram "Set A", 'a' using 1, '' using 2, '' using 3, \
     newhistogram "Set B", 'b' using 1, '' using 2, '' using 3
```

The labels "Set A" and "Set B" will appear beneath the respective sets of histograms, under the overall x axis label.

The newhistogram command can also be used to force histogram coloring to begin with a specific color (linetype). By default colors will continue to increment successively even across histogram boundaries. Here is an example using the same coloring for multiple histograms

```
plot newhistogram "Set A" lt 4, 'a' using 1, '' using 2, '' using 3, \
     newhistogram "Set B" lt 4, 'b' using 1, '' using 2, '' using 3
```

Similarly you can force the next histogram to begin with a specified fillstyle. If the fillstyle is set to **pattern**, then the pattern used for filling will be incremented automatically.

The **at <x-coord>** option sets the x coordinate position of the following histogram to <x-coord>. For example

```
set style histogram cluster
set style data histogram
set style fill solid 1.0 border -1
set xtic 1 offset character 0,0.3
plot newhistogram "Set A", \
     'file.dat' u 1 t 1, '' u 2 t 2, \
     newhistogram "Set B" at 8, \
     'file.dat' u 2 t 2, '' u 2 t 2
```

will position the second histogram to start at x=8.

Automated iteration over multiple columns

If you want to create a histogram from many columns of data in a single file, it is very convenient to use the plot iteration feature. See **plot for** (p. 91). For example, to create stacked histograms of the data in columns 3 through 8

```
set style histogram columnstacked
plot for [i=3:8] "datafile" using i title columnhead
```

Image

The **image**, **rgbimage**, and **rgbalpha** plotting styles all project a uniformly sampled grid of data values onto a plane in either 2D or 3D. The input data may be an actual bitmapped image, perhaps converted from a standard format such as PNG, or a simple array of numerical values.

This figure illustrates generation of a heat map from an array of scalar values. The current palette is used to map each value onto the color assigned to the corresponding pixel.

```
plot '-' matrix with image
5 4 3 1 0
2 2 0 0 1
0 0 0 1 0
0 1 2 4 3
e
e
```

Each pixel (data point) of the input 2D image will become a rectangle or parallelepiped in the plot. The coordinates of each data point will determine the center of the parallelepiped. That is, an $M \times N$ set of data will form an image with $M \times N$ pixels. This is different from the **pm3d** plotting style, where an $M \times N$ set of data will form a surface of $(M-1) \times (N-1)$ elements. The scan directions for a binary image data grid can be further controlled by additional keywords. See **binary keywords flipx** (p. 77), **keywords center** (p. 77), and **keywords rotate** (p. 77).

Image data can be scaled to fill a particular rectangle within a 2D plot coordinate system by specifying the x and y extent of each pixel. See **binary keywords dx** (p. 77) and **dy** (p. 77). To generate the figure at the right, the same input image was placed multiple times, each with a specified dx , dy , and origin. The input PNG image of a building is 50x128 pixels. The tall building was drawn by mapping this using **dx=0.5 dy=1.5**. The short building used a mapping **dx=0.5 dy=0.35**.

The **image** style handles input pixels containing a grayscale or color palette value. Thus 2D plots (**plot** command) require 3 columns of data ($x,y,value$), while 3D plots (**splot** command) require 4 columns of data ($x,y,z,value$).

The **rgbimage** style handles input pixels that are described by three separate values for the red, green, and blue components. Thus 5D data (x,y,r,g,b) is needed for **plot** and 6D data (x,y,z,r,g,b) for **splot**. The individual red, green, and blue components are assumed to lie in the range [0:255].

The **rgbalpha** style handles input pixels that contain alpha channel (transparency) information in addition to the red, green, and blue components. Thus 6D data (x,y,r,g,b,a) is needed for **plot** and 7D data (x,y,z,r,g,b,a) for **splot**. The r , g , b , and a components are assumed to lie in the range [0:255].

Transparency

The **rgbalpha** plotting style assumes that each pixel of input data contains an alpha value in the range [0:255]. A pixel with $\alpha = 0$ is purely transparent and does not alter the underlying contents of the plot. A pixel with $\alpha = 255$ is purely opaque. All terminal types can handle these two extreme cases. A pixel with $0 < \alpha < 255$ is partially transparent. Only a few terminal types can handle this correctly; other terminals will approximate this by treating α as being either 0 or 255.

Image pixels

Some terminals use device- or library-specific optimizations to render image data within a rectangular 2D area. This sometimes produces undesirable output, e.g. bad clipping or scaling, missing edges. The **pixels** keyword tells gnuplot to use generic code that renders the image pixel-by-pixel instead. This rendering mode is slower and may result in much larger output files, but should produce a consistent rendered view on all terminals. (The **pixels** options was called **failsafe** mode in previous gnuplot versions.) Example:

```
plot 'data' with image pixels
```

Impulses

The **impulses** style displays a vertical line from $y=0$ to the y value of each point (2D) or from $z=0$ to the z value of each point (3D). Note that the y or z values may be negative. Data from additional columns can be used to control the color of each impulse. To use this style effectively in 3D plots, it is useful to choose thick lines (`linewidth > 1`). This approximates a 3D bar chart.

```
1 column:    y
2 columns:   x y      # line from [x,0] to [x,y] (2D)
3 columns:   x y z    # line from [x,y,0] to [x,y,z] (3D)
```

Labels

The **labels** style reads coordinates and text from a data file and places the text string at the corresponding 2D or 3D position. 3 or 4 input columns of basic data are required. Additional input columns may be used to provide information such as variable font size or text color (see **rgbcolor variable** (p. 35)).

```
3 columns:   x y string # 2D version
4 columns:   x y z string # 3D version
```

The font, color, rotation angle and other properties of the printed text may be specified as additional command options (see **set label** (p. 125)). The example below generates a 2D plot with text labels constructed from the city whose name is taken from column 1 of the input file, and whose geographic coordinates are in columns 4 and 5. The font size is calculated from the value in column 3, in this case the population.

```
CityName(String,Size) = sprintf("{/= %d %s}", Scale(Size), String)
plot 'cities.dat' using 5:4:($3) using ($3) with labels
```

If we did not want to adjust the font size to a different size for each city name, the command would be much simpler:

```
plot 'cities.dat' using 5:4:1 with labels font "Times,8"
```

If the labels are marked as **hypertext** then the text only appears if the mouse is hovering over the corresponding anchor point. See **hypertext** (p. 127). In this case you must enable the label's **point** attribute so that there is a point to act as the hypertext anchor:

```
plot 'cities.dat' using 5:4:1 with labels hypertext point pt 7
```

The **labels** style can also be used in place of the **points** style when the set of predefined point symbols is not suitable or not sufficiently flexible. For example, here we define a set of chosen single-character symbols and assign one of them to each point in a plot based on the value in data column 3:

```
set encoding utf8
symbol(z) = "●◻+⊕♠♣♥♦"[int(z):int(z)]
splot 'file' using 1:2:($3) using ($3) with labels
```

See also **datastrings** (p. 21), **set style data** (p. 154).

Lines

The **lines** style connects adjacent points with straight line segments. It may be used in either 2D or 3D plots. The basic form requires 1, 2, or 3 columns of input data. Additional input columns may be used to provide information such as variable line color (see **rgbcolor variable** (p. 35)).

2D form

```
1 column:    y          # implicit x from row number
2 columns:   x y
```

3D form

```
1 column:  z      # implicit x from row, y from index
3 columns: x  y  z
```

See also **linetype** (p. 128), **linewidth** (p. 156), and **linestyle** (p. 156).

Linespoints

The **linespoints** style (short form **lp**) connects adjacent points with straight line segments and then goes back to draw a small symbol at each point. Points are drawn with the default size determined by **set pointsize** unless a specific point size is given in the plot command or a variable point size is provided in an additional column of input data. Additional input columns may also be used to provide information such as variable line color. See **lines** (p. 54) and **points** (p. 55).

The **pointinterval** (short form **pi**) property of the linetype can be used to control whether or not every point in the plot is given a symbol. For example, 'with lp pi 3' will draw line segments through every data point, but will only place a symbol on every 3rd point. A negative value for **pointinterval** will erase the portion of line segment that passes underneath the symbol. The size of the erased portion is controlled by **set pointintervalbox**.

Parallelexes

Parallel axis plots can highlight correlation in a multidimensional data set. Each input column is associated with a separately scaled vertical axis. The column values read from each line of input are connected by line segments drawn from axis 1 to axis 2 to axis 3 and so on. That is, each line of input is represented by a separate line in the parallel axes plot. It is common to use some discrete categorization to assign line colors, allowing visual exploration of the correlation between this categorization and the axis dimensions. By default gnuplot will automatically determine the range and scale of the individual axes from the input data, but the usual **set axis range** commands can be used to customize this. See **set paxis** (p. 140).

The maximum number of parallel axes is fixed at the time the program is built. The maximum for this copy of gnuplot is reported by "show version long".

Points

The **points** style displays a small symbol at each point. The command **set pointsize** may be used to change the default size of the points. 1 or 2 columns of basic input data are required in 2D plots; 1 or 3 columns are required in 3D plots. See **style lines** (p. 54). Additional input columns may be used to provide information such as variable point size or variable point color.

The first 8 point types are shared by all terminals. Individual terminals may provide a much larger number of distinct point types. Use the **test** command to show what is provided by the current terminal. Alternatively any single printable character may be given instead of a point type, as in the example below. Longer strings may be plotted using the plot style **labels** rather than **points**.

Polar

Polar plots are not really a separate plot style but are listed here for completeness. The option **set polar** tells gnuplot to interpret input 2D coordinates as <angle>,<radius> rather than <x>,<y>. Many, but not all, 2D plotting styles work in polar mode. The figure shows a combination of plot styles **lines** and **filledcurves**. See **set polar** (p. 149), **set rrange** (p. 151), **set size square** (p. 151).

Steps

The **steps** style is only relevant to 2D plotting. It connects consecutive points with two line segments: the first from (x1,y1) to (x2,y1) and the second from (x2,y1) to (x2,y2). The input column requires are the same as for plot styles **lines** and **points**. The difference between **fsteps** and **steps** is that **fsteps** traces first the change in y and then the change in x. **steps** traces first the change in x and then the change in y. To fill the area between the curve and the baseline at y=0, use **fillsteps**. See also

[steps demo](#).

Rgbalpha

See [image](#) (p. 53).

Rgbimage

See [image](#) (p. 53).

Vectors

The 2D **vectors** style draws a vector from (x,y) to (x+xdelta,y+ydelta). The 3D **vectors** style is similar, but requires six columns of basic data. A small arrowhead is drawn at the end of each vector.

```
4 columns:  x  y  xdelta  ydelta
6 columns:  x  y  z  xdelta  ydelta  zdelta
```

In both cases, an additional input column (5th in 2D, 7th in 3D) may be used to provide variable (per-datapoint) color information. (see [linecolor](#) (p. 33) and [rgbcolor variable](#) (p. 35)).

splot with vectors is supported only for **set mapping cartesian**.

The keywords "with vectors" may be followed by an inline arrow style specifications, a reference to a predefined arrow style, or a request to read the index of the desired arrow style for each vector from a separate column. Note: If you choose "arrowstyle variable" it will fill in all arrow properties at the time the corresponding vector is drawn; you cannot mix this keyword with other line or arrow style qualifiers in the plot command.

```
plot ... with vectors filled heads
plot ... with vectors arrowstyle 3
plot ... using 1:2:3:4:5 with vectors arrowstyle variable
```

Example:

```
plot 'file.dat' using 1:2:3:4 with vectors head filled lt 2
splot 'file.dat' using 1:2:3:(1):(1):(1) with vectors filled head lw 2
```

set clip one and **set clip two** affect vectors drawn in 2D. See [set clip](#) (p. 105) and [arrowstyle](#) (p. 153).

Xerrorbars

The **xerrorbars** style is only relevant to 2D data plots. **xerrorbars** is like **points**, except that a horizontal error bar is also drawn. At each point (x,y), a line is drawn from (xlow,y) to (xhigh,y) or from (x-xdelta,y) to (x+xdelta,y), depending on how many data columns are provided. A tic mark is placed at the ends of the error bar (unless **set bars** is used — see [set bars](#) (p. 102) for details). The basic style requires either 3 or 4 columns:

```
3 columns:  x  y  xdelta
4 columns:  x  y  xlow  xhigh
```

An additional input column (4th or 5th) may be used to provide information such as variable point color.

Xyerrorbars

The **xyerrorbars** style is only relevant to 2D data plots. **xyerrorbars** is like **points**, except that horizontal and vertical error bars are also drawn. At each point (x,y), lines are drawn from (x,y-ydelta) to (x,y+ydelta) and from (x-xdelta,y) to (x+xdelta,y) or from (x,ylow) to (x,yhigh) and from (xlow,y) to (xhigh,y), depending upon the number of data columns provided. A tic mark is placed at the ends of the error bar (unless **set bars** is used — see **set bars** (p. 102) for details). Either 4 or 6 input columns are required.

```
4 columns:  x  y  xdelta ydelta
6 columns:  x  y  xlow  xhigh ylow yhigh
```

If data are provided in an unsupported mixed form, the **using** filter on the **plot** command should be used to set up the appropriate form. For example, if the data are of the form (x,y,xdelta,ylow,yhigh), then you can use

```
plot 'data' using 1:2:($1-$3):($1+$3):4:5 with xyerrorbars
```

An additional input column (5th or 7th) may be used to provide variable (per-datapoint) color information.

Yerrorbars

The **yerrorbars** (or **errorbars**) style is only relevant to 2D data plots. **yerrorbars** is like **points**, except that a vertical error bar is also drawn. At each point (x,y), a line is drawn from (x,y-ydelta) to (x,y+ydelta) or from (x,ylow) to (x,yhigh), depending on how many data columns are provided. A tic mark is placed at the ends of the error bar (unless **set bars** is used — see **set bars** (p. 102) for details).

```
2 columns:  [implicit x] y ydelta
3 columns:  x  y  ydelta
4 columns:  x  y  ylow  yhigh
```

An additional input column (4th or 5th) may be used to provide information such as variable point color.

See also

[errorbar demo.](#)

Xerrorlines

The **xerrorlines** style is only relevant to 2D data plots. **xerrorlines** is like **linespoints**, except that a horizontal error line is also drawn. At each point (x,y), a line is drawn from (xlow,y) to (xhigh,y) or from (x-xdelta,y) to (x+xdelta,y), depending on how many data columns are provided. A tic mark is placed at the ends of the error bar (unless **set bars** is used — see **set bars** (p. 102) for details). The basic style requires either 3 or 4 columns:

```
3 columns:  x  y  xdelta
4 columns:  x  y  xlow  xhigh
```

An additional input column (4th or 5th) may be used to provide information such as variable point color.

Xyerrorlines

The **xyerrorlines** style is only relevant to 2D data plots. **xyerrorlines** is like **linespoints**, except that horizontal and vertical error bars are also drawn. At each point (x,y), lines are drawn from (x,y-ydelta) to (x,y+ydelta) and from (x-xdelta,y) to (x+xdelta,y) or from (x,ylow) to (x,yhigh) and from (xlow,y) to (xhigh,y), depending upon the number of data columns provided. A tic mark is placed at the ends of the error bar (unless **set bars** is used — see **set bars** (p. 102) for details). Either 4 or 6 input columns are required.

```

4 columns:  x  y  xdelta ydelta
6 columns:  x  y  xlow  xhigh ylow  yhigh

```

If data are provided in an unsupported mixed form, the **using** filter on the **plot** command should be used to set up the appropriate form. For example, if the data are of the form (x,y,xdelta,ylow,yhigh), then you can use

```
plot 'data' using 1:2:($1-$3):($1+$3):4:5 with xyerrorlines
```

An additional input column (5th or 7th) may be used to provide variable (per-datapoint) color information.

Yerrorlines

The **yerrorlines** (or **errorlines**) style is only relevant to 2D data plots. **yerrorlines** is like **linespoints**, except that a vertical error line is also drawn. At each point (x,y), a line is drawn from (x,y-ydelta) to (x,y+ydelta) or from (x,ylow) to (x,yhigh), depending on how many data columns are provided. A tic mark is placed at the ends of the error bar (see **set bars** (p. 102) for details). Either 3 or 4 input columns are required.

```

3 columns:  x  y  ydelta
4 columns:  x  y  ylow  yhigh

```

An additional input column (4th or 5th) may be used to provide information such as variable point color.

See also

[errorbar demo.](#)

3D (surface) plots

Surface plots are generated using the **splot** command rather than the **plot** command. The style **with lines** draws a surface made from a grid of lines. Solid surfaces can be drawn using the style **with pm3d**. Usually the surface is displayed at some arbitrary viewing angle, such that it clearly represents a 3D surface. In this case the X, Y, and Z axes are all visible in the plot. The illusion of 3D is enhanced by choosing hidden line removal or depth-sorted surface elements. See **hidden3d** (p. 120) and the **depthorder** (p. 142) option of **set pm3d** (p. 141). The **splot** command can also calculate and draw contour lines corresponding to constant Z values. These contour lines may be drawn onto the surface itself, or projected onto the XY plane. See **set contour** (p. 108).

2D projection (set view map)

An important special case of the **splot** command is to map the Z coordinate onto a 2D surface by projecting the plot along the Z axis. See **set view map** (p. 164). This plot mode can be used to generate contour plots and heat maps. This figure shows contours plotted once with plot style **lines**, once with style **labels**.

Part III

Commands

This section lists the commands acceptable to **gnuplot** in alphabetical order. Printed versions of this document contain all commands; the text available interactively may not be complete. Indeed, on some systems there may be no commands at all listed under this heading.

Note that in most cases unambiguous abbreviations for command names and their options are permissible, i.e., "**p f(x) w li**" instead of "**plot f(x) with lines**".

In the syntax descriptions, braces ({}) denote optional arguments and a vertical bar (|) separates mutually exclusive choices.

Cd

The **cd** command changes the working directory.

Syntax:

```
cd '<directory-name>'
```

The directory name must be enclosed in quotes.

Examples:

```
cd 'subdir'
cd ".."
```

It is recommended that Windows users use single-quotes, because backslash [\\] has special significance inside double-quotes and has to be escaped. For example,

```
cd "c:\\newdata"
```

fails, but

```
cd 'c:\\newdata'
cd "c:\\\\newdata"
```

work as expected.

Call

The **call** command is identical to the **load** command with one exception: the name of the file being loaded may be followed by up to nine parameters.

```
call "inputfile" <param-1> <param-2> <param-3> ... <param-9>
```

Previous versions of gnuplot performed macro-like substitution of the special tokens \$0, \$1, ... \$9 with the literal contents of these parameters. This mechanism is now deprecated (see **call old-style (p. 60)**).

Gnuplot now provides a set of string variables ARG0, ARG1, ..., ARG9 and an integer variable ARGC. When a **call** command is executed ARG0 is set to the name of the input file, ARGC is set to the number of parameters present, and ARG1 to ARG9 are loaded from the parameters that follow it on the command line. Any existing contents of the ARG variables are saved and restored across a **call** command.

Because the parameters are stored in ordinary string variables, they may be dereferenced by macro expansion (analogous to the old-style deprecated syntax). However in many cases it is more natural to use them as you would any other variable.

Example

```

Call site
  MYFILE = "script1.gp"
  FUNC = "sin(x)"
  call MYFILE FUNC 1.23 "This is a plot title"
Upon entry to the called script
  ARG0 holds "script1.gp"
  ARG1 holds the string "sin(x)"
  ARG2 holds the string "1.23"
  ARG3 holds the string "This is a plot title"
  ARGV is 3
The script itself can now execute
  plot @ARG1 with lines title ARG3
  print ARG2 * 4.56, @ARG2 * 4.56
  print "This plot produced by script ", ARG0

```

Notice that ARG1 must be dereferenced as a macro, but ARG2 may be dereferenced either as a macro (yielding a numerical constant) or a variable (yielding that same numerical value after auto-promotion of the string "1.23" to a real).

The same result could be obtained directly from a shell script by invoking gnuplot with the **-c** command line option:

```
gnuplot -persist -c "script1.gp" "sin(x)" 1.23 "This is a plot title"
```

Old-style

This describes the call mechanism used by previous versions of gnuplot, now deprecated.

```
call "<input-file>" <param-0> <param-1> ... <param-9>
```

The name of the input file must be enclosed in quotes. As each line is read from the input file, it is scanned for the following special character sequences: \$0 \$1 \$2 \$3 \$4 \$5 \$6 \$7 \$8 \$9 \$#. If found, the sequence \$+digit is replaced by the corresponding parameter from the **call** command line. Quote characters are not copied and string variable substitution is not performed. The character sequence \$# is replaced by the number of passed parameters. \$ followed by any other character is treated as an escape sequence; use \$\$ to get a single \$.

Example:

If the file 'calltest.gp' contains the line:

```
print "argc=$# p0=$0 p1=$1 p2=$2 p3=$3 p4=$4 p5=$5 p6=$6 p7=x$7x"
```

entering the command:

```
call 'calltest.gp' "abcd" 1.2 + "'quoted'" -- "$2"
```

will display:

```
argc=7 p0=abcd p1=1.2 p2=+ p3='quoted' p4=- p5=- p6=$2 p7=xx
```

NOTES: This use of the \$ character conflicts both with gnuplot's own syntax for datafile columns and with the use of \$ to indicate environmental variables in a unix-like shell. The special sequence \$# was mis-interpreted as a comment delimiter in gnuplot versions 4.5 through 4.6.3. Quote characters are ignored during substitution, so string constants are easily corrupted.

Clear

The **clear** command erases the current screen or output device as specified by **set output**. This usually generates a formfeed on hardcopy devices. Use **set terminal** to set the device type.

For some terminals **clear** erases only the portion of the plotting surface defined by **set size**, so for these it can be used in conjunction with **set multiplot** to create an inset.

Example:

```
set multiplot
plot sin(x)
set origin 0.5,0.5
set size 0.4,0.4
clear
plot cos(x)
unset multiplot
```

Please see **set multiplot** (p. 133), **set size** (p. 151), and **set origin** (p. 139) for details of these commands.

Do

Syntax:

```
do for <iteration-spec> {
    <commands>
    <commands>
}
```

Execute a sequence of commands multiple times. The commands must be enclosed in curly brackets, and the opening "{" must be on the same line as the **do** keyword. This command cannot be used with old-style (un-bracketed) if/else statements. See **if** (p. 70). For examples of iteration specifiers, see **iteration** (p. 33).

Example:

```
set multiplot layout 2,2
do for [name in "A B C D"] {
    filename = name . ".dat"
    set title sprintf("Condition %s",name)
    plot filename title name
}
unset multiplot
```

Evaluate

The **evaluate** command executes the commands given as an argument string. Newline characters are not allowed within the string.

Syntax:

```
eval <string expression>
```

This is especially useful for a repetition of similar commands.

Example:

```
set_label(x, y, text) \
= sprintf("set label '%s' at %f, %f point pt 5", text, x, y)
eval set_label(1., 1., 'one/one')
eval set_label(2., 1., 'two/one')
eval set_label(1., 2., 'one/two')
```

Please see **substitution macros** (p. 40) for another way to execute commands from a string.

Exit

The commands **exit** and **quit**, as well as the END-OF-FILE character (usually Ctrl-D) terminate input from the current input stream: terminal session, pipe, or file input (pipe). If input streams are nested (inherited **load** scripts), then reading will continue in the parent stream. When the top level stream is closed, the program itself will exit.

The command **exit gnuplot** will immediately and unconditionally cause gnuplot to exit even if the input stream is multiply nested. In this case any open output files may not be completed cleanly. Example of use:

```
bind "ctrl-x" "unset output; exit gnuplot"
```

The command **exit error "error message"** simulates a program error. In interactive mode it prints the error message and returns to the command line, breaking out of all nested loops or calls. In non-interactive mode the program will exit.

See help for **batch/interactive** (p. 19) for more details.

Fit

The **fit** command fits a user-supplied real-valued expression to a set of data points, using the nonlinear least-squares Marquardt-Levenberg algorithm. There can be up to 12 independent variables, there is always 1 dependent variable, and any number of parameters can be fitted. Optionally, error estimates can be input for weighting the data points.

The basic use of **fit** is best explained by a simple example:

```
f(x) = a + b*x + c*x**2
fit f(x) 'measured.dat' using 1:2 via a,b,c
plot 'measured.dat' u 1:2, f(x)
```

Syntax:

```
fit {<ranges>} <expression>
    '<datafile>' {datafile-modifiers}
    [{unitweights} | {y|xy|z}error | errors <var1>{,<var2>,...}]
    via '<parameter file>' | <var1>{,<var2>,...}
```

Ranges may be specified to filter the data used in fitting. Out-of-range data points are ignored. The syntax is

```
[{dummy_variable=}{<min>}{:<max>}],
```

analogous to **plot**; see **plot ranges** (p. 90).

<expression> can be any valid **gnuplot** expression, although the most common is a previously user-defined function of the form $f(x)$ or $f(x,y)$. It must be real-valued. The names of the independent variables are set by the **set dummy** command, or in the <ranges> part of the command (see below); by default, the first two are called x and y . Furthermore, the expression should depend on one or more variables whose value is to be determined by the fitting procedure.

<datafile> is treated as in the **plot** command. All the **plot datafile** modifiers (**using**, **every**,...) except **smooth** are applicable to **fit**. See **plot datafile** (p. 78).

The datafile contents can be interpreted flexibly by providing a **using** qualifier as with plot commands. For example to generate the independent variable x as the sum of columns 2 and 3, while taking z from column 6 and requesting equal weights:

```
fit ... using ($2+$3):6
```

In the absence of a **using** specification, the fit implicitly assumes there is only a single independent variable. If the file itself, or the using specification, contains only a single column of data, the line number is taken as the independent variable. If a **using** specification is given, there can be up to 12 independent variables (and more if specially configured at compile time).

The **unitweights** option, which is the default, causes all data points to be weighted equally. This can be changed by using the **errors** keyword to read error estimates of one or more of the variables from the data file. These error estimates are interpreted as the standard deviation s of the corresponding variable value and used to compute a weight for the datum as $1/s^2$.

In case of error estimates of the independent variables, these weights are further multiplied by fitting function derivatives according to the "effective variance method" (Jay Orear, Am. J. Phys., Vol. 50, 1982).

The **errors** keyword is to be followed by a comma-separated list of one or more variable names for which errors are to be input; the dependent variable z must always be among them, while independent variables are optional. For each variable in this list, an additional column will be read from the file, containing that variable's error estimate. Again, flexible interpretation is possible by providing the **using** qualifier. Note that the number of independent variables is thus implicitly given by the total number of columns in the **using** qualifier, minus 1 (for the dependent variable), minus the number of variables in the **errors** qualifier.

As an example, if one has 2 independent variables, and errors for the first independent variable and the dependent variable, one uses the **errors x,z** qualifier, and a **using** qualifier with 5 columns, which are interpreted as $x:y:z:sx:sz$ (where x and y are the independent variables, z the dependent variable, and sx and sz the standard deviations of x and z).

A few shorthands for the **errors** qualifier are available: **yerrors** (for fits with 1 column of independent variable), and **zerrors** (for the general case) are all equivalent to **errors z**, indicating that there is a single extra column with errors of the dependent variable.

xyerrors, for the case of 1 independent variable, indicates that there are two extra columns, with errors of both the independent and the dependent variable. In this case the errors on x and y are treated by Orear's effective variance method.

Note that **yerror** and **xyerror** are similar in both form and interpretation to the **yerrorlines** and **xyerrorlines** 2D plot styles.

With the command **set fit v4** the fit command syntax is compatible with **gnuplot** version 4 and before. Then there must be two more **using** qualifiers (z and s) than there are independent variables, unless there is only one variable. **gnuplot** then uses the following formats, depending on the number of columns given in the **using** specification:

z	# 1 independent variable (line number)
$x:z$	# 1 independent variable (1st column)
$x:z:s$	# 1 independent variable (3 columns total)
$x:y:z:s$	# 2 independent variables (4 columns total)
$x1:x2:x3:z:s$	# 3 independent variables (5 columns total)
$x1:x2:x3:...:xN:z:s$	# N independent variables (N+2 columns total)

Please beware that this means that you have to supply z -errors s in a fit with two or more independent variables. If you want unit weights you need to supply them explicitly by using e.g. then format $x:y:z:(1)$.

The dummy variable names may be changed when specifying a range as noted above. The first range corresponds to the first **using** spec, and so on. A range may also be given for z (the dependent variable), in which case data points for which $f(x,...)$ is out of the z range will not contribute to the residual being minimized.

Multiple datasets may be simultaneously fit with functions of one independent variable by making y a 'pseudo-variable', e.g., the dataline number, and fitting as two independent variables. See **fit multi-branch** (p. 68).

The **via** qualifier specifies which parameters are to be optimized, either directly, or by referencing a parameter file.

Examples:

```
f(x) = a*x**2 + b*x + c
g(x,y) = a*x**2 + b*y**2 + c*x*y
set fit limit 1e-6
fit f(x) 'measured.dat' via 'start.par'
fit f(x) 'measured.dat' using 3:($-5) via 'start.par'
fit f(x) './data/trash.dat' using 1:2:3 yerror via a, b, c
```

```

fit g(x,y) 'surface.dat' using 1:2:3 via a, b, c
fit a0 + a1*x/(1 + a2*x/(1 + a3*x)) 'measured.dat' via a0,a1,a2,a3
fit a*x + b*y 'surface.dat' using 1:2:3 via a,b
fit [*:*][yaks=*:~] a*x+b*yaks 'surface.dat' u 1:2:3 via a,b

fit [] [] [t=*:~] a*x + b*y + c*t 'foo.dat' using 1:2:3:4 via a,b,c

set dummy x1, x2, x3, x4, x5
h(x1,x2,x3,x4,s5) = a*x1 + b*x2 + c*x3 + d*x4 + e*x5
fit h(x1,x2,x3,x4,x5) 'foo.dat' using 1:2:3:4:5:6 via a,b,c,d,e

```

After each iteration step, detailed information about the current state of the fit is written to the display. The same information about the initial and final states is written to a log file, "fit.log". This file is always appended to, so as to not lose any previous fit history; it should be deleted or renamed as desired. By using the command **set fit logfile**, the name of the log file can be changed.

If activated by using **set fit errorvariables**, the error for each fitted parameter will be stored in a variable named like the parameter, but with "_err" appended. Thus the errors can be used as input for further computations.

If **set fit prescale** is activated, fit parameters are prescaled by their initial values. This helps the Marquardt-Levenberg routine converge more quickly and reliably in cases where parameters differ in size by several orders of magnitude.

The fit may be interrupted by pressing Ctrl-C (Ctrl-Break in wgnuplot). After the current iteration completes, you have the option to (1) stop the fit and accept the current parameter values, (2) continue the fit, (3) execute a **gnuplot** command as specified by **set fit script** or the environment variable **FIT_SCRIPT**. The default is **replot**, so if you had previously plotted both the data and the fitting function in one graph, you can display the current state of the fit.

Once **fit** has finished, the **update** command may be used to store final values in a file for subsequent use as a parameter file. See **update** (p. 183) for details.

Adjustable parameters

There are two ways that **via** can specify the parameters to be adjusted, either directly on the command line or indirectly, by referencing a parameter file. The two use different means to set initial values.

Adjustable parameters can be specified by a comma-separated list of variable names after the **via** keyword. Any variable that is not already defined is created with an initial value of 1.0. However, the fit is more likely to converge rapidly if the variables have been previously declared with more appropriate starting values.

In a parameter file, each parameter to be varied and a corresponding initial value are specified, one per line, in the form

```
varname = value
```

Comments, marked by '#', and blank lines are permissible. The special form

```
varname = value      # FIXED
```

means that the variable is treated as a 'fixed parameter', initialized by the parameter file, but not adjusted by **fit**. For clarity, it may be useful to designate variables as fixed parameters so that their values are reported by **fit**. The keyword **# FIXED** has to appear in exactly this form.

Short introduction

fit is used to find a set of parameters that 'best' fits your data to your user-defined function. The fit is judged on the basis of the sum of the squared differences or 'residuals' (SSR) between the input data points and the function values, evaluated at the same places. This quantity is often called 'chisquare' (i.e., the Greek letter chi, to the power of 2). The algorithm attempts to minimize SSR, or more precisely, WSSR, as the residuals are 'weighted' by the input data errors (or 1.0) before being squared; see **fit error_estimates** (p. 65) for details.

That's why it is called 'least-squares fitting'. Let's look at an example to see what is meant by 'non-linear', but first we had better go over some terms. Here it is convenient to use z as the dependent variable for user-defined functions of either one independent variable, $z=f(x)$, or two independent variables, $z=f(x,y)$. A parameter is a user-defined variable that **fit** will adjust, i.e., an unknown quantity in the function declaration. Linearity/non-linearity refers to the relationship of the dependent variable, z , to the parameters which **fit** is adjusting, not of z to the independent variables, x and/or y . (To be technical, the second {and higher} derivatives of the fitting function with respect to the parameters are zero for a linear least-squares problem).

For linear least-squares (LLS), the user-defined function will be a sum of simple functions, not involving any parameters, each multiplied by one parameter. NLLS handles more complicated functions in which parameters can be used in a large number of ways. An example that illustrates the difference between linear and nonlinear least-squares is the Fourier series. One member may be written as

$$z=a*\sin(c*x) + b*\cos(c*x).$$

If a and b are the unknown parameters and c is constant, then estimating values of the parameters is a linear least-squares problem. However, if c is an unknown parameter, the problem is nonlinear.

In the linear case, parameter values can be determined by comparatively simple linear algebra, in one direct step. However LLS is a special case which is also solved along with more general NLLS problems by the iterative procedure that **gnuplot** uses. **fit** attempts to find the minimum by doing a search. Each step (iteration) calculates WSSR with a new set of parameter values. The Marquardt-Levenberg algorithm selects the parameter values for the next iteration. The process continues until a preset criterion is met, either (1) the fit has "converged" (the relative change in WSSR is less than a certain limit, see **set fit limit** (p. 115)), or (2) it reaches a preset iteration count limit (see **set fit maxiter** (p. 115)). The fit may also be interrupted and subsequently halted from the keyboard (see **fit** (p. 62)). The user variable `FIT_CONVERGED` contains 1 if the previous fit command terminated due to convergence; it contains 0 if the previous fit terminated for any other reason. `FIT_NITER` contains the number of iterations that were done during the last fit.

Often the function to be fitted will be based on a model (or theory) that attempts to describe or predict the behaviour of the data. Then **fit** can be used to find values for the free parameters of the model, to determine how well the data fits the model, and to estimate an error range for each parameter. See **fit error_estimates** (p. 65).

Alternatively, in curve-fitting, functions are selected independent of a model (on the basis of experience as to which are likely to describe the trend of the data with the desired resolution and a minimum number of parameters*functions.) The **fit** solution then provides an analytic representation of the curve.

However, if all you really want is a smooth curve through your data points, the **smooth** option to **plot** may be what you've been looking for rather than **fit**.

Error estimates

In **fit**, the term "error" is used in two different contexts, data error estimates and parameter error estimates.

Data error estimates are used to calculate the relative weight of each data point when determining the weighted sum of squared residuals, WSSR or chisquare. They can affect the parameter estimates, since they determine how much influence the deviation of each data point from the fitted function has on the final values. Some of the **fit** output information, including the parameter error estimates, is more meaningful if accurate data error estimates have been provided.

The **statistical overview** describes some of the **fit** output and gives some background for the 'practical guidelines'.

Statistical overview

The theory of non-linear least-squares (NLLS) is generally described in terms of a normal distribution of errors, that is, the input data is assumed to be a sample from a population having a given mean and a Gaussian (normal) distribution about the mean with a given standard deviation. For a sample of sufficiently large size, and knowing the population standard deviation, one can use the statistics of the chisquare distribution to describe a "goodness of fit" by looking at the variable often called "chisquare". Here, it is

sufficient to say that a reduced chisquare (chisquare/degrees of freedom, where degrees of freedom is the number of datapoints less the number of parameters being fitted) of 1.0 is an indication that the weighted sum of squared deviations between the fitted function and the data points is the same as that expected for a random sample from a population characterized by the function with the current value of the parameters and the given standard deviations.

If the standard deviation for the population is not constant, as in counting statistics where variance = counts, then each point should be individually weighted when comparing the observed sum of deviations and the expected sum of deviations.

At the conclusion **fit** reports 'stdfit', the standard deviation of the fit, which is the rms of the residuals, and the variance of the residuals, also called 'reduced chisquare' when the data points are weighted. The number of degrees of freedom (the number of data points minus the number of fitted parameters) is used in these estimates because the parameters used in calculating the residuals of the datapoints were obtained from the same data. If the data points have weights, **gnuplot** calculates the so-called p-value, i.e. one minus the cumulative distribution function of the chisquare-distribution for the number of degrees of freedom and the resulting chisquare, see **practical guidelines** (p. 66). These values are exported to the variables

```
FIT_NDF = Number of degrees of freedom
FIT_WSSR = Weighted sum-of-squares residual
FIT_STDFIT = sqrt(WSSR/NDF)
FIT_P = p-value
```

To estimate confidence levels for the parameters, one can use the minimum chisquare obtained from the fit and chisquare statistics to determine the value of chisquare corresponding to the desired confidence level, but considerably more calculation is required to determine the combinations of parameters which produce such values.

Rather than determine confidence intervals, **fit** reports parameter error estimates which are readily obtained from the variance-covariance matrix after the final iteration. By convention, these estimates are called "standard errors" or "asymptotic standard errors", since they are calculated in the same way as the standard errors (standard deviation of each parameter) of a linear least-squares problem, even though the statistical conditions for designating the quantity calculated to be a standard deviation are not generally valid for the NLLS problem. The asymptotic standard errors are generally over-optimistic and should not be used for determining confidence levels, but are useful for qualitative purposes.

The final solution also produces a correlation matrix indicating correlation of parameters in the region of the solution; The main diagonal elements, autocorrelation, are always 1; if all parameters were independent, the off-diagonal elements would be nearly 0. Two variables which completely compensate each other would have an off-diagonal element of unit magnitude, with a sign depending on whether the relation is proportional or inversely proportional. The smaller the magnitudes of the off-diagonal elements, the closer the estimates of the standard deviation of each parameter would be to the asymptotic standard error.

Practical guidelines

If you have a basis for assigning weights to each data point, doing so lets you make use of additional knowledge about your measurements, e.g., take into account that some points may be more reliable than others. That may affect the final values of the parameters.

Weighting the data provides a basis for interpreting the additional **fit** output after the last iteration. Even if you weight each point equally, estimating an average standard deviation rather than using a weight of 1 makes WSSR a dimensionless variable, as chisquare is by definition.

Each fit iteration will display information which can be used to evaluate the progress of the fit. (An '*' indicates that it did not find a smaller WSSR and is trying again.) The 'sum of squares of residuals', also called 'chisquare', is the WSSR between the data and your fitted function; **fit** has minimized that. At this stage, with weighted data, chisquare is expected to approach the number of degrees of freedom (data points minus parameters). The WSSR can be used to calculate the reduced chisquare (WSSR/ndf) or stdfit, the standard deviation of the fit, sqrt(WSSR/ndf). Both of these are reported for the final WSSR.

If the data are unweighted, stdfit is the rms value of the deviation of the data from the fitted function, in user units.

If you supplied valid data errors, the number of data points is large enough, and the model is correct, the reduced chisquare should be about unity. (For details, look up the 'chi-squared distribution' in your favorite statistics reference.) If so, there are additional tests, beyond the scope of this overview, for determining how well the model fits the data.

A reduced chisquare much larger than 1.0 may be due to incorrect data error estimates, data errors not normally distributed, systematic measurement errors, 'outliers', or an incorrect model function. A plot of the residuals, e.g., **plot 'datafile' using 1:(\$2-f(\$1))**, may help to show any systematic trends. Plotting both the data points and the function may help to suggest another model.

Similarly, a reduced chisquare less than 1.0 indicates WSSR is less than that expected for a random sample from the function with normally distributed errors. The data error estimates may be too large, the statistical assumptions may not be justified, or the model function may be too general, fitting fluctuations in a particular sample in addition to the underlying trends. In the latter case, a simpler function may be more appropriate.

The p-value of the fit is one minus the cumulative distribution function of the chisquare-distribution for the number of degrees of freedom and the resulting chisquare. This can serve as a measure of the goodness-of-fit. The range of the p-value is between zero and one. A very small or large p-value indicates that the model does not describe the data and its errors well. As described above, this might indicate a problem with the data, its errors or the model, or a combination thereof. A small p-value might indicate that the errors have been underestimated and the errors of the final parameters should thus be scaled. See also **set fit errorscale** (p. 115).

You'll have to get used to both **fit** and the kind of problems you apply it to before you can relate the standard errors to some more practical estimates of parameter uncertainties or evaluate the significance of the correlation matrix.

Note that **fit**, in common with most NLLS implementations, minimizes the weighted sum of squared distances $(y-f(x))^2$. It does not provide any means to account for "errors" in the values of x, only in y. Also, any "outliers" (data points outside the normal distribution of the model) will have an exaggerated effect on the solution.

Control

Settings of the **fit** command are controlled by **set fit**. The old **gnuplot** user variables are deprecated as of version 5, see **fit control variables** (p. 67).

There are a number of environment variables that can be defined to affect **fit** before starting **gnuplot**, see **fit control environment** (p. 68).

Control variables

The user defined variables described here are deprecated, see **set fit** (p. 115).

The default epsilon limit (1e-5) may be changed by declaring a value for

FIT_LIMIT

When the sum of squared residuals changes between two iteration steps by a factor less than this number (epsilon), the fit is considered to have 'converged'.

The maximum number of iterations may be limited by declaring a value for

FIT_MAXITER

A value of 0 (or not defining it at all) means that there is no limit.

If you need even more control about the algorithm, and know the Marquardt-Levenberg algorithm well, there are some more variables to influence it. The startup value of **lambda** is normally calculated automatically from the ML-matrix, but if you want to, you may provide your own one with

FIT_START_LAMBDA

Specifying **FIT_START_LAMBDA** as zero or less will re-enable the automatic selection. The variable

FIT_LAMBDA_FACTOR

gives the factor by which **lambda** is increased or decreased whenever the chi-squared target function increased or decreased significantly. Setting FIT_LAMBDA_FACTOR to zero re-enables the default factor of 10.0.

Other variables with the FIT_ prefix may be added to **fit**, so it is safer not to use that prefix for user-defined variables.

The variables FIT_SKIP and FIT_INDEX were used by earlier releases of **gnuplot** with a 'fit' patch called **gnufit** and are no longer available. The datafile **every** modifier provides the functionality of FIT_SKIP. FIT_INDEX was used for multi-branch fitting, but multi-branch fitting of one independent variable is now done as a pseudo-3D fit in which the second independent variable and **using** are used to specify the branch. See **fit multi-branch** (p. 68).

Environment variables

The environment variables must be defined before **gnuplot** is executed; how to do so depends on your operating system.

FIT_LOG

changes the name (and/or path) of the file to which the fit log will be written from the default of "fit.log" in the working directory. The default value can be overwritten using the command **set fit logfile**.

FIT_SCRIPT

specifies a command that may be executed after an user interrupt. The default is **replot**, but a **plot** or **load** command may be useful to display a plot customized to highlight the progress of the fit. This setting can also be changed using **set fit script**.

Multi-branch

In multi-branch fitting, multiple data sets can be simultaneously fit with functions of one independent variable having common parameters by minimizing the total WSSR. The function and parameters (branch) for each data set are selected by using a 'pseudo-variable', e.g., either the dataline number (a 'column' index of -1) or the datafile index (-2), as the second independent variable.

Example: Given two exponential decays of the form, $z=f(x)$, each describing a different data set but having a common decay time, estimate the values of the parameters. If the datafile has the format $x:z:s$, then

```
f(x,y) = (y==0) ? a*exp(-x/tau) : b*exp(-x/tau)
fit f(x,y) 'datafile' using 1:-2:2:3 via a, b, tau
```

For a more complicated example, see the file "hexa.fnc" used by the "fit.dem" demo.

Appropriate weighting may be required since unit weights may cause one branch to predominate if there is a difference in the scale of the dependent variable. Fitting each branch separately, using the multi-branch solution as initial values, may give an indication as to the relative effect of each branch on the joint solution.

Starting values

Nonlinear fitting is not guaranteed to converge to the global optimum (the solution with the smallest sum of squared residuals, SSR), and can get stuck at a local minimum. The routine has no way to determine that; it is up to you to judge whether this has happened.

fit may, and often will get "lost" if started far from a solution, where SSR is large and changing slowly as the parameters are varied, or it may reach a numerically unstable region (e.g., too large a number causing a floating point overflow) which results in an "undefined value" message or **gnuplot** halting.

To improve the chances of finding the global optimum, you should set the starting values at least roughly in the vicinity of the solution, e.g., within an order of magnitude, if possible. The closer your starting values are to the solution, the less chance of stopping at another minimum. One way to find starting values is to plot data and the fitting function on the same graph and change parameter values and **replot** until reasonable

similarity is reached. The same plot is also useful to check whether the fit stopped at a minimum with a poor fit.

Of course, a reasonably good fit is not proof there is not a "better" fit (in either a statistical sense, characterized by an improved goodness-of-fit criterion, or a physical sense, with a solution more consistent with the model.) Depending on the problem, it may be desirable to **fit** with various sets of starting values, covering a reasonable range for each parameter.

Tips

Here are some tips to keep in mind to get the most out of **fit**. They're not very organized, so you'll have to read them several times until their essence has sunk in.

The two forms of the **via** argument to **fit** serve two largely distinct purposes. The **via "file"** form is best used for (possibly unattended) batch operation, where you just supply the startup values in a file and can later use **update** to copy the results back into another (or the same) parameter file.

The **via var1, var2, ...** form is best used interactively, where the command history mechanism may be used to edit the list of parameters to be fitted or to supply new startup values for the next try. This is particularly useful for hard problems, where a direct fit to all parameters at once won't work without good starting values. To find such, you can iterate several times, fitting only some of the parameters, until the values are close enough to the goal that the final fit to all parameters at once will work.

Make sure that there is no mutual dependency among parameters of the function you are fitting. For example, don't try to fit $a \cdot \exp(x+b)$, because $a \cdot \exp(x+b) = a \cdot \exp(b) \cdot \exp(x)$. Instead, fit either $a \cdot \exp(x)$ or $\exp(x+b)$.

A technical issue: The larger the ratio of the largest and the smallest absolute parameter values, the slower the fit will converge. If the ratio is close to or above the inverse of the machine floating point precision, it may take next to forever to converge, or refuse to converge at all. You will either have to adapt your function to avoid this, e.g., replace 'parameter' by '1e9*parameter' in the function definition, and divide the starting value by 1e9 or use **set fit prescale** which does this internally according to the parameter starting values.

If you can write your function as a linear combination of simple functions weighted by the parameters to be fitted, by all means do so. That helps a lot, because the problem is no longer nonlinear and should converge with only a small number of iterations, perhaps just one.

Some prescriptions for analysing data, given in practical experimentation courses, may have you first fit some functions to your data, perhaps in a multi-step process of accounting for several aspects of the underlying theory one by one, and then extract the information you really wanted from the fitting parameters of those functions. With **fit**, this may often be done in one step by writing the model function directly in terms of the desired parameters. Transforming data can also quite often be avoided, though sometimes at the cost of a more difficult fit problem. If you think this contradicts the previous paragraph about simplifying the fit function, you are correct.

A "singular matrix" message indicates that this implementation of the Marquardt-Levenberg algorithm can't calculate parameter values for the next iteration. Try different starting values, writing the function in another form, or a simpler function.

Finally, a nice quote from the manual of another fitting package (fudgit), that kind of summarizes all these issues: "Nonlinear fitting is an art!"

Help

The **help** command displays built-in help. To specify information on a particular topic use the syntax:

```
help {<topic>}
```

If <topic> is not specified, a short message is printed about **gnuplot**. After help for the requested topic is given, a menu of subtopics is given; help for a subtopic may be requested by typing its name, extending the help request. After that subtopic has been printed, the request may be extended again or you may go back one level to the previous topic. Eventually, the **gnuplot** command line will return.

If a question mark (?) is given as the topic, the list of topics currently available is printed on the screen.

History

The **history** command print or saves previous commands in the history list, or reexecutes an previous entry in the list. To modify the behavior of this command, see **set history** (p. 122).

Examples:

```
history                # show the complete history
history 5              # show last 5 entries in the history
history quiet 5        # show last 5 entries without entry numbers
history "hist.gp"      # write the complete history to file hist.gp
history "hist.gp" append # append the complete history to file hist.gp
history 10 "hist.gp"   # write last 10 commands to file hist.gp
history 10 "|head -5 >>diary.gp" # write 5 history commands using pipe
history ?load          # show all history entries starting with "load"
history ?"set c"        # like above, several words enclosed in quotes
hi !reread             # execute last entry starting with "reread"
hist !"set xr"         # like above, several words enclosed in quotes
hist !55               # reexecute the command at history entry 55
```

If

New syntax:

```
if (<condition>) { <commands>;
    <commands>
    <commands>
} else {
    <commands>
}
```

Old syntax:

```
if (<condition>) <command-line> [; else if (<condition>) ...; else ...]
```

This version of gnuplot supports block-structured if/else statements. If the keyword **if** or **else** is immediately followed by an opening "{", then conditional execution applies to all statements, possibly on multiple input lines, until a matching "}" terminates the block. If commands may be nested.

The old single-line if/else syntax is still supported, but can not be mixed with the new block-structured syntax. See **if-old** (p. 70).

If-old

Through gnuplot version 4.4, the scope of the if/else commands was limited to a single input line. Now a multi-line clause may be enclosed in curly brackets. The old syntax is still honored but cannot be used inside a bracketed clause.

If no opening "{" follows the **if** keyword, the command(s) in <command-line> will be executed if <condition> is true (non-zero) or skipped if <condition> is false (zero). Either case will consume commands on the input line until the end of the line or an occurrence of **else**. Note that use of ; to allow multiple commands on the same line will *not* end the conditionalized commands.

Examples:

```
pi=3
if (pi!=acos(-1)) print "?Fixing pi!"; pi=acos(-1); print pi
```

will display:

```
?Fixing pi!
3.14159265358979
```

but

```
if (1==2) print "Never see this"; print "Or this either"
```

will not display anything.

else:

```
v=0
v=v+1; if (v%2) print "2" ; else if (v%3) print "3"; else print "fred"
```

(repeat the last line repeatedly!)

For

The **plot**, **splot**, **set** and **unset** commands may optionally contain an iteration for clause. This has the effect of executing the basic command multiple times, each time re-evaluating any expressions that make use of the iteration control variable. Iteration of arbitrary command sequences can be requested using the **do** command. Two forms of iteration clause are currently supported:

```
for [intvar = start:end{:increment}]
for [stringvar in "A B C D"]
```

Examples:

```
plot for [filename in "A.dat B.dat C.dat"] filename using 1:2 with lines
plot for [basename in "A B C"] basename.".dat" using 1:2 with lines
set for [i = 1:10] style line i lc rgb "blue"
unset for [tag = 100:200] label tag
```

Nested iteration is supported:

```
set for [i=1:9] for [j=1:9] label i*10+j sprintf("%d",i*10+j) at i,j
```

See additional documentation for **iteration** (p. 33), **do** (p. 61).

Import

The **import** command associates a user-defined function name with a function exported by an external shared object. This constitutes a plugin mechanism that extends the set of functions available in gnuplot.

Syntax:

```
import func(x[,y,z,...]) from "sharedobj[:symbol]"
```

Examples:

```
# make the function myfun, exported by "mylib.so" or "mylib.dll"
# available for plotting or numerical calculation in gnuplot
import myfun(x) from "mylib"
import myfun(x) from "mylib:myfun"      # same as above

# make the function theirfun, defined in "theirlib.so" or "theirlib.dll"
# available under a different name
import myfun(x,y,z) from "theirlib:theirfun"
```

The program extends the name given for the shared object by either ".so" or ".dll" depending on the operating system, and searches for it first as a full path name and then as a path relative to the current directory. The operating system itself may also search any directories in LD_LIBRARY_PATH or DYLD_LIBRARY_PATH.

Load

The **load** command executes each line of the specified input file as if it had been typed in interactively. Files created by the **save** command can later be **loaded**. Any text file containing valid commands can be created and then executed by the **load** command. Files being **loaded** may themselves contain **load** or **call** commands. See **comments** (p. 21) for information about comments in commands. To **load** with arguments, see **call** (p. 59).

Syntax:

```
load "<input-file>"
```

The name of the input file must be enclosed in quotes.

The special filename "-" may be used to **load** commands from standard input. This allows a **gnuplot** command file to accept some commands from standard input. Please see help for **batch/interactive** (p. 19) for more details.

On some systems which support a `popen` function (Unix), the load file can be read from a pipe by starting the file name with a '<'.

Examples:

```
load 'work.gnu'
load "func.dat"
load "< loadfile_generator.sh"
```

The **load** command is performed implicitly on any file names given as arguments to **gnuplot**. These are loaded in the order specified, and then **gnuplot** exits.

Lower

Syntax:

```
lower {plot_window_nb}
```

The **lower** command lowers (opposite to **raise**) plot window(s) associated with the interactive terminal of your gnuplot session, i.e. **pm**, **win**, **wxt** or **x11**. It puts the plot window to bottom in the z-order windows stack of the window manager of your desktop.

As **x11** and **wxt** support multiple plot windows, then by default they lower these windows in descending order of most recently created on top to the least recently created on bottom. If a plot number is supplied as an optional parameter, only the associated plot window will be lowered if it exists.

The optional parameter is ignored for single plot-window terminals, i.e. **pm** and **win**.

Pause

The **pause** command displays any text associated with the command and then waits a specified amount of time or until the carriage return is pressed. **pause** is especially useful in conjunction with **load** files.

Syntax:

```
pause <time> {"<string>"}
pause mouse {<endcondition>}{, <endcondition>} {"<string>"}

```

<time> may be any constant or expression. Choosing -1 will wait until a carriage return is hit, zero (0) won't pause at all, and a positive number will wait the specified number of seconds. The time is rounded to an integer number of seconds if subsecond time resolution is not supported by the given platform. **pause 0** is synonymous with **print**.

If the current terminal supports **mousing**, then **pause mouse** will terminate on either a mouse click or on ctrl-C. For all other terminals, or if mousing is not active, **pause mouse** is equivalent to **pause -1**.

If one or more end conditions are given after **pause mouse**, then any one of the conditions will terminate the pause. The possible end conditions are **keypress**, **button1**, **button2**, **button3**, **close**, and **any**. If the pause terminates on a keypress, then the ascii value of the key pressed is returned in `MOUSE_KEY`. The character itself is returned as a one character string in `MOUSE_CHAR`. Hotkeys (bind command) are disabled if keypress is one of the end conditions. Zooming is disabled if button3 is one of the end conditions.

In all cases the coordinates of the mouse are returned in variables `MOUSE_X`, `MOUSE_Y`, `MOUSE_X2`, `MOUSE_Y2`. See **mouse variables** (p. 38).

Note: Since **pause** communicates with the operating system rather than the graphics, it may behave differently with different device drivers (depending upon how text and graphics are mixed).

Examples:

```
pause -1    # Wait until a carriage return is hit
pause 3     # Wait three seconds
pause -1    "Hit return to continue"
pause 10    "Isn't this pretty? It's a cubic spline."
pause mouse "Click any mouse button on selected data point"
pause mouse keypress "Type a letter from A-F in the active window"
pause mouse button1,keypress
pause mouse any "Any key or button will terminate"
```

The variant "pause mouse key" will resume after any keypress in the active plot window. If you want to wait for a particular key to be pressed, you can use a reread loop such as:

```
print "I will resume after you hit the Tab key in the plot window"
load "wait_for_tab"
```

File "wait_for_tab" contains the lines

```
pause mouse key
if (MOUSE_KEY != 9) reread
```

Plot

plot is the primary command for drawing plots with **gnuplot**. It offers many different graphical representations for functions and data. **plot** is used to draw 2D functions and data. **splot** draws 2D projections of 3D surfaces and data.

Syntax:

```
plot {<ranges>} <plot-element> {, <plot-element>, <plot-element>}
```

Each plot element consists of a definition, a function, or a data source together with optional properties or modifiers:

```
plot-element:
    {<iteration>}
    <definition> | {<sampling-range>} <function> | <data source>
    {axes <axes>} {<title-spec>}
    {with <style>}
```

The graphical representation of each plot element is determined by the keyword **with**, e.g. **with lines** or **with boxplot**. See **plotting styles** (p. 44).

The data to be plotted is either generated by a function (two functions if in parametric mode), read from a data file, or read from a named data block that was defined previously. Multiple datafiles, data blocks, and/or functions may be plotted in a single plot command separated by commas. See **data** (p. 78), **inline data** (p. 81), **functions** (p. 89).

A plot-element that contains the definition of a function or variable does not create any visible output, see third example below.

Examples:

```

plot sin(x)
plot sin(x), cos(x)
plot f(x) = sin(x*a), a = .2, f(x), a = .4, f(x)
plot "datafile.1" with lines, "datafile.2" with points
plot [t=1:10] [-pi:pi*2] tan(t), \
    "data.1" using (tan($2)):(($3/$4) smooth csplines \
        axes x1y2 notitle with lines 5
plot for [datafile in "spinach.dat broccoli.dat"] datafile

```

See also **show plot** (p. 140).

Axes

There are four possible sets of axes available; the keyword `<axes>` is used to select the axes for which a particular line should be scaled. **x1y1** refers to the axes on the bottom and left; **x2y2** to those on the top and right; **x1y2** to those on the bottom and right; and **x2y1** to those on the top and left. Ranges specified on the **plot** command apply only to the first set of axes (bottom left).

Binary

BINARY DATA FILES:

It is necessary to provide the keyword **binary** after the filename. Adequate details of the file format must be given on the command line or extracted from the file itself for a supported binary **filetype**. In particular, there are two structures for binary files, binary matrix format and binary general format.

The **binary matrix** format contains a two dimensional array of 32 bit IEEE float values plus an additional column and row of coordinate values. In the **using** specifier of a plot command, column 1 refers to the matrix row coordinate, column 2 refers to the matrix column coordinate, and column 3 refers to the value stored in the array at those coordinates.

The **binary general** format contains an arbitrary number of columns for which information must be specified at the command line. For example, **array**, **record**, **format** and **using** can indicate the size, format and dimension of data. There are a variety of useful commands for skipping file headers and changing endianness. There are a set of commands for positioning and translating data since often coordinates are not part of the file when uniform sampling is inherent in the data. Unlike reading from a text or matrix binary file, general binary does not treat the generated columns as 1, 2 or 3 in the **using** list. Instead column 1 refers to column 1 of the file, or as specified in the **format** list.

There are global default settings for the various binary options which may be set using the same syntax as the options when used as part of the **(s)plot <filename> binary ...** command. This syntax is **set datafile binary ...**. The general rule is that common command-line specified parameters override file-extracted parameters which override default parameters.

Binary matrix is the default binary format when no keywords specific to **binary general** are given, i.e., **array**, **record**, **format**, **filetype**.

General binary data can be entered at the command line via the special file name `'-'`. However, this is intended for use through a pipe where programs can exchange binary data, not for keyboards. There is no "end of record" character for binary data. Gnuplot continues reading from a pipe until it has read the number of points declared in the **array** qualifier. See **binary matrix** (p. 177) or **binary general** (p. 75) for more details.

The **index** keyword is not supported, since the file format allows only one surface per file. The **every** and **using** filters are supported. **using** operates as if the data were read in the above triplet form.

[Binary File Splot Demo.](#)

General

The **binary** keyword appearing alone indicates a binary data file that contains both coordinate information describing a non-uniform grid and the value of each grid point (see **binary matrix** (p. 177)). Binary data in any other format requires additional keywords to describe the layout of the data. Unfortunately the syntax of these required additional keywords is convoluted. Nevertheless the general binary mode is particularly useful for application programs sending large amounts of data to gnuplot.

Syntax:

```
plot '<file_name>' {binary <binary list>} ...
splot '<file_name>' {binary <binary list>} ...
```

General binary format is activated by keywords in <binary list> pertaining to information about file structure, i.e., **array**, **record**, **format** or **filetype**. Otherwise, non-uniform matrix binary format is assumed. (See **binary matrix** (p. 177) for more details.)

NB: In previous versions of gnuplot there have been some differences between the interpretation of binary data keywords by **plot** and **splot**. Where the meanings differ, one or both may change in a future gnuplot version.

Gnuplot knows how to read a few standard binary file types that are fully self-describing, e.g. PNG images. Type **show datafile binary** at the command line for a list. Apart from these, you can think of binary data files as conceptually the same as text data. Each point has columns of information which are selected via the **using** specification. If no **format** string is specified, gnuplot will read in a number of binary values equal to the largest column given in the <using list>. For example, **using 1:3** will result in three columns being read, of which the second will be ignored. Certain plot types have an associated default using specification. For example, **with image** has a default of **using 1**, while **with rgbimage** has a default of **using 1:2:3**.

Array

Describes the sampling array dimensions associated with the binary file. The coordinates will be generated by gnuplot. A number must be specified for each dimension of the array. For example, **array=(10,20)** means the underlying sampling structure is two-dimensional with 10 points along the first (x) dimension and 20 points along the second (y) dimension. A negative number indicates that data should be read until the end of file. If there is only one dimension, the parentheses may be omitted. A colon can be used to separate the dimensions for multiple records. For example, **array=25:35** indicates there are two one-dimensional records in the file.

Note: Gnuplot version 4.2 used the syntax **array=128x128** rather than **array=(128,128)**. The older syntax is now deprecated.

Record

This keyword serves the same function as **array** and has the same syntax. However, **record** causes gnuplot to not generate coordinate information. This is for the case where such information may be included in one of the columns of the binary data file.

Skip

This keyword allows you to skip sections of a binary file. For instance, if the file contains a 1024 byte header before the start of the data region you would probably want to use

```
plot '<file_name>' binary skip=1024 ...
```

If there are multiple records in the file, you may specify a leading offset for each. For example, to skip 512 bytes before the 1st record and 256 bytes before the second and third records

```
plot '<file_name>' binary record=356:356:356 skip=512:256:256 ...
```

Format

The default binary format is a float. For more flexibility, the format can include details about variable sizes. For example, **format="%uchar%int%float"** associates an unsigned character with the first using column, an int with the second column and a float with the third column. If the number of size specifications is less than the greatest column number, the size is implicitly taken to be similar to the last given variable size.

Furthermore, similar to the **using** specification, the format can include discarded columns via the ***** character and have implicit repetition via a numerical repeat-field. For example, **format="%*2int%3float"** causes gnuplot to discard two ints before reading three floats. To list variable sizes, type **show datafile binary datasizes**. There are a group of names that are machine dependent along with their sizes in bytes for the particular compilation. There is also a group of names which attempt to be machine independent.

Endian

Often the endianness of binary data in the file does not agree with the endianness used by the platform on which gnuplot is running. Several words can direct gnuplot how to arrange bytes. For example **endian=little** means treat the binary file as having byte significance from least to greatest. The options are

```

    little:  least significant to greatest significance
           big:  greatest significance to least significance
    default:  assume file endianness is the same as compiler
    swap (swab): Interchange the significance.  (If things
                don't look right, try this.)

```

Gnuplot can support "middle" ("pdp") endian if it is compiled with that option.

Filetype

For some standard binary file formats gnuplot can extract all the necessary information from the file in question. As an example, "format=edf" will read ESRF Header File format files. For a list of the currently supported file formats, type **show datafile binary filetypes**.

There is a special file type called **auto** for which gnuplot will check if the binary file's extension is a quasi-standard extension for a supported format.

Command line keywords may be used to override settings extracted from the file. The settings from the file override any defaults. See **set datafile binary** (p. 111).

Avs **avs** is one of the automatically recognized binary file types for images. AVS is an extremely simple format, suitable mostly for streaming between applications. It consists of 2 longs (xwidth, ywidth) followed by a stream of pixels, each with four bytes of information alpha/red/green/blue.

Edf **edf** is one of the automatically recognized binary file types for images. EDF stands for ESRF Data Format, and it supports both edf and ehf formats (the latter means ESRF Header Format). More information on specifications can be found at

<http://www.edfplus.info/specs>

Png If gnuplot was configured to use the libgd library for png/gif/jpeg output, then it can also be used to read these same image types as binary files. You can use an explicit command

```
plot 'file.png' binary filetype=png
```

Or the file type will be recognized automatically from the extension if you have previously requested

```
set datafile binary filetype=auto
```

Keywords

The following keywords apply only when generating coordinates from binary data files. That is, the control mapping the individual elements of a binary array, matrix, or image to specific x/y/z positions.

Scan A great deal of confusion can arise concerning the relationship between how gnuplot scans a binary file and the dimensions seen on the plot. To lessen the confusion, conceptually think of gnuplot *always* scanning the binary file point/line/plane or fast/medium/slow. Then this keyword is used to tell gnuplot how to map this scanning convention to the Cartesian convention shown in plots, i.e., x/y/z. The qualifier for scan is a two or three letter code representing where point is assigned (first letter), line is assigned (second letter), and plane is assigned (third letter). For example, **scan=yx** means the fastest, point-by-point, increment should be mapped along the Cartesian y dimension and the middle, line-by-line, increment should be mapped along the x dimension.

When the plotting mode is **plot**, the qualifier code can include the two letters x and y. For **splot**, it can include the three letters x, y and z.

There is nothing restricting the inherent mapping from point/line/plane to apply only to Cartesian coordinates. For this reason there are cylindrical coordinate synonyms for the qualifier codes where t (theta), r and z are analogous to the x, y and z of Cartesian coordinates.

Transpose Shorthand notation for **scan=yx** or **scan=yxz**.

Dx, dy, dz When gnuplot generates coordinates, it uses the spacing described by these keywords. For example **dx=10 dy=20** would mean space samples along the x dimension by 10 and space samples along the y dimension by 20. **dy** cannot appear if **dx** does not appear. Similarly, **dz** cannot appear if **dy** does not appear. If the underlying dimensions are greater than the keywords specified, the spacing of the highest dimension given is extended to the other dimensions. For example, if an image is being read from a file and only **dx=3.5** is given gnuplot uses a delta x and delta y of 3.5.

The following keywords also apply only when generating coordinates. However they may also be used with matrix binary files.

Flipx, flipy, flipz Sometimes the scanning directions in a binary datafile are not consistent with that assumed by gnuplot. These keywords can flip the scanning direction along dimensions x, y, z.

Origin When gnuplot generates coordinates based upon transposition and flip, it attempts to always position the lower left point in the array at the origin, i.e., the data lies in the first quadrant of a Cartesian system after transpose and flip.

To position the array somewhere else on the graph, the **origin** keyword directs gnuplot to position the lower left point of the array at a point specified by a tuple. The tuple should be a double for **plot** and a triple for **splot**. For example, **origin=(100,100):(100,200)** is for two records in the file and intended for plotting in two dimensions. A second example, **origin=(0,0,3.5)**, is for plotting in three dimensions.

Center Similar to **origin**, this keyword will position the array such that its center lies at the point given by the tuple. For example, **center=(0,0)**. Center does not apply when the size of the array is **Inf**.

Rotate The transpose and flip commands provide some flexibility in generating and orienting coordinates. However, for full degrees of freedom, it is possible to apply a rotational vector described by a rotational angle in two dimensions.

The **rotate** keyword applies to the two-dimensional plane, whether it be **plot** or **splot**. The rotation is done with respect to the positive angle of the Cartesian plane.

The angle can be expressed in radians, radians as a multiple of pi, or degrees. For example, **rotate=1.5708**, **rotate=0.5pi** and **rotate=90deg** are equivalent.

If **origin** is specified, the rotation is done about the lower left sample point before translation. Otherwise, the rotation is done about the array **center**.

Perpendicular For **splot**, the concept of a rotational vector is implemented by a triple representing the vector to be oriented normal to the two-dimensional x-y plane. Naturally, the default is (0,0,1). Thus specifying both rotate and perpendicular together can orient data myriad ways in three-space.

The two-dimensional rotation is done first, followed by the three-dimensional rotation. That is, if R' is the rotational 2×2 matrix described by an angle, and P is the 3×3 matrix projecting (0,0,1) to (xp,yp,zp), let R be constructed from R' at the upper left sub-matrix, 1 at element 3,3 and zeros elsewhere. Then the matrix formula for translating data is $v' = P R v$, where v is the 3×1 vector of data extracted from the data file. In cases where the data of the file is inherently not three-dimensional, logical rules are used to place the data in three-space. (E.g., usually setting the z-dimension value to zero and placing 2D data in the x-y plane.)

Data

Discrete data contained in a file can be displayed by specifying the name of the data file (enclosed in single or double quotes) on the **plot** command line.

Syntax:

```
plot '<file_name>' {binary <binary list>}
                        {{nonuniform} matrix}
                        {index <index list> | index "<name>"}
                        {every <every list>}
                        {skip <number-of-lines>}
                        {using <using list>}
                        {smooth <option>}
                        {volatile} {noautoscale}
```

The modifiers **binary**, **index**, **every**, **skip**, **using**, and **smooth** are discussed separately. In brief, **binary** allows data entry from a binary file, **index** selects which data sets in a multi-data-set file are to be plotted, **every** specifies which points within a single data set are to be plotted, **using** determines how the columns within a single record are to be interpreted, and **smooth** allows for simple interpolation and approximation. **splot** has a similar syntax, but does not support the **smooth** option.

The **noautoscale** keyword means that the points making up this plot will be ignored when automatically determining axis range limits.

TEXT DATA FILES:

Data files should contain at least one data point per record (**using** can select one data point from the record). Records beginning with **#** (and also with **!** on VMS) will be treated as comments and ignored. Each data point represents an (x,y) pair. For **plots** with error bars or error bars with lines (see **errorbars** (p. 88) or **errorlines** (p. 89)), each data point is (x,y,ydelta), (x,y,ylow,yhigh), (x,y,xdelta), (x,y,xlow,xhigh), or (x,y,xlow,xhigh,ylow,yhigh).

In all cases, the numbers of each record of a data file must be separated by white space (one or more blanks or tabs) unless a format specifier is provided by the **using** option. This white space divides each record into columns. However, whitespace inside a pair of double quotes is ignored when counting columns, so the following datafile line has three columns:

```
1.0 "second column" 3.0
```

Data may be written in exponential format with the exponent preceded by the letter e or E. The fortran exponential specifiers d, D, q, and Q may also be used if the command **set datafile fortran** is in effect.

Only one column (the y value) need be provided. If x is omitted, **gnuplot** provides integer values starting at 0.

In datafiles, blank records (records with no characters other than blanks and a newline and/or carriage return) are significant.

Single blank records designate discontinuities in a **plot**; no line will join points separated by a blank records (if they are plotted with a line style).

Two blank records in a row indicate a break between separate data sets. See **index** (p. 80).

If autoscaling has been enabled (**set autoscale**), the axes are automatically extended to include all data-points, with a whole number of tic marks if tics are being drawn. This has two consequences: i) For **splot**, the corner of the surface may not coincide with the corner of the base. In this case, no vertical line is drawn. ii) When plotting data with the same x range on a dual-axis graph, the x coordinates may not coincide if the x2tics are not being drawn. This is because the x axis has been autoextended to a whole number of tics, but the x2 axis has not. The following example illustrates the problem:

```
reset; plot '-', '-' axes x2y1
1 1
19 19
e
1 1
19 19
e
```

To avoid this, you can use the **fixmin/fixmax** feature of the **set autoscale** command, which turns off the automatic extension of the axis range up to the next tic mark.

Label coordinates and text can also be read from a data file (see **labels** (p. 54)).

Every

The **every** keyword allows a periodic sampling of a data set to be plotted.

In the discussion a "point" is a datum defined by a single record in the file; "block" here will mean the same thing as "datablock" (see **glossary** (p. 32)).

Syntax:

```
plot 'file' every {<point_incr>
                  {:{<block_incr>}
                  {:{<start_point>}
                  {:{<start_block>}
                  {:{<end_point>}
                  {:{<end_block>}}}}}
```

The data points to be plotted are selected according to a loop from **<start_point>** to **<end_point>** with increment **<point_incr>** and the blocks according to a loop from **<start_block>** to **<end_block>** with increment **<block_incr>**.

The first datum in each block is numbered '0', as is the first block in the file.

Note that records containing unplotable information are counted.

Any of the numbers can be omitted; the increments default to unity, the start values to the first point or block, and the end values to the last point or block. ':' at the end of the **every** option is not permitted. If **every** is not specified, all points in all lines are plotted.

Examples:

```
every :::3::3      # selects just the fourth block ('0' is first)
every :::::9       # selects the first 10 blocks
every 2:2          # selects every other point in every other block
every ::5::15      # selects points 5 through 15 in each block
```

See

[simple plot demos \(simple.dem\)](#)

,

[Non-parametric splot demos](#)

, and

[Parametric splot demos](#)

Example datafile

This example plots the data in the file "population.dat" and a theoretical curve:

```
pop(x) = 103*exp((1965-x)/10)
set xrange [1960:1990]
plot 'population.dat', pop(x)
```

The file "population.dat" might contain:

```
# Gnu population in Antarctica since 1965
1965  103
1970   55
1975   34
1980   24
1985   10
```

Binary examples:

```
# Selects two float values (second one implicit) with a float value
# discarded between them for an indefinite length of 1D data.
plot '<file_name>' binary format="%float%*float" using 1:2 with lines

# The data file header contains all details necessary for creating
# coordinates from an EDF file.
plot '<file_name>' binary filetype=edf with image
plot '<file_name>.edf' binary filetype=auto with image

# Selects three unsigned characters for components of a raw RGB image
# and flips the y-dimension so that typical image orientation (start
# at top left corner) translates to the Cartesian plane. Pixel
# spacing is given and there are two images in the file. One of them
# is translated via origin.
plot '<file_name>' binary array=(512,1024):(1024,512) format='%uchar' \
    dx=2:1 dy=1:2 origin=(0,0):(1024,1024) flipy u 1:2:3 w rgbimage

# Four separate records in which the coordinates are part of the
# data file. The file was created with a endianness different from
# the system on which gnuplot is running.
splot '<file_name>' binary record=30:30:29:26 endian=swap u 1:2:3

# Same input file, but this time we skip the 1st and 3rd records
splot '<file_name>' binary record=30:26 skip=360:348 endian=swap u 1:2:3
```

See also [binary matrix](#) (p. [177](#)).

Index

The **index** keyword allows you to select specific data sets in a multi-data-set file for plotting.

Syntax:

```
plot 'file' index { <m>{:<n>{:<p>}} | "<name>" }
```


Data sets are separated by pairs of blank records. **index** <m> selects only set <m>; **index** <m>:<n> selects sets in the range <m> to <n>; and **index** <m>:<n>:<p> selects indices <m>, <m>+<p>, <m>+2<p>, etc., but stopping at <n>. Following C indexing, the index 0 is assigned to the first data set in the file. Specifying too large an index results in an error message. If <p> is specified but <n> is left blank then every <p>-th dataset is read until the end of the file. If **index** is not specified, the entire file is plotted as a single data set.

Example:

```
plot 'file' index 4:5
```

For each point in the file, the index value of the data set it appears in is available via the pseudo-column **column(-2)**. This leads to an alternative way of distinguishing individual data sets within a file as shown below. This is more awkward than the **index** command if all you are doing is selecting one data set for plotting, but is very useful if you want to assign different properties to each data set. See **pseudocolumns** (p. 87), **lc variable** (p. 34).

Example:

```
plot 'file' using 1:(column(-2)==4 ? $2 : NaN)          # very awkward
plot 'file' using 1:2:(column(-2)) linecolor variable # very useful!
```

index '<name>' selects the data set with name '<name>'. Names are assigned to data sets in comment lines. The comment character and leading white space are removed from the comment line. If the resulting line starts with <name>, the following data set is now named <name> and can be selected.

Example:

```
plot 'file' index 'Population'
```

Please note that every comment that starts with <name> will name the following data set. To avoid problems it may be useful to choose a naming scheme like '== Population ==' or '[Population]'.

Inline data

There are two mechanisms for embedding data into a stream of gnuplot commands. If the special filename '-' appears in a plot command, then the lines immediately following the plot command are interpreted as inline data. See **special-filenames** (p. 83). Data provided in this way can only be used once, by the plot command it follows.

The second mechanism defines a named data block as a here-document. The named data is persistent and may be referred to by more than one plot command. Example:

```
$Mydata << EOD
11 22 33 first line of data
44 55 66 second line of data
# comments work just as in a data file
77 88 99
EOD
stats $Mydata using 1:3
plot $Mydata using 1:3 with points, $Mydata using 1:2 with impulses
```

Data block names must begin with a \$ character, which distinguishes them from other types of persistent variables. The end-of-data delimiter (EOD in the example) may be any sequence of alphanumeric characters.

The storage associated with named data blocks can be released using **undefine** command. **undefine \$*** frees all named data blocks at once.

Skip

The **skip** keyword tells the program to skip lines at the start of a text (i.e. not binary) data file. The lines that are skipped do not count toward the line count used in processing the **every** keyword. Note that **skip N** skips lines only at the start of the file, whereas **every ::N** skips lines at the start of every data block in the file. See also **binary skip** (p. 75) for a similar option that applies to binary data files.

Smooth

gnuplot includes a few general-purpose routines for interpolation and approximation of data; these are grouped under the **smooth** option. More sophisticated data processing may be performed by preprocessing the data externally or by using **fit** with an appropriate model.

Syntax:

```
smooth {unique | frequency | cumulative | cnormal | kdensity {bandwidth}
      | csplines | acsplines | mcsplines | bezier | sbezier
      | unwrap}
```

unique, **frequency**, **cumulative** and **cnormal** plot the data after making them monotonic. **unwrap** manipulates the data to avoid jumps of more than π by adding or subtracting multiples of 2π . Each of the other routines uses the data to determine the coefficients of a continuous curve between the endpoints of the data. This curve is then plotted in the same manner as a function, that is, by finding its value at uniform intervals along the abscissa (see **set samples** (p. 151)) and connecting these points with straight line segments (if a line style is chosen).

If **autoscale** is in effect, the ranges will be computed such that the plotted curve lies within the borders of the graph.

If **autoscale** is not in effect, and the smooth option is either **acspline** or **cspline**, the sampling of the generated curve is done across the intersection of the x range covered by the input data and the fixed abscissa range as defined by **set xrange**.

If too few points are available to allow the selected option to be applied, an error message is produced. The minimum number is one for **unique** and **frequency**, four for **acsplines**, and three for the others.

The **smooth** options have no effect on function plots.

Acsplines The **acsplines** option approximates the data with a "natural smoothing spline". After the data are made monotonic in x (see **smooth unique** (p. 83)), a curve is piecewise constructed from segments of cubic polynomials whose coefficients are found by fitting to the individual data points weighted by the value, if any, given in the third column of the using spec. The default is equivalent to

```
plot 'data-file' using 1:2:(1.0) smooth acsplines
```

Qualitatively, the absolute magnitude of the weights determines the number of segments used to construct the curve. If the weights are large, the effect of each datum is large and the curve approaches that produced by connecting consecutive points with natural cubic splines. If the weights are small, the curve is composed of fewer segments and thus is smoother; the limiting case is the single segment produced by a weighted linear least squares fit to all the data. The smoothing weight can be expressed in terms of errors as a statistical weight for a point divided by a "smoothing factor" for the curve so that (standard) errors in the file can be used as smoothing weights.

Example:

```
sw(x,S)=1/(x*x*S)
plot 'data_file' using 1:2:(sw($3,100)) smooth acsplines
```

Bezier The **bezier** option approximates the data with a Bezier curve of degree n (the number of data points) that connects the endpoints.

Csplines The **csplines** option connects consecutive points by natural cubic splines after rendering the data monotonic (see **smooth unique** (p. 83)).

Mcsplines The **mcsplines** option connects consecutive points by cubic splines constrained such that the smoothed function preserves the monotonicity and convexity of the original data points. FN Fritsch & RE Carlson (1980) "Monotone Piecewise Cubic Interpolation", SIAM Journal on Numerical Analysis 17: 238-246.

Sbezier The **sbezier** option first renders the data monotonic (**unique**) and then applies the **bezier** algorithm.

Unique The **unique** option makes the data monotonic in x; points with the same x-value are replaced by a single point having the average y-value. The resulting points are then connected by straight line segments.

Unwrap The **unwrap** option modifies the input data so that any two successive points will not differ by more than π ; a point whose y value is outside this range will be incremented or decremented by multiples of 2π until it falls within π of the previous point. This operation is useful for making wrapped phase measurements continuous over time.

Frequency The **frequency** option makes the data monotonic in x; points with the same x-value are replaced by a single point having the summed y-values. To plot a histogram of the number of data values in equal size bins, set the y-value to 1.0 so that the sum is a count of occurrences in that bin: Example:

```
binwidth = <something> # set width of x values in each bin
bin(val) = binwidth * floor(val/binwidth)
plot "datafile" using (bin(column(1))):(1.0) smooth frequency
```

See also

[smooth.dem](#)

Cumulative The **cumulative** option makes the data monotonic in x; points with the same x-value are replaced by a single point containing the cumulative sum of y-values of all data points with lower x-values (i.e. to the left of the current data point). This can be used to obtain a cumulative distribution function from data. See also

[smooth.dem](#)

Cnormal The **cnormal** option makes the data monotonic in x and normalises the y-values onto the range [0:1]. Points with the same x-value are replaced by a single point containing the cumulative sum of y-values of all data points with lower x-values (i.e. to the left of the current data point) divided by the total sum of all y-values. This can be used to obtain a normalised cumulative distribution function from data (useful when comparing sets of samples with differing numbers of members). See also

[smooth.dem](#)

Kdensity The **kdensity** option is a way to plot a kernel density estimate (which is a smooth histogram) for a random collection of points, using Gaussian kernels. A Gaussian is placed at the location of each point in the first column and the sum of all these Gaussians is plotted as a function. The value in the second column is taken as weight of the Gaussian. To obtain a normalized histogram, this should be $1/\text{number-of-points}$. By default gnuplot calculates and uses the bandwidth which would be optimal for normally distributed data.

```
default_bandwidth = sigma * (4/3N) ** (0.2)
```

This will usually be a very conservative, i.e. broad bandwidth. Alternatively, you can provide an explicit bandwidth.

```
plot $DATA smooth kdensity bandwidth <value> with boxes
```

The bandwidth used in the previous plot is stored in variable GPVAL_KDENSITY_BANDWIDTH.

Special-filenames

There are a few filenames that have a special meaning: `''`, `'-'`, `'+'` and `'++'`.

The empty filename `''` tells gnuplot to re-use the previous input file in the same plot command. So to plot two columns from the same input file:

```
plot 'filename' using 1:2, '' using 1:3
```

The filename can also be reused over subsequent plot commands, however **save** then only records the name in a comment.

The special filenames '+' and '++' are a mechanism to allow the full range of **using** specifiers and plot styles with inline functions. Normally a function plot can only have a single y (or z) value associated with each sampled point. The pseudo-file '+' treats the sampled points as column 1, and allows additional column values to be specified via a **using** specification, just as for a true input file. By default samples are generated over the full range as set by **set xrange**, with the sampling controlled via **set samples**.

```
plot '+' using ($1):(sin($1)):(sin($1)**2) with filledcurves
```

An independent sampling range can be provided immediately before the '+'. Like in normal function plots, a name can be assigned to the independent variable. If given for the first plot element, the sampling range specifier has to be preceded by the **sample** keyword (see also **plot sampling (p. 91)**).

```
plot sample [beta=0:2*pi] '+' using (sin(beta)):(cos(beta)) with lines
```

Similarly the pseudo-file '++' returns 2 columns of data forming a regular grid of [x,y] coordinates with the number of points along x controlled by **set samples** and the number of points along y controlled by **set isosamples**. In parametric mode the samples are along u and v rather than along x and y. You must set **xrange** and **yrange** (or **urange** and **vrange**) before plotting '++'. Examples:

```
splot '++' using 1:2:(sin($1)*sin($2)) with pm3d
plot '++' using 1:2:(sin($1)*sin($2)) with image
```

The special filename '-' specifies that the data are inline; i.e., they follow the command. Only the data follow the command; **plot** options like filters, titles, and line styles remain on the **plot** command line. This is similar to << in unix shell script, and \$DECK in VMS DCL. The data are entered as though they are being read from a file, one data point per record. The letter "e" at the start of the first column terminates data entry. The **using** option can be applied to these data — using it to filter them through a function might make sense, but selecting columns probably doesn't!

'-' is intended for situations where it is useful to have data and commands together, e.g., when **gnuplot** is run as a sub-process of some front-end application. Some of the demos, for example, might use this feature. While **plot** options such as **index** and **every** are recognized, their use forces you to enter data that won't be used. For example, while

```
plot '-' index 0, '-' index 1
2
4
6

10
12
14
e
2
4
6

10
12
14
e
```

does indeed work,

```
plot '-', '-'
2
4
```

```
6
e
10
12
14
e
```

is a lot easier to type.

If you use '-' with **replot**, you may need to enter the data more than once. See **replot** (p. 96), **refresh** (p. 96).

A blank filename (') specifies that the previous filename should be reused. This can be useful with things like

```
plot 'a/very/long/filename' using 1:2, '' using 1:3, '' using 1:4
```

(If you use both '-' and '' on the same **plot** command, you'll need to have two sets of inline data, as in the example above.)

On systems with a popen function, the datafile can be piped through a shell command by starting the file name with a '<'. For example,

```
pop(x) = 103*exp(-x/10)
plot "< awk '{print $1-1965, $2}' population.dat", pop(x)
```

would plot the same information as the first population example but with years since 1965 as the x axis. If you want to execute this example, you have to delete all comments from the data file above or substitute the following command for the first part of the command above (the part up to the comma):

```
plot "< awk '$0 !~ /^#/ {print $1-1965, $2}' population.dat"
```

While this approach is most flexible, it is possible to achieve simple filtering with the **using** keyword.

On systems with an fdopen() function, data can be read from an arbitrary file descriptor attached to either a file or pipe. To read from file descriptor **n** use '<&n'. This allows you to easily pipe in several data files in a single call from a POSIX shell:

```
$ gnuplot -p -e "plot '<&3', '<&4'" 3<data-3 4<data-4
$ ./gnuplot 5< <(myprogram -with -options)
gnuplot> plot '<&5'
```

Thru

The **thru** keyword is deprecated.

Old syntax:

```
plot 'file' thru f(x)
```

Current syntax:

```
plot 'file' using 1:(f($2))
```

Using

The most common datafile modifier is **using**. It tells the program which columns of data in the input file are to be plotted.

Syntax:

```
plot 'file' using <entry> {:<entry> {:<entry> ...}} {'format'}
```

If a format is specified, it is used to read in each datafile record using the C library 'scanf' function. Otherwise the record is interpreted as consisting of columns (fields) of data separated by whitespace (spaces and/or tabs), but see **datafile separator** (p. 110).

Each <entry> may be a simple column number that selects the value from one field of the input file, a string that matches a column label in the first line of a data set, an expression enclosed in parentheses, or a special function not enclosed in parentheses such as `xticlabels(2)`.

If the entry is an expression in parentheses, then the function `column(N)` may be used to indicate the value in column N. That is, `column(1)` refers to the first item read, `column(2)` to the second, and so on. The special symbols \$1, \$2, ... are shorthand for `column(1)`, `column(2)` ... The function **valid(N)** tests whether the value in the Nth column is a valid number. If each column of data in the input file contains a label in the first row rather than a data value, this label can be used to identify the column on input and/or in the plot legend. The `column()` function can be used to select an input column by label rather than by column number. For example, if the data file contains

```
Height    Weight    Age
val1      val1      val1
...        ...        ...
```

then the following plot commands are all equivalent

```
plot 'datafile' using 3:1, '' using 3:2
plot 'datafile' using (column("Age")):(column(1)), \
    '' using (column("Age")):(column(2))
plot 'datafile' using "Age":"Height", '' using "Age":"Weight"
```

The full string must match. Comparison is case-sensitive. To use the column labels in the plot legend, use **set key autotitle columnhead**.

In addition to the actual columns 1...N in the input data file, gnuplot presents data from several "pseudo-columns" that hold bookkeeping information. E.g. \$0 or `column(0)` returns the sequence number of this data record within a dataset. Please see **pseudocolumns** (p. 87).

An empty <entry> will default to its order in the list of entries. For example, **using ::4** is interpreted as **using 1:2:4**.

If the **using** list has only a single entry, that <entry> will be used for y and the data point number (pseudo-column \$0) is used for x; for example, **"plot 'file' using 1"** is identical to **"plot 'file' using 0:1"**. If the **using** list has two entries, these will be used for x and y. See **set style** (p. 152) and **fit** (p. 62) for details about plotting styles that make use of data from additional columns of input.

'scanf' accepts several numerical specifications but **gnuplot** requires all inputs to be double-precision floating-point variables, so "%lf" is essentially the only permissible specifier. A format string given by the user must contain at least one such input specifier, and no more than seven of them. 'scanf' expects to see white space — a blank, tab ("\t"), newline ("\n"), or formfeed ("\f") — between numbers; anything else in the input stream must be explicitly skipped.

Note that the use of "\t", "\n", or "\f" requires use of double-quotes rather than single-quotes.

Using examples This creates a plot of the sum of the 2nd and 3rd data against the first: The format string specifies comma- rather than space-separated columns. The same result could be achieved by specifying **set datafile separator comma**.

```
plot 'file' using 1:($2+$3) '%lf,%lf,%lf'
```

In this example the data are read from the file "MyData" using a more complicated format:

```
plot 'MyData' using "%*lf%lf%*20[^\n]%lf"
```

The meaning of this format is:

```
%*lf      ignore a number
%lf        read a double-precision number (x by default)
%*20[^\n]  ignore 20 non-newline characters
%lf        read a double-precision number (y by default)
```

One trick is to use the ternary `?:` operator to filter data:

```
plot 'file' using 1:($3>10 ? $2 : 1/0)
```

which plots the datum in column two against that in column one provided the datum in column three exceeds ten. `1/0` is undefined; **gnuplot** quietly ignores undefined points, so unsuitable points are suppressed. Or you can use the pre-defined variable `NaN` to achieve the same result.

In fact, you can use a constant expression for the column number, provided it doesn't start with an opening parenthesis; constructs like **using** `0+(complicated expression)` can be used. The crucial point is that the expression is evaluated once if it doesn't start with a left parenthesis, or once for each data point read if it does.

If timeseries data are being used, the time can span multiple columns. The starting column should be specified. Note that the spaces within the time must be included when calculating starting columns for other data. E.g., if the first element on a line is a time with an embedded space, the y value should be specified as column three.

It should be noted that **plot 'file'**, **plot 'file' using 1:2**, and **plot 'file' using (\$1):(\$2)** can be subtly different: 1) if **file** has some lines with one column and some with two, the first will invent x values when they are missing, the second will quietly ignore the lines with one column, and the third will store an undefined value for lines with one point (so that in a plot with lines, no line joins points across the bad point); 2) if a line contains text at the first column, the first will abort the plot on an error, but the second and third should quietly skip the garbage.

In fact, it is often possible to plot a file with lots of lines of garbage at the top simply by specifying

```
plot 'file' using 1:2
```

However, if you want to leave text in your data files, it is safer to put the comment character (`#`) in the first column of the text lines.

Pseudocolumns Expressions in the **using** clause of a plot statement can refer to additional bookkeeping values in addition to the actual data values contained in the input file. These are contained in "pseudo-columns".

<code>column(0)</code>	The sequential order of each point within a data set. The counter starts at 0 and is reset by two sequential blank records. The shorthand form <code>\$0</code> is available.
<code>column(-1)</code>	This counter starts at 0 and is reset by a single blank line. This corresponds to the data line in array or grid data.
<code>column(-2)</code>	The index number of the current data set within a file that contains multiple data sets. See 'index'.

Xticlabels Axis tick labels can be generated via a string function, usually taking a data column as an argument. The simplest form uses the data column itself as a string. That is, `xticlabels(N)` is shorthand for `xticlabels(stringcolumn(N))`. This example uses the contents of column 3 as x-axis tick labels.

```
plot 'datafile' using <xcol>:<ycol>:xticlabels(3) with <plotstyle>
```

Axis tick labels may be generated for any of the plot axes: `x` `x2` `y` `y2` `z`. The **ticlabels(<labelcol>)** specifiers must come after all of the data coordinate specifiers in the **using** portion of the command. For each data point which has a valid set of `X,Y[,Z]` coordinates, the string value given to `xticlabels()` is added to the list of xtic labels at the same `X` coordinate as the point it belongs to. **xticlabels()** may be shortened to **xtic()** and so on.

Example:

```
splot "data" using 2:4:6:xtic(1):ytic(3):ztic(6)
```

In this example the x and y axis tic labels are taken from different columns than the x and y coordinate values. The z axis tics, however, are generated from the z coordinate of the corresponding point.

Example:

```
plot "data" using 1:2:xtic( $3 > 10. ? "A" : "B" )
```

This example shows the use of a string-valued function to generate x-axis tick labels. Each point in the data file generates a tick mark on x labeled either "A" or "B" depending on the value in column 3.

X2ticlabels See **plot using xticlabels** (p. 87).

Yticlabels See **plot using xticlabels** (p. 87).

Y2ticlabels See **plot using xticlabels** (p. 87).

Zticlabels See **plot using xticlabels** (p. 87).

Volatile

The **volatile** keyword in a plot command indicates that the data previously read from the input stream or file may not be available for re-reading. This tells the program to use **refresh** rather than **replot** commands whenever possible. See **refresh** (p. 96).

Errorbars

Error bars are supported for 2D data file plots by reading one to four additional columns (or **using** entries); these additional values are used in different ways by the various errorbar styles.

In the default situation, **gnuplot** expects to see three, four, or six numbers on each line of the data file — either

```
(x, y, ydelta),
(x, y, ylow, yhigh),
(x, y, xdelta),
(x, y, xlow, xhigh),
(x, y, xdelta, ydelta), or
(x, y, xlow, xhigh, ylow, yhigh).
```

The x coordinate must be specified. The order of the numbers must be exactly as given above, though the **using** qualifier can manipulate the order and provide values for missing columns. For example,

```
plot 'file' with errorbars
plot 'file' using 1:2:(sqrt($1)) with xerrorbars
plot 'file' using 1:2:($1-$3):($1+$3):4:5 with xyerrorbars
```

The last example is for a file containing an unsupported combination of relative x and absolute y errors. The **using** entry generates absolute x min and max from the relative error.

The y error bar is a vertical line plotted from (x, ylow) to (x, yhigh). If ydelta is specified instead of ylow and yhigh, ylow = y - ydelta and yhigh = y + ydelta are derived. If there are only two numbers on the record, yhigh and ylow are both set to y. The x error bar is a horizontal line computed in the same fashion. To get lines plotted between the data points, **plot** the data file twice, once with errorbars and once with lines (but remember to use the **notitle** option on one to avoid two entries in the key). Alternately, use the **errorlines** command (see **errorlines** (p. 89)).

The error bars have crossbars at each end unless **set bars** is used (see **set bars** (p. 102) for details).

If autoscaling is on, the ranges will be adjusted to include the error bars.

See also

[errorbar demos](#).

See **plot using** (p. 85), **plot with** (p. 93), and **set style** (p. 152) for more information.

Errorlines

Lines with error bars are supported for 2D data file plots by reading one to four additional columns (or **using** entries); these additional values are used in different ways by the various errorlines styles.

In the default situation, **gnuplot** expects to see three, four, or six numbers on each line of the data file — either

```
(x, y, ydelta),
(x, y, ylow, yhigh),
(x, y, xdelta),
(x, y, xlow, xhigh),
(x, y, xdelta, ydelta), or
(x, y, xlow, xhigh, ylow, yhigh).
```

The x coordinate must be specified. The order of the numbers must be exactly as given above, though the **using** qualifier can manipulate the order and provide values for missing columns. For example,

```
plot 'file' with errorlines
plot 'file' using 1:2:(sqrt($1)) with xerrorlines
plot 'file' using 1:2:($1-$3):($1+$3):4:5 with xyerrorlines
```

The last example is for a file containing an unsupported combination of relative x and absolute y errors. The **using** entry generates absolute x min and max from the relative error.

The y error bar is a vertical line plotted from (x, ylow) to (x, yhigh). If ydelta is specified instead of ylow and yhigh, ylow = y - ydelta and yhigh = y + ydelta are derived. If there are only two numbers on the record, yhigh and ylow are both set to y. The x error bar is a horizontal line computed in the same fashion.

The error bars have crossbars at each end unless **set bars** is used (see **set bars** (p. 102) for details).

If autoscaling is on, the ranges will be adjusted to include the error bars.

See **plot using** (p. 85), **plot with** (p. 93), and **set style** (p. 152) for more information.

Functions

Built-in or user-defined functions can be displayed by the **plot** and **splot** commands in addition to, or instead of, data read from a file. The requested function is evaluated by sampling at regular intervals spanning the independent axis range[s]. See **set samples** (p. 151) and **set isosamples** (p. 122). Example:

```
approx(ang) = ang - ang**3 / (3*2)
plot sin(x) title "sin(x)", approx(x) title "approximation"
```

To set a default plot style for functions, see **set style function** (p. 155). For information on built-in functions, see **expressions functions** (p. 24). For information on defining your own functions, see **user-defined** (p. 30).

Parametric

When in parametric mode (**set parametric**) mathematical expressions must be given in pairs for **plot** and in triplets for **splot**.

Examples:

```
plot sin(t),t**2
splot cos(u)*cos(v),cos(u)*sin(v),sin(u)
```

Data files are plotted as before, except any preceding parametric function must be fully specified before a data file is given as a plot. In other words, the x parametric function (**sin(t)** above) and the y parametric function (**t**2** above) must not be interrupted with any modifiers or data functions; doing so will generate a syntax error stating that the parametric function is not fully specified.

Other modifiers, such as **with** and **title**, may be specified only after the parametric function has been completed:

```
plot sin(t),t**2 title 'Parametric example' with linespoints
```

See also

[Parametric Mode Demos.](#)

Ranges

This section describes only the optional axis ranges that may appear as the very first items in a **plot** command. If present, these ranges override any range limits established by a previous **set range** statement. For optional ranges elsewhere in a **plot** command that limit sampling of an individual plot component see **sampling** (p. 91).

Syntax:

```
[{<dummy-var>=}{{<min>}:{<max>}}]
[{{<min>}:{<max>}}]
```

The first form applies to the independent variable (**xrange** or **trange**, if in parametric mode). The second form applies to dependent variables. <dummy-var> optionally establishes a new name for the independent variable. (The default name may be changed with **set dummy**.)

In non-parametric mode, ranges must be given in the order

```
plot [<xrange>] [<yrange>] [<x2range>] [<y2range>] ...
```

In parametric mode, ranges must be given in the order

```
plot [<trange>] [<xrange>] [<yrange>] [<x2range>] [<y2range>] ...
```

The following **plot** command shows setting **trange** to $[-\pi:\pi]$, **xrange** to $[-1.3:1.3]$ and **yrange** to $[-1:1]$ for the duration of the graph:

```
plot [-pi:pi] [-1.3:1.3] [-1:1] sin(t),t**2
```

* can be used to allow autoscaling of either of min and max. Use an empty range `[]` as a placeholder if necessary.

Ranges specified on the **plot** or **splot** command line affect only that one graph; use the **set xrange**, **set yrange**, etc., commands to change the default ranges for future graphs.

For time data you must provide the range in quotes, using the same format used to read time from the datafile. See **set timefmt** (p. 161).

Examples:

This uses the current ranges:

```
plot cos(x)
```

This sets the x range only:

```
plot [-10:30] sin(pi*x)/(pi*x)
```

This is the same, but uses t as the dummy-variable:

```
plot [t = -10 :30] sin(pi*t)/(pi*t)
```

This sets both the x and y ranges:

```
plot [-pi:pi] [-3:3] tan(x), 1/x
```

This sets only the y range:

```
plot [ ] [-2:sin(5)*-8] sin(x)**besj0(x)
```

This sets xmax and ymin only:

```
plot [:200] [-pi:] $mydata using 1:2
```

This sets the x range for a timeseries:

```
set timefmt "%d/%m/%y %H:%M"
plot ["1/6/93 12:00":"5/6/93 12:00"] 'timedata.dat'
```

Sampling

By default, computed functions or data generated for the pseudo-file "+" are sampled over the entire range of the plot. This range may have been specified by a prior **set xrange** command, by an explicit global range specifier at the very start of the plot command, or by autoscaling of the range to span data seen in all the elements of this plot command. However, individual plot components can be assigned a more restricted sampling range.

Examples:

This establishes a total range on x running from 0 to 1000 and then plots data from a file and two functions each spanning a portion of the total range:

```
plot [0:1000] 'datafile', [0:200] func1(x), [200:500] func2(x)
```

This is similar except that the total range is established by the contents of the data file. In this case the sampled functions may or may not be entirely contained in the plot:

```
set autoscale x
plot 'datafile', [0:200] func1(x), [200:500] func2(x)
```

This command is ambiguous. The initial range will be interpreted as applying to the entire plot, not solely to the sampling of the first function as was probably the intent:

```
plot [0:10] f(x), [10:20] g(x), [20:30] h(x)
```

This command removes the ambiguity of the previous example by inserting the keyword **sample** so that the range is not applied to the entire plot:

```
plot sample [0:10] f(x), [10:20] g(x), [20:30] h(x)
```

This example shows one way of tracing out a helix in a 3D plot

```
splot [-2:2][-2:2] sample [h=1:10] '+' using (cos(h)):(sin(h)):(h)
```

For loops in plot command

If many similar files or functions are to be plotted together, it may be convenient to do so by iterating over a shared plot command.

Syntax:

```
plot for [<variable> = <start> : <end> {:<increment>}]
plot for [<variable> in "string of words"]
```

The scope of an iteration ends at the next comma or the end of the command, whichever comes first. Therefore iteration does not work for plots in parametric mode.

This will plot one curve, $\sin(3x)$, because iteration ends at the comma

```
plot for [i=1:3] j=i, sin(j*x)
```

This will plot three curves because there is no comma after the definition of j

```
plot for [i=1:3] j=i sin(j*x)
```

Example:

```
plot for [dataset in "apples bananas"] dataset."dat" title dataset
```

In this example iteration is used both to generate a file name and a corresponding title.

Example:

```
file(n) = sprintf("dataset_%d.dat",n)
splot for [i=1:10] file(i) title sprintf("dataset %d",i)
```

This example defines a string-valued function that generates file names, and plots ten such files together. The iteration variable ('i' in this example) is treated as an integer, and may be used more than once.

Example:

```
set key left
plot for [n=1:4] x**n sprintf("%d",n)
```

This example plots a family of functions.

Example:

```
list = "apple banana cabbage daikon eggplant"
item(n) = word(list,n)
plot for [i=1:words(list)] item[i]."dat" title item(i)
list = "new stuff"
replot
```

This example steps through a list and plots once per item. Because the items are retrieved dynamically, you can change the list and then replot.

Example:

```
list = "apple banana cabbage daikon eggplant"
plot for [i in list] i."dat" title i
list = "new stuff"
replot
```

This example does exactly the same thing as the previous example, but uses the string iterator form of the command rather than an integer iterator.

Title

By default each plot is listed in the key by the corresponding function or file name. You can give an explicit plot title instead using the **title** option.

Syntax:

```
title <text> | notitle [<ignored text>]
title columnheader | title columnheader(N)
      {at {beginning|end}}
```

where <text> is a quoted string or an expression that evaluates to a string. The quotes will not be shown in the key.

There is also an option that will interpret the first entry in a column of input data (i.e. the column header) as a text field, and use it as the key title. See **datastrings** (p. 21). This can be made the default by specifying **set key autotitle columnhead**.

The line title and sample can be omitted from the key by using the keyword **notitle**. A null title (**title ''**) is equivalent to **notitle**. If only the sample is wanted, use one or more blanks (**title ' '**). If **notitle** is followed by a string this string is ignored.

If **key autotitles** is set (which is the default) and neither **title** nor **notitle** are specified the line title is the function name or the file name as it appears on the **plot** command. If it is a file name, any datafile modifiers specified will be included in the default title.

The layout of the key itself (position, title justification, etc.) can be controlled by **set key**. Please see **set key** (p. 122) for details.

If you want the title of a plotted line to be placed immediately before or after that line in the graph itself, use **at {beginning|end}**. This option may be useful when plotting **with lines** but makes little sense for some other plot styles.

Examples:

This plots $y=x$ with the title 'x':

```
plot x
```

This plots x squared with title " x^2 " and file "data.1" with title "measured data":

```
plot x**2 title "x^2", 'data.1' t "measured data"
```

This puts an untitled circular border around a polar graph:

```
set polar; plot my_function(t), 1 notitle
```

Plot multiple columns of data, each of which contains its own title on the first line of the file. Place the titles after the corresponding lines rather than in a separate key:

```
unset key
set offset 0, graph 0.1
plot for [i=1:4] 'data' using i with lines title columnhead at end
```

With

Functions and data may be displayed in one of a large number of styles. The **with** keyword provides the means of selection.

Syntax:

```
with <style> { {linestyle | ls <line_style>}
              | {{linetype | lt <line_type>}
                {linewidth | lw <line_width>}
                {linecolor | lc <colorspec>}
                {pointtype | pt <point_type>}
                {pointsize | ps <point_size>}
                {fill | fs <fillstyle>}
                {nohidden3d} {nocontours} {nosurface}
                {palette}}
            }
```

where <style> is one of

lines	dots	steps	errorbars	xerrorbar	xyerrorlines
points	impulses	fsteps	errorlines	xerrorlines	yerrorbars
linespoints	labels	histeps	financebars	xyerrorbars	yerrorlines
surface	vectors	parallellaxes			

or

boxes	boxplot	ellipses	image
boxerrorbars	candlesticks	filledcurves	rgbimage
boxxyerrorbars	circles	histograms	rgbalpha pm3d

The first group of styles have associated line, point, and text properties. The second group of styles also have fill properties. See **fillstyle** (p. 155). Some styles have further sub-styles. See **plotting styles** (p. 44) for details of each.

A default style may be chosen by **set style function** and **set style data**.

By default, each function and data file will use a different line type and point type, up to the maximum number of available types. All terminal drivers support at least six different point types, and re-use them, in order, if more are required. To see the complete set of line and point types available for the current terminal, type **test** (p. 182).

If you wish to choose the line or point type for a single plot, <line_type> and <point_type> may be specified. These are positive integer constants (or expressions) that specify the line type and point type to be used for the plot. Use **test** to display the types available for your terminal.

You may also scale the line width and point size for a plot by using <line_width> and <point_size>, which are specified relative to the default values for each terminal. The pointsize may also be altered globally — see **set pointsize** (p. 149) for details. But note that both <point_size> as set here and as set by **set pointsize** multiply the default point size — their effects are not cumulative. That is, **set pointsize 2; plot x w p ps 3** will use points three times default size, not six.

It is also possible to specify **pointsize variable** either as part of a line style or for an individual plot. In this case one extra column of input is required, i.e. 3 columns for a 2D plot and 4 columns for a 3D plot.

The size of each individual point is determined by multiplying the global `pointsize` by the value read from the data file.

If you have defined specific line type/width and point type/size combinations with **set style line**, one of these may be selected by setting `<line_style>` to the index of the desired style.

If gnuplot was built with **pm3d** support, the special keyword **palette** is allowed for smooth color change of lines, points and dots in **splots**. The color is chosen from a smooth palette which was set previously with the command **set palette**. The color value corresponds to the z-value of the point coordinates or to the color coordinate if specified by the 4th parameter in **using**. Both 2D and 3D plots (**plot** and **splot** commands) can use palette colors as specified by either their fractional value or the corresponding value mapped to the colorbox range. A palette color value can also be read from an explicitly specified input column in the **using** specifier. See **colors** (p. 33), **set palette** (p. 144), **linetype** (p. 128).

The keyword **nohidden3d** applies only to plots made with the **splot** command. Normally the global option **set hidden3d** applies to all plots in the graph. You can attach the **nohidden3d** option to any individual plots that you want to exclude from the hidden3d processing. The individual elements other than surfaces (i.e. lines, dots, labels, ...) of a plot marked **nohidden3d** will all be drawn, even if they would normally be obscured by other plot elements.

Similarly, the keyword **nocontours** will turn off contouring for an individual plot even if the global property **set contour** is active.

Similarly, the keyword **nosurface** will turn off the 3D surface for an individual plot even if the global property **set surface** is active.

The keywords may be abbreviated as indicated.

Note that the **linewidth**, **pointsize** and **palette** options are not supported by all terminals.

Examples:

This plots $\sin(x)$ with impulses:

```
plot sin(x) with impulses
```

This plots x with points, x^2 with the default:

```
plot x w points, x**2
```

This plots $\tan(x)$ with the default function style, file "data.1" with lines:

```
plot [ ] [-2:5] tan(x), 'data.1' with l
```

This plots "leastsq.dat" with impulses:

```
plot 'leastsq.dat' w i
```

This plots the data file "population" with boxes:

```
plot 'population' with boxes
```

This plots "exper.dat" with errorbars and lines connecting the points (errorbars require three or four columns):

```
plot 'exper.dat' w lines, 'exper.dat' notitle w errorbars
```

Another way to plot "exper.dat" with errorlines (errorbars require three or four columns):

```
plot 'exper.dat' w errorlines
```

This plots $\sin(x)$ and $\cos(x)$ with linespoints, using the same line type but different point types:

```
plot sin(x) with linesp lt 1 pt 3, cos(x) with linesp lt 1 pt 4
```

This plots file "data" with points of type 3 and twice usual size:

```
plot 'data' with points pointtype 3 pointsize 2
```

This plots file "data" with variable pointsize read from column 4

```
plot 'data' using 1:2:4 with points pt 5 pointsize variable
```

This plots two data sets with lines differing only by weight:

```
plot 'd1' t "good" w l lt 2 lw 3, 'd2' t "bad" w l lt 2 lw 1
```

This plots filled curve of $x*x$ and a color stripe:

```
plot x*x with filledcurve closed, 40 with filledcurve y=10
```

This plots $x*x$ and a color box:

```
plot x*x, (x>=-5 && x<=5 ? 40 : 1/0) with filledcurve y=10 lt 8
```

This plots a surface with color lines:

```
splot x*x-y*y with line palette
```

This plots two color surfaces at different altitudes:

```
splot x*x-y*y with pm3d, x*x+y*y with pm3d at t
```

Print

The **print** command prints the value of `<expression>` to the screen. It is synonymous with **pause 0**. `<expression>` may be anything that **gnuplot** can evaluate that produces a number, or it can be a string.

Syntax:

```
print <expression> {, <expression>, ...}
```

See **expressions** (p. 24). The output file can be set with **set print**. See also **printerr** (p. 95).

Printerr

printerr is the same as **print** except that output is always sent to stderr even if a prior **set print** command remains in effect.

Pwd

The **pwd** command prints the name of the working directory to the screen.

Note that if you wish to store the current directory into a string variable or use it in string expressions, then you can use variable `GPVAL_PWD`, see **show variables all** (p. 163).

Quit

The **exit** and **quit** commands and END-OF-FILE character will exit **gnuplot**. Each of these commands will clear the output device (as does the **clear** command) before exiting.

Raise

Syntax:

```
raise {plot_window_nb}
```

The **raise** command raises (opposite to **lower**) plot window(s) associated with the interactive terminal of your gnuplot session, i.e. **pm**, **win**, **wxt** or **x11**. It puts the plot window to front (top) in the z-order windows stack of the window manager of your desktop.

As **x11** and **wxt** support multiple plot windows, then by default they raise these windows in descending order of most recently created on top to the least recently created on bottom. If a plot number is supplied as an optional parameter, only the associated plot window will be raised if it exists.

The optional parameter is ignored for single plot-windows terminal, i.e. **pm** and **win**.

If the window is not raised under X11, then perhaps the plot window is running in a different X11 session (telnet or ssh session, for example), or perhaps raising is blocked by your window manager policy setting.

Refresh

The **refresh** command is similar to **replot**, with two major differences. **refresh** reformats and redraws the current plot using the data already read in. This means that you can use **refresh** for plots with inline data (pseudo-device '-') and for plots from datafiles whose contents are volatile. You cannot use the **refresh** command to add new data to an existing plot.

Mousing operations, in particular zoom and unzoom, will use **refresh** rather than **replot** if appropriate. Example:

```
plot 'datafile' volatile with lines, '-' with labels
100 200 "Special point"
e
# Various mousing operations go here
set title "Zoomed in view"
set term post
set output 'zoom.ps'
refresh
```

Replot

The **replot** command without arguments repeats the last **plot** or **splot** command. This can be useful for viewing a plot with different **set** options, or when generating the same plot for several devices.

Arguments specified after a **replot** command will be added onto the last **plot** or **splot** command (with an implied ',' separator) before it is repeated. **replot** accepts the same arguments as the **plot** and **splot** commands except that ranges cannot be specified. Thus you can use **replot** to plot a function against the second axes if the previous command was **plot** but not if it was **splot**.

N.B. — use of

```
plot '-' ; ... ; replot
```

is not recommended, because it will require that you type in the data all over again. In most cases you can use the **refresh** command instead, which will redraw the plot using the data previously read in.

Note that in multiplot mode, **replot** can only reproduce the most recent component plot, not the full set.

See also **command-line-editing** (p. 20) for ways to edit the last **plot** (p. 73) (**splot** (p. 176)) command.

See also **show plot** (p. 140) to show the whole current plotting command, and the possibility to copy it into the **history** (p. 70).

Reread

The **reread** command causes the current **gnuplot** command file, as specified by a **load** command or on the command line, to be reset to its starting point before further commands are read from it. This essentially implements an endless loop of the commands from the beginning of the command file to the **reread** command. (But this is not necessarily a disaster — **reread** can be very useful when used in conjunction with **if**.) The **reread** command has no effect if input from standard input.

Examples:

Suppose the file "looper" contains the commands


```

a=a+1
plot sin(x*a)
pause -1
if(a<5) reread

```

and from within **gnuplot** you submit the commands

```

a=0
load 'looper'

```

The result will be five plots (separated by the **pause** message).

Suppose the file "data" contains six columns of numbers with a total yrange from 0 to 10; the first is x and the next are five different functions of x. Suppose also that the file "plotter" contains the commands

```

c_p = c_p+1
plot "$0" using 1:c_p with lines linetype c_p
if(c_p < n_p) reread

```

and from within **gnuplot** you submit the commands

```

n_p=6
c_p=1
unset key
set yrange [0:10]
set multiplot
call 'plotter' 'data'
unset multiplot

```

The result is a single graph consisting of five plots. The yrange must be set explicitly to guarantee that the five separate graphs (drawn on top of each other in multiplot mode) will have exactly the same axes. The linetype must be specified; otherwise all the plots would be drawn with the same type. See `animate.dem` in `demo` directory for an animated example.

Reset

The **reset** command causes all graph-related options that can be set with the **set** command to return to their default values. This command can be used to restore the default settings after executing a loaded command file, or to return to a defined state after lots of settings have been changed.

The following are *not* affected by **reset**.

```

'set term' 'set output' 'set loadpath' 'set fontpath' 'set linetype'
'set encoding' 'set decimalsign' 'set locale' 'set psdir' 'set fit'

```

Note that **reset** does not necessarily return settings to the state they were in at program entry, because the default values may have been altered by commands in the initialization files `gnuplotrc` or `$HOME/.gnuplot`. However, these commands can be re-executed by using the variant command **reset session**.

reset session deletes any user-defined variables and functions, restores default settings, and then re-executes the system-wide `gnuplotrc` initialization file and any private `$HOME/.gnuplot` preferences file. See **initialization** (p. 39).

reset errors clears only the error state variables `GPVAL_ERRNO` and `GPVAL_ERRMSG`.

reset bind restores all hotkey bindings to their default state.

Save

The **save** command saves user-defined functions, variables, the **set term** status, all **set** options, or all of these, plus the last **plot** (**splot**) command to the specified file.

Syntax:

```
save {<option>} '<filename>'
```

where <option> is **functions**, **variables**, **terminal** or **set**. If no option is used, **gnuplot** saves functions, variables, **set** options and the last **plot** (**splot**) command.

saved files are written in text format and may be read by the **load** command. For **save** with the **set** option or without any option, the **terminal** choice and the **output** filename are written out as a comment, to get an output file that works in other installations of gnuplot, without changes and without risk of unwillingly overwriting files.

save terminal will write out just the **terminal** status, without the comment marker in front of it. This is mainly useful for switching the **terminal** setting for a short while, and getting back to the previously set terminal, afterwards, by loading the saved **terminal** status. Note that for a single gnuplot session you may rather use the other method of saving and restoring current terminal by the commands **set term push** and **set term pop**, see **set term** (p. 159).

The filename must be enclosed in quotes.

The special filename "-" may be used to **save** commands to standard output. On systems which support a popen function (Unix), the output of save can be piped through an external program by starting the file name with a '|'. This provides a consistent interface to **gnuplot**'s internal settings to programs which communicate with **gnuplot** through a pipe. Please see help for **batch/interactive** (p. 19) for more details.

Examples:

```
save 'work.gnu'
save functions 'func.dat'
save var 'var.dat'
save set 'options.dat'
save term 'myterm.gnu'
save '-'
save '|grep title >t.gp'
```

Set-show

The **set** command can be used to set *lots* of options. No screen is drawn, however, until a **plot**, **splot**, or **replot** command is given.

The **show** command shows their settings; **show all** shows all the settings.

Options changed using **set** can be returned to the default state by giving the corresponding **unset** command. See also the **reset** (p. 97) command, which returns all settable parameters to default values.

The **set** and **unset** commands may optionally contain an iteration clause. See **plot for** (p. 91).

Angles

By default, **gnuplot** assumes the independent variable in polar graphs is in units of radians. If **set angles degrees** is specified before **set polar**, then the default range is [0:360] and the independent variable has units of degrees. This is particularly useful for plots of data files. The angle setting also applies to 3D mapping as set via the **set mapping** command.

Syntax:

```
set angles {degrees | radians}
show angles
```

The angle specified in **set grid polar** is also read and displayed in the units specified by **set angles**.

set angles also affects the arguments of the machine-defined functions sin(x), cos(x) and tan(x), and the outputs of asin(x), acos(x), atan(x), atan2(x), and arg(x). It has no effect on the arguments of hyperbolic functions or Bessel functions. However, the output arguments of inverse hyperbolic functions of complex arguments are affected; if these functions are used, **set angles radians** must be in effect to maintain consistency between input and output arguments.

```
x={1.0,0.1}
set angles radians
y=sinh(x)
print y      #prints {1.16933, 0.154051}
print asinh(y) #prints {1.0, 0.1}
```

but

```
set angles degrees
y=sinh(x)
print y      #prints {1.16933, 0.154051}
print asinh(y) #prints {57.29578, 5.729578}
```

See also

[poldat.dem: polar plot using set angles demo.](#)

Arrow

Arbitrary arrows can be placed on a plot using the **set arrow** command.

Syntax:

```
set arrow {<tag>} from <position> to <position>
set arrow {<tag>} from <position> rto <position>
set arrow {<tag>} from <position> length <coord> angle <ang>
set arrow <tag> arrowstyle | as <arrow_style>
set arrow <tag> {nohead | head | backhead | heads}
                  {size <headlength>,<headangle>{,<backangle>}}
                  {filled | empty | nofilled | noborder}
                  {front | back}
                  {linestyle <line_style>}
                  {linetype <line_type>} {linewidth <line_width>}

unset arrow {<tag>}
show arrow {<tag>}
```

<tag> is an integer that identifies the arrow. If no tag is given, the lowest unused tag value is assigned automatically. The tag can be used to delete or change a specific arrow. To change any attribute of an existing arrow, use the **set arrow** command with the appropriate tag and specify the parts of the arrow to be changed.

The position of the first end point of the arrow is always specified by "from". The other end point can be specified using any of three different mechanisms. The <position>s are specified by either x,y or x,y,z, and may be preceded by **first**, **second**, **graph**, **screen**, or **character** to select the coordinate system. Unspecified coordinates default to 0. See **coordinates** (p. 21) for details. A coordinate system specifier does not carry over from the first endpoint description the second.

- 1) "to <position>" specifies the absolute coordinates of the other end.
- 2) "rto <position>" specifies an offset to the "from" position. For linear axes, **graph** and **screen** coordinates, the distance between the start and the end point corresponds to the given relative coordinate. For logarithmic axes, the relative given coordinate corresponds to the factor of the coordinate between start and end point. Thus, a negative relative value or zero are not allowed for logarithmic axes.
- 3) "length <coordinate> angle <angle>" specifies the orientation of the arrow in the plane of the graph. Again any of the coordinate systems can be used to specify the length. The angle is always in degrees.

Other characteristics of the arrow can either be specified as a pre-defined arrow style or by providing them in **set arrow** command. For a detailed explanation of arrow characteristics, see **arrowstyle** (p. 153).

Examples:

To set an arrow pointing from the origin to (1,2) with user-defined linestyle 5, use:

```
set arrow to 1,2 ls 5
```

To set an arrow from bottom left of plotting area to (-5,5,3), and tag the arrow number 3, use:

```
set arrow 3 from graph 0,0 to -5,5,3
```

To change the preceding arrow to end at 1,1,1, without an arrow head and double its width, use:

```
set arrow 3 to 1,1,1 nohead lw 2
```

To draw a vertical line from the bottom to the top of the graph at x=3, use:

```
set arrow from 3, graph 0 to 3, graph 1 nohead
```

To draw a vertical arrow with T-shape ends, use:

```
set arrow 3 from 0,-5 to 0,5 heads size screen 0.1,90
```

To draw an arrow relatively to the start point, where the relative distances are given in graph coordinates, use:

```
set arrow from 0,-5 rto graph 0.1,0.1
```

To draw an arrow with relative end point in logarithmic x axis, use:

```
set logscale x
set arrow from 100,-5 rto 10,10
```

This draws an arrow from 100,-5 to 1000,5. For the logarithmic x axis, the relative coordinate 10 means "factor 10" while for the linear y axis, the relative coordinate 10 means "difference 10".

To delete arrow number 2, use:

```
unset arrow 2
```

To delete all arrows, use:

```
unset arrow
```

To show all arrows (in tag order), use:

```
show arrow
```

```
arrows demos.
```

Autoscale

Autoscaling may be set individually on the x, y or z axis or globally on all axes. The default is to autoscale all axes. If you want to autoscale based on a subset of the plots in the figure, you can mark the other ones with the flag **noautoscale**. See **datafile** (p. 78).

Syntax:

```
set autoscale {<axes>{|min|max|fixmin|fixmax|fix} | fix | keepfix}
set autoscale noextend
unset autoscale {<axes>}
show autoscale
```

where <axes> is either **x**, **y**, **z**, **cb**, **x2**, **y2** or **xy**. A keyword with **min** or **max** appended (this cannot be done with **xy**) tells **gnuplot** to autoscale just the minimum or maximum of that axis. If no keyword is given, all axes are autoscaled.

By default autoscaling sets the axis range limits to the nearest tic label position that includes all the plot data. Keywords **fixmin**, **fixmax**, **fix** or **noextend** tell **gnuplot** to disable extension of the axis range to the next tic mark position. In this case the axis range limit exactly matches the coordinate of the most extreme data point. **set autoscale noextend** is a synonym for **set autoscale fix**. Range extension for a single axis can be disabled by appending the **noextend** keyword to the corresponding range command, e.g.

```
set yrange [0:~] noextend
```

set autoscale keepfix autoscales all axes while leaving the fix settings unchanged.

When autoscaling, the axis range is automatically computed and the dependent axis (y for a **plot** and z for **splot**) is scaled to include the range of the function or data being plotted.

If autoscaling of the dependent axis (y or z) is not set, the current y or z range is used.

Autoscaling the independent variables (x for **plot** and x,y for **splot**) is a request to set the domain to match any data file being plotted. If there are no data files, autoscaling an independent variable has no effect. In other words, in the absence of a data file, functions alone do not affect the x range (or the y range if plotting $z = f(x,y)$).

Please see **set xrange** (p. 167) for additional information about ranges.

The behavior of autoscaling remains consistent in parametric mode, (see **set parametric** (p. 140)). However, there are more dependent variables and hence more control over x, y, and z axis scales. In parametric mode, the independent or dummy variable is t for **plots** and u,v for **splots**. **autoscale** in parametric mode, then, controls all ranges (t, u, v, x, y, and z) and allows x, y, and z to be fully autoscaled.

Autoscaling works the same way for polar mode as it does for parametric mode for **plot**, with the extension that in polar mode **set dummy** can be used to change the independent variable from t (see **set dummy** (p. 113)).

When tics are displayed on second axes but no plot has been specified for those axes, x2range and y2range are inherited from xrange and yrange. This is done *before* applying offsets or autoextending the ranges to a whole number of tics, which can cause unexpected results. To prevent this you can explicitly link the secondary axis range to the primary axis range. See **set link** (p. 128).

Examples:

This sets autoscaling of the y axis (other axes are not affected):

```
set autoscale y
```

This sets autoscaling only for the minimum of the y axis (the maximum of the y axis and the other axes are not affected):

```
set autoscale ymin
```

This disables extension of the x2 axis tics to the next tic mark, thus keeping the exact range as found in the plotted data and functions:

```
set autoscale x2fixmin
set autoscale x2fixmax
```

This sets autoscaling of the x and y axes:

```
set autoscale xy
```

This sets autoscaling of the x, y, z, x2 and y2 axes:

```
set autoscale
```

This disables autoscaling of the x, y, z, x2 and y2 axes:

```
unset autoscale
```

This disables autoscaling of the z axis only:

```
unset autoscale z
```

Parametric mode

When in parametric mode (**set parametric**), the xrange is as fully scalable as the y range. In other words, in parametric mode the x axis can be automatically scaled to fit the range of the parametric function that is being plotted. Of course, the y axis can also be automatically scaled just as in the non-parametric case. If autoscaling on the x axis is not set, the current x range is used.

Data files are plotted the same in parametric and non-parametric mode. However, there is a difference in mixed function and data plots: in non-parametric mode with autoscaled x, the x range of the datafile controls the x range of the functions; in parametric mode it has no influence.

For completeness a last command **set autoscale t** is accepted. However, the effect of this "scaling" is very minor. When **gnuplot** determines that the *t* range would be empty, it makes a small adjustment if autoscaling is true. Otherwise, **gnuplot** gives an error. Such behavior may, in fact, not be very useful and the command **set autoscale t** is certainly questionable.

splot extends the above ideas as you would expect. If autoscaling is set, then *x*, *y*, and *z* ranges are computed and each axis scaled to fit the resulting data.

Polar mode

When in polar mode (**set polar**), the *xrange* and the *yrange* may be left in autoscale mode. If **set rrange** is used to limit the extent of the polar axis, then *xrange* and *yrange* will adjust to match this automatically. However, explicit *xrange* and *yrange* commands can later be used to make further adjustments. See **set rrange** (p. 151). The *trange* may also be autoscaled. Note that if the *trange* is contained within one quadrant, autoscaling will produce a polar plot of only that single quadrant.

Explicitly setting one or two ranges but not others may lead to unexpected results. See also

[polar demos](#).

Bars

The **set bars** command controls the tics at the ends of error bars, and also at the end of the whiskers belonging to a boxplot.

Syntax:

```
set bars {small | large | fullwidth | <size>} {front | back}
unset bars
show bars
```

small is a synonym for 0.0, and **large** for 1.0. The default is 1.0 if no size is given.

The keyword **fullwidth** is relevant only to boxplots and to histograms with errorbars. It sets the width of the errorbar ends to be the same as the width of the associated box. It does not change the width of the box itself.

The **front** and **back** keywords are relevant only to errorbars attached to filled rectangles (boxes, candlesticks, histograms).

Bind

Show the current state of all hotkey bindings. See **bind** (p. 36).

Bmargin

The command **set bmargin** sets the size of the bottom margin. Please see **set margin** (p. 131) for details.

Border

The **set border** and **unset border** commands control the display of the graph borders for the **plot** and **splot** commands. Note that the borders do not necessarily coincide with the axes; with **plot** they often do, but with **splot** they usually do not.

Syntax:

```
set border {<integer>}
        {front | back | behind} {linewidth | lw <line_width>}
        {{linestyle | ls <line_style>} | {linetype | lt <line_type>}}
unset border
show border
```

With a **splot** displayed in an arbitrary orientation, like **set view 56,103**, the four corners of the x-y plane can be referred to as "front", "back", "left" and "right". A similar set of four corners exist for the top surface, of course. Thus the border connecting, say, the back and right corners of the x-y plane is the "bottom right back" border, and the border connecting the top and bottom front corners is the "front vertical". (This nomenclature is defined solely to allow the reader to figure out the table that follows.)

The borders are encoded in a 12-bit integer: the four low bits control the border for **plot** and the sides of the base for **splot**; the next four bits control the verticals in **splot**; the four high bits control the edges on top of an **splot**. The border settings is thus the sum of the appropriate entries from the following table:

Graph Border Encoding		
Bit	plot	splot
1	bottom	bottom left front
2	left	bottom left back
4	top	bottom right front
8	right	bottom right back
16	no effect	left vertical
32	no effect	back vertical
64	no effect	right vertical
128	no effect	front vertical
256	no effect	top left back
512	no effect	top right back
1024	no effect	top left front
2048	no effect	top right front

The default setting is 31, which is all four sides for **plot**, and base and z axis for **splot**.

In 2D plots the border is normally drawn on top of all plots elements (**front**). If you want the border to be drawn behind the plot elements, use **set border back**.

In hidden3d plots the lines making up the border are normally subject to the same hidden3d processing as the plot elements. **set border behind** will override this default.

Using the optional <line_style>, <line_type> and <line_width> specifiers, the way the border lines are drawn can be influenced (limited by what the current terminal driver supports).

For **plot**, tics may be drawn on edges other than bottom and left by enabling the second axes – see **set xtics** (p. 168) for details.

If a **splot** draws only on the base, as is the case with "**unset surface; set contour base**", then the verticals and the top are not drawn even if they are specified.

The **set grid** options 'back', 'front' and 'layerdefault' also control the order in which the border lines are drawn with respect to the output of the plotted data.

Examples:

Draw default borders:

```
set border
```

Draw only the left and bottom (**plot**) or both front and back bottom left (**splot**) borders:

```
set border 3
```

Draw a complete box around a **splot**:

```
set border 4095
```

Draw a topless box around a **splot**, omitting the front vertical:

```
set border 127+256+512 # or set border 1023-128
```

Draw only the top and right borders for a **plot** and label them as axes:

```
unset xtics; unset ytics; set x2tics; set y2tics; set border 12
```

Boxwidth

The **set boxwidth** command is used to set the default width of boxes in the **boxes**, **boxerrorbars**, **boxplot**, **candlesticks** and **histograms** styles.

Syntax:

```
set boxwidth {<width>} {absolute|relative}
show boxwidth
```

By default, adjacent boxes are extended in width until they touch each other. A different default width may be specified using the **set boxwidth** command. **Relative** widths are interpreted as being a fraction of this default width.

An explicit value for the boxwidth is interpreted as being a number of units along the current x axis (**absolute**) unless the modifier **relative** is given. If the x axis is a log-scale (see **set log** (p. 129)) then the value of boxwidth is truly "absolute" only at $x=1$; this physical width is maintained everywhere along the axis (i.e. the boxes do not become narrower the value of x increases). If the range spanned by a log scale x axis is far from $x=1$, some experimentation may be required to find a useful value of boxwidth.

The default is superseded by explicit width information taken from an extra data column in styles **boxes** or **boxerrorbars**. In a four-column data set, the fourth column will be interpreted as the box width unless the width is set to -2.0, in which case the width will be calculated automatically. See **style boxes** (p. 44) and **style boxerrorbars** (p. 44) for more details.

To set the box width to automatic use the command

```
set boxwidth
```

or, for four-column data,

```
set boxwidth -2
```

The same effect can be achieved with the **using** keyword in **plot**:

```
plot 'file' using 1:2:3:4:(-2)
```

To set the box width to half of the automatic size use

```
set boxwidth 0.5 relative
```

To set the box width to an absolute value of 2 use

```
set boxwidth 2 absolute
```

Color

Gnuplot supports two alternative sets of linetypes. The default set uses a different color for each linetype, although it also allows you to draw dotted or dashed lines in that color. The alternative monochrome set uses only dot/dash pattern or linewidth to distinguish linetypes. The **set color** command selects the color linetypes. See **set monochrome** (p. 131), **set linetype** (p. 128), and **set colorsequence** (p. 104).

Colorsequence

Syntax:

```
set colorsequence {default|classic|podo}
```

set colorsequence default selects a terminal-independent repeating sequence of eight colors. See **set linetype** (p. 128), **colors** (p. 33).

set colorsequence classic lets each separate terminal type provide its own sequence of line colors. The number provided varies from 4 to more than 100, but most start with red/green/blue/magenta/cyan/yellow. This was the default behaviour of earlier gnuplot versions.

set colorsequence podo selects eight colors drawn from a set recommended by Wong (2011) [Nature Methods 8:441] as being easily distinguished by color-blind viewers with either protanopia or deuteranopia.

In each case you can further customize the length of the sequence and the colors used. See **set linetype** (p. 128), **colors** (p. 33).

Clabel

This command is obsolete. Use **set cntrlabel** instead. **unset clabel** is replaced by **set cntrlabel onecolor**. **set clabel "format"** is replaced by **set cntrlabel format "format"**.

Clip

gnuplot can clip data points and lines that are near the boundaries of a graph.

Syntax:

```
set clip <clip-type>
unset clip <clip-type>
show clip
```

Three clip types for points and lines are supported by **gnuplot**: **points**, **one**, and **two**. One, two, or all three clip types may be active for a single graph. Note that clipping of color filled quadrangles drawn by **pm3d** maps and surfaces is not controlled by this command, but by **set pm3d cliplin** and **set pm3d clip4in**.

The **points** clip type forces **gnuplot** to clip (actually, not plot at all) data points that fall within but too close to the boundaries. This is done so that large symbols used for points will not extend outside the boundary lines. Without clipping points near the boundaries, the plot may look bad. Adjusting the x and y ranges may give similar results.

Setting the **one** clip type causes **gnuplot** to draw a line segment which has only one of its two endpoints within the graph. Only the in-range portion of the line is drawn. The alternative is to not draw any portion of the line segment.

Some lines may have both endpoints out of range, but pass through the graph. Setting the **two** clip-type allows the visible portion of these lines to be drawn.

In no case is a line drawn outside the graph.

The defaults are **noclip points**, **clip one**, and **noclip two**.

To check the state of all forms of clipping, use

```
show clip
```

For backward compatibility with older versions, the following forms are also permitted:

```
set clip
unset clip
```

set clip is synonymous with **set clip points**; **unset clip** turns off all three types of clipping.

Cntrlabel

Syntax:

```
set cntrlabel {format "format"} {font "font"}
set cntrlabel {start <int>} {interval <int>}
set cntrlabel onecolor
```

set cntrlabel controls the labeling of contours, either in the key (default) or on the plot itself in the case of **plot ... with labels**. In the latter case labels are placed along each contour line according to the "pointinterval" property of the label descriptor. By default a label is placed on the 5th line segment making up the contour line and repeated every 20th segment. These defaults are equivalent to

```
set cntrlabel start 5 interval 20
```

They can be changed either via the **set cntrlabel** command or by specifying the interval in the **splot** command itself

```
set contours; splot $F00 with labels point pointinterval -1
```

Setting the interval to a negative value means that the label appear only once per contour line. However if **set samples** or **set isosamples** is large then many contour lines may be created, each with a single label.

A contour label is placed in the plot key for each linetype used. By default each contour level is given its own linetype, so a separate label appears for each. The command **set cntrlabel onecolor** causes all contours to be drawn using the same linetype, so only one label appears in the plot key. This command replaces an older command **unset clabel**.

Cntrparam

set cntrparam controls the generation of contours and their smoothness for a contour plot. **show contour** displays current settings of **cntrparam** as well as **contour**.

Syntax:

```
set cntrparam { { linear
                | cubicspline
                | bspline
                | points <n>
                | order <n>
                | levels { auto {<n>} | <n>
                        | discrete <z1> {,<z2>{,<z3>...}}
                        | incremental <start>, <incr> {,<end>}
                        }
                }
            }
show contour
```

This command has two functions. First, it sets the values of z for which contour points are to be determined (by linear interpolation between data points or function isosamples.) Second, it controls the way contours are drawn between the points determined to be of equal z . $\langle n \rangle$ should be an integral constant expression and $\langle z1 \rangle$, $\langle z2 \rangle$... any constant expressions. The parameters are:

linear, **cubicspline**, **bspline** — Controls type of approximation or interpolation. If **linear**, then straight line segments connect points of equal z magnitude. If **cubicspline**, then piecewise-linear contours are interpolated between the same equal z points to form somewhat smoother contours, but which may undulate. If **bspline**, a guaranteed-smoother curve is drawn, which only approximates the position of the points of equal- z .

points — Eventually all drawings are done with piecewise-linear strokes. This number controls the number of line segments used to approximate the **bspline** or **cubicspline** curve. Number of cubicspline or bspline segments (strokes) = **points** * number of linear segments.

order — Order of the bspline approximation to be used. The bigger this order is, the smoother the resulting contour. (Of course, higher order bspline curves will move further away from the original piecewise linear data.) This option is relevant for **bspline** mode only. Allowed values are integers in the range from 2 (linear) to 10.

levels — Selection of contour levels, controlled by **auto** (default), **discrete**, **incremental**, and $\langle n \rangle$, number of contour levels.

For **auto**, $\langle n \rangle$ specifies a nominal number of levels; the actual number will be adjusted to give simple labels. If the surface is bounded by z_{\min} and z_{\max} , contours will be generated at integer multiples of dz between z_{\min} and z_{\max} , where dz is 1, 2, or 5 times some power of ten (like the step between two tic marks).

For **levels discrete**, contours will be generated at $z = \langle z1 \rangle$, $\langle z2 \rangle$... as specified; the number of discrete levels sets the number of contour levels. In **discrete** mode, any **set cntrparam levels** $\langle n \rangle$ are ignored.

For **incremental**, contours are generated at values of z beginning at $\langle \text{start} \rangle$ and increasing by $\langle \text{increment} \rangle$, until the number of contours is reached. $\langle \text{end} \rangle$ is used to determine the number of contour levels, which will be changed by any subsequent **set cntrparam levels** $\langle n \rangle$. If the z axis is logarithmic, $\langle \text{increment} \rangle$ will be interpreted as a factor, just like in **set ztics**.

If the command **set cntrparam** is given without any arguments specified, the defaults are used: linear, 5

points, order 4, 5 auto levels.

Examples:

```
set cntrparam bspline
set cntrparam points 7
set cntrparam order 10
```

To select levels automatically, 5 if the level increment criteria are met:

```
set cntrparam levels auto 5
```

To specify discrete levels at .1, .37, and .9:

```
set cntrparam levels discrete .1,1/exp(1),.9
```

To specify levels from 0 to 4 with increment 1:

```
set cntrparam levels incremental 0,1,4
```

To set the number of levels to 10 (changing an incremental end or possibly the number of auto levels):

```
set cntrparam levels 10
```

To set the start and increment while retaining the number of levels:

```
set cntrparam levels incremental 100,50
```

See also **set contour** (p. 108) for control of where the contours are drawn, and **set cntrlabel** (p. 105) for control of the format of the contour labels and linetypes.

See also

```
contours demo (contours.dem)
```

and

```
contours with user defined levels demo (discrete.dem).
```

Color box

The color scheme, i.e. the gradient of the smooth color with `min_z` and `max_z` values of **pm3d**'s **palette**, is drawn in a color box unless **unset colorbox**.

```
set colorbox
set colorbox {
    { vertical | horizontal }
    { default | user }
    { origin x, y }
    { size x, y }
    { front | back }
    { noborder | bdefault | border [line style] }
}
show colorbox
unset colorbox
```

Color box position can be **default** or **user**. If the latter is specified the values as given with the **origin** and **size** subcommands are used. The box can be drawn after (**front**) or before (**back**) the graph or the surface.

The orientation of the color gradient can be switched by options **vertical** and **horizontal**.

origin x, y and **size x, y** are used only in combination with the **user** option. The x and y values are interpreted as screen coordinates by default, and this is the only legal option for 3D plots. 2D plots, including splot with **set view map**, allow any coordinate system to be specified. Try for example:

```
set colorbox horiz user origin .1,.02 size .8,.04
```

which will draw a horizontal gradient somewhere at the bottom of the graph.

border turns the border on (this is the default). **noborder** turns the border off. If an positive integer argument is given after **border**, it is used as a line style tag which is used for drawing the border, e.g.:

```
set style line 2604 linetype -1 linewidth .4
set colorbox border 2604
```

will use line style **2604**, a thin line with the default border color (-1) for drawing the border. **bdefault** (which is the default) will use the default border line style for drawing the border of the color box.

The axis of the color box is called **cb** and it is controlled by means of the usual axes commands, i.e. **set/unset/show** with **cbrange**, **[m]cbtics**, **format cb**, **grid [m]cb**, **cblabel**, and perhaps even **cbdata**, **[no]cbdtics**, **[no]cbmtics**.

set colorbox without any parameter switches the position to default. **unset colorbox** resets the default parameters for the colorbox and switches the colorbox off.

See also help for **set pm3d** (p. 141), **set palette** (p. 144), **x11 pm3d** (p. ??), and **set style line** (p. 156).

Colornames

Gnuplot knows a limited number of color names. You can use these to define the color range spanned by a pm3d palette, or to assign a terminal-independent color to a particular linetype or linestyle. To see the list of known color names, use the command **show colornames** (p. 108). Example:

```
set style line 1 linecolor "sea-green"
```

Contour

set contour enables contour drawing for surfaces. This option is available for **splot** only. It requires grid data, see **grid_data** (p. 179) for more details. If contours are desired from non-grid data, **set dgrid3d** can be used to create an appropriate grid.

Syntax:

```
set contour {base | surface | both}
unset contour
show contour
```

The three options specify where to draw the contours: **base** draws the contours on the grid base where the x/yticks are placed, **surface** draws the contours on the surfaces themselves, and **both** draws the contours on both the base and the surface. If no option is provided, the default is **base**.

See also **set cntrparam** (p. 106) for the parameters that affect the drawing of contours, and **set cntrlabel** (p. 105) for control of labeling of the contours.

The surface can be switched off (see **unset surface** (p. 158)), giving a contour-only graph. Though it is possible to use **set size** to enlarge the plot to fill the screen, more control over the output format can be obtained by writing the contour information to a datablock, and rereading it as a 2D datafile plot:

```
unset surface
set contour
set cntrparam ...
set table $datablock
splot ...
unset table
# contour info now in $datablock
set term <whatever>
plot $datablock
```

In order to draw contours, the data should be organized as "grid data". In such a file all the points for a single y-isoline are listed, then all the points for the next y-isoline, and so on. A single blank line (a line containing no characters other than blank spaces and a carriage return and/or a line feed) separates one y-isoline from the next. See also **splot datafile** (p. 176).

See also

```
contours demo (contours.dem)
```

and

```
contours with user defined levels demo (discrete.dem).
```

Dashtype

The **set dashtype** command allows you to define a dash pattern that can then be referred to by its index. This is purely a convenience, as anywhere that would accept the dashtype by its numerical index would also accept an explicit dash pattern. Example:

```
set dashtype 5 (2,4,2,6)  # define or redefine dashtype number 5
plot f1(x) dt 5           # plot using the new dashtype
plot f1(x) dt (2,4,2,6)  # exactly the same plot as above
set linetype 5 dt 5       # always use this dash pattern with linetype 5
set dashtype 66 "..-"    # define a new dashtype using a string
```

See also **dashtype** (p. 35).

Data style

This form of the command is deprecated. Please see **set style data** (p. 154).

Datafile

The **set datafile** command options control interpretation of fields read from input data files by the **plot**, **splot**, and **fit** commands. Six such options are currently implemented.

Set datafile fortran

The **set datafile fortran** command enables a special check for values in the input file expressed as Fortran D or Q constants. This extra check slows down the input process, and should only be selected if you do in fact have datafiles containing Fortran D or Q constants. The option can be disabled again using **unset datafile fortran**.

Set datafile nofpe_trap

The **set datafile nofpe_trap** command tells gnuplot not to re-initialize a floating point exception handler before every expression evaluation used while reading data from an input file. This can significantly speed data input from very large files at the risk of program termination if a floating-point exception is generated.

Set datafile missing

The **set datafile missing** command tells **gnuplot** there is a special string used in input data files to denote a missing data entry. There is no default character for **missing**, but in many cases any non-parseable string of characters found where a numerical value is expected will cause the line to be treated as missing data. There is a distinction between missing data and invalid data (e.g. "NaN", 1/0.). Invalid data causes a gap in the line drawn through the points; missing data does not.

Syntax:

```
set datafile missing "<string>"
show datafile missing
unset datafile
```

Note: The treatment of certain cases has changed in this version of gnuplot. The example below shows the difference. Example:

```

set style data linespoints
plot '-' title "(a)"
  1 10
  2 20
  3 ?
  4 40
  5 50
  e
set datafile missing "?"
plot '-' title "(b)"
  1 10
  2 20
  3 ?
  4 40
  5 50
  e
plot '-' using 1:2 title "(c)"
  1 10
  2 20
  3 NaN
  4 40
  5 50
  e
plot '-' using 1:($2) title "(d)"
  1 10
  2 20
  3 NaN
  4 40
  5 50
  e

```

Plot (a) differs because the third line contains only one valid number. Old gnuplot versions switched to a single-datum-on-a-line convention that the line number is "x" and the datum is "y", erroneously placing the point at(2,3).

Both the old and new gnuplot versions handle the same data correctly if the '?' character is designated as a marker for missing data (b).

Old gnuplot versions handled NaN differently depending of the form of the **using** clause, as shown in plots (c) and (d). Gnuplot now handles NaN the same whether the input column was specified as N or (\$N). See also the

[imageNaN demo](#).

Set datafile separator

The command **set datafile separator** tells **gnuplot** that data fields in subsequent input files are separated by a specific character rather than by whitespace. The most common use is to read in csv (comma-separated value) files written by spreadsheet or database programs. By default data fields are separated by whitespace.

Syntax:

```
set datafile separator {whitespace | tab | comma | "<chars>"}
```

Examples:

```

# Input file contains tab-separated fields
set datafile separator "\t"

# Input file contains comma-separated values fields
set datafile separator comma

```

```
# Input file contains fields separated by either * or |
set datafile separator "*" |
```

Set datafile commentschars

The **set datafile commentschars** tells **gnuplot** what characters are used in a data file to begin comment lines. If the first non-blank character on a line is one of the specified characters then the rest of the input line is ignored. Default value of the string is "#!" on VMS and "#" otherwise.

Syntax:

```
set datafile commentschars {"<string>"}
show datafile commentschars
unset commentschars
```

Then, the following line in a data file is completely ignored

```
# 1 2 3 4
```

but the following

```
1 # 3 4
```

produces a rather unexpected plot unless

```
set datafile missing '#'
```

is specified as well.

Example:

```
set datafile commentschars "#!%"
```

Set datafile binary

The **set datafile binary** command is used to set the defaults when reading binary data files. The syntax matches precisely that used for commands **plot** and **splot**. See **binary matrix** (p. 177) and **binary general** (p. 75) for details about the keywords that can be present in <binary list>.

Syntax:

```
set datafile binary <binary list>
show datafile binary
show datafile
unset datafile
```

Examples:

```
set datafile binary filetype=auto
set datafile binary array=(512,512) format="%uchar"

show datafile binary    # list current settings
```

Decimalsign

The **set decimalsign** command selects a decimal sign for numbers printed into tic labels or **set label** strings.

Syntax:

```
set decimalsign {<value> | locale {"<locale>"}}
unset decimalsign
show decimalsign
```

The argument `<value>` is a string to be used in place of the usual decimal point. Typical choices include the period, `'.'`, and the comma, `','`, but others may be useful, too. If you omit the `<value>` argument, the decimal separator is not modified from the usual default, which is a period. Unsetting `decimalsign` has the same effect as omitting `<value>`.

Example:

Correct typesetting in most European countries requires:

```
set decimalsign ','
```

Please note: If you set an explicit string, this affects only numbers that are printed using gnuplot's `gprintf()` formatting routine, include axis tics. It does not affect the format expected for input data, and it does not affect numbers printed with the `sprintf()` formatting routine. To change the behavior of both input and output formatting, instead use the form

```
set decimalsign locale
```

This instructs the program to use both input and output formats in accordance with the current setting of the `LC_ALL`, `LC_NUMERIC`, or `LANG` environmental variables.

```
set decimalsign locale "foo"
```

This instructs the program to format all input and output in accordance with locale `"foo"`, which must be installed. If locale `"foo"` is not found then an error message is printed and the decimal sign setting is unchanged. On linux systems you can get a list of the locales installed on your machine by typing `"locale -a"`. A typical linux locale string is of the form `"sl.SI.UTF-8"`. A typical Windows locale string is of the form `"Slovenian.Slovenia.1250"` or `"slovenian"`. Please note that interpretation of the locale settings is done by the C library at runtime. Older C libraries may offer only partial support for locale settings such as the thousands grouping separator character.

```
set decimalsign locale; set decimalsign "."
```

This sets all input and output to use whatever decimal sign is correct for the current locale, but over-rides this with an explicit `'.'` in numbers formatted using gnuplot's internal `gprintf()` function.

Dgrid3d

The **set dgrid3d** command enables, and can set parameters for, non-grid to grid data mapping. See **plot grid_data** (p. 179) for more details about the grid data structure.

Syntax:

```
set dgrid3d {<rows>} {,{<cols>}}
    { splines |
      qnorm {<norm>} |
      (gauss | cauchy | exp | box | hann)
      {kdensity} {<dx>} {,<dy>} }
unset dgrid3d
show dgrid3d
```

By default **dgrid3d** is disabled. When enabled, 3D data read from a file are always treated as a scattered data set. A grid with dimensions derived from a bounding box of the scattered data and size as specified by the `row/col_size` parameters is created for plotting and contouring. The grid is equally spaced in *x* (rows) and in *y* (columns); the *z* values are computed as weighted averages or spline interpolations of the scattered points' *z* values. In other words, a regularly spaced grid is created and the a smooth approximation to the raw data is evaluated for all grid points. This approximation is plotted in place of the raw data.

The number of columns defaults to the number of rows, which defaults to 10.

Several algorithms are available to calculate the approximation from the raw data. Some of these algorithms can take additional parameters. These interpolations are such the closer the data point is to a grid point, the more effect it has on that grid point.

The **splines** algorithm calculates an interpolation based on "thin plate splines". It does not take additional parameters.

The **qnorm** algorithm calculates a weighted average of the input data at each grid point. Each data point is weighted inversely by its distance from the grid point raised to the norm power. (Actually, the weights are given by the inverse of $dx^{\text{norm}} + dy^{\text{norm}}$, where dx and dy are the components of the separation of the grid point from each data point. For some norms that are powers of two, specifically 4, 8, and 16, the computation is optimized by using the Euclidean distance in the weight calculation, $(dx^2 + dy^2)^{\text{norm}/2}$. However, any non-negative integer can be used.) The power of the norm can be specified as a single optional parameter. This algorithm is the default.

Finally, several smoothing kernels are available to calculate weighted averages: $z = \text{Sum}_i w(d_i) * z_i / \text{Sum}_i w(d_i)$, where z_i is the value of the i -th data point and d_i is the distance between the current grid point and the location of the i -th data point. All kernels assign higher weights to data points that are close to the current grid point and lower weights to data points further away.

The following kernels are available:

```
gauss :    w(d) = exp(-d*d)
cauchy :    w(d) = 1/(1 + d*d)
exp :      w(d) = exp(-d)
box :      w(d) = 1                if d<1
           = 0                    otherwise
hann :      w(d) = 0.5*(1-cos(2*pi*d)) if d<1
           w(d) = 0                    otherwise
```

When using one of these five smoothing kernels, up to two additional numerical parameters can be specified: dx and dy . These are used to rescale the coordinate differences when calculating the distance: $d_i = \sqrt{((x-x_i)/dx)^2 + ((y-y_i)/dy)^2}$, where x,y are the coordinates of the current grid point and x_i,y_i are the coordinates of the i -th data point. The value of dy defaults to the value of dx , which defaults to 1. The parameters dx and dy make it possible to control the radius over which data points contribute to a grid point IN THE UNITS OF THE DATA ITSELF.

The optional keyword **kdensity**, which must come after the name of the kernel, but before the (optional) scale parameters, modifies the algorithm so that the values calculated for the grid points are not divided by the sum of the weights ($z = \text{Sum}_i w(d_i) * z_i$). If all z_i are constant, this effectively plots a bivariate kernel density estimate: a kernel function (one of the five defined above) is placed at each data point, the sum of these kernels is evaluated at every grid point, and this smooth surface is plotted instead of the original data. This is similar in principle to + what the **smooth kdensity** option does to 1D datasets. (See `kdensity2d.dem` for usage demo)

A slightly different syntax is also supported for reasons of backwards compatibility. If no interpolation algorithm has been explicitly selected, the **qnorm** algorithm is assumed. Up to three comma-separated, optional parameters can be specified, which are interpreted as the the number of rows, the number of columns, and the norm value, respectively.

The **dgrid3d** option is a simple scheme which replaces scattered data with weighted averages on a regular grid. More sophisticated approaches to this problem exist and should be used to preprocess the data outside **gnuplot** if this simple solution is found inadequate.

See also

`dgrid3d.dem: dgrid3d demo.`

and

`scatter.dem: dgrid3d demo.`

Dummy

The **set dummy** command changes the default dummy variable names.

Syntax:

```
set dummy {<dummy-var>} {,<dummy-var>}
show dummy
```

By default, **gnuplot** assumes that the independent, or "dummy", variable for the **plot** command is "t" if in parametric or polar mode, or "x" otherwise. Similarly the independent variables for the **splot** command are "u" and "v" in parametric mode (**splot** cannot be used in polar mode), or "x" and "y" otherwise.

It may be more convenient to call a dummy variable by a more physically meaningful or conventional name. For example, when plotting time functions:

```
set dummy t
plot sin(t), cos(t)
```

Examples:

```
set dummy u,v
set dummy ,s
```

The second example sets the second variable to s. To reset the dummy variable names to their default values, use

```
unset dummy
```

Encoding

The **set encoding** command selects a character encoding.

Syntax:

```
set encoding {<value>}
set encoding locale
show encoding
```

Valid values are

default	- tells a terminal to use its default encoding
iso_8859_1	- the most common Western European encoding used by many Unix workstations and by MS-Windows. This encoding is known in the PostScript world as 'ISO-Latin1'.
iso_8859_15	- a variant of iso_8859_1 that includes the Euro symbol
iso_8859_2	- used in Central and Eastern Europe
iso_8859_9	- used in Turkey (also known as Latin5)
koi8r	- popular Unix cyrillic encoding
koi8u	- Ukrainian Unix cyrillic encoding
cp437	- codepage for MS-DOS
cp850	- codepage for OS/2, Western Europe
cp852	- codepage for OS/2, Central and Eastern Europe
cp950	- MS version of Big5 (emf terminal only)
cp1250	- codepage for MS Windows, Central and Eastern Europe
cp1251	- codepage for 8-bit Russian, Serbian, Bulgarian, Macedonian
cp1252	- codepage for MS Windows, Western Europe
cp1254	- codepage for MS Windows, Turkish (superset of Latin5)
sjis	- shift-JIS Japanese encoding
utf8	- variable-length (multibyte) representation of Unicode entry point for each character

The command **set encoding locale** is different from the other options. It attempts to determine the current locale from the runtime environment. On most systems this is controlled by the environmental variables LC_ALL, LC_CTYPE, or LANG. This mechanism is necessary, for example, to pass multibyte character encodings such as UTF-8 or EUC_JP to the wxt and cairopdf terminals. This command does not affect the locale-specific representation of dates or numbers. See also **set locale** (p. 129) and **set decimalsign** (p. 111).

Generally you must set the encoding before setting the terminal type. Note that encoding is not supported by all terminal drivers and that the device must be able to produce the desired non-standard characters.

Fit

The **set fit** command controls the options for the **fit** command.

Syntax:

```
set fit {nolog | logfile {"<filename>"|default}}
      {{no}quiet|results|brief|verbose}
      {{no}errorvariables}
      {{no}covariancevariables}
      {{no}errorscaling}
      {{no}prescale}
      {maxiter <value>|default}
      {limit <epsilon>|default}
      {limit_abs <epsilon_abs>}
      {start-lambda <value>|default}
      {lambda-factor <value>|default}
      {script {"<command>"|default}}
      {v4 | v5}

unset fit
show fit
```

The **logfile** option defines where the **fit** command writes its output. The **<filename>** argument must be enclosed in single or double quotes. If no filename is given or **unset fit** is used the log file is reset to its default value "fit.log" or the value of the environmental variable **FIT_LOG**. If the given logfile name ends with a / or \, it is interpreted to be a directory name, and the actual filename will be "fit.log" in that directory.

By default the information written to the log file is also echoed to the terminal session. **set fit quiet** turns off the echo, whereas **results** prints only final results. **brief** gives one line summaries for every iteration of the fit in addition. **verbose** yields detailed iteration reports as it was the default before version 5.

If the **errorvariables** option is turned on, the error of each fitted parameter computed by **fit** will be copied to a user-defined variable whose name is formed by appending "_err" to the name of the parameter itself. This is useful mainly to put the parameter and its error onto a plot of the data and the fitted function, for reference, as in:

```
set fit errorvariables
fit f(x) 'datafile' using 1:2 via a, b
print "error of a is:", a_err
set label 1 sprintf("a=%6.2f +/- %6.2f", a, a_err)
plot 'datafile' using 1:2, f(x)
```

If the **errorscaling** option is specified, which is the default, the calculated parameter errors are scaled with the reduced chi square. This is equivalent to providing data errors equal to the calculated standard deviation of the fit (FIT_STDFIT) resulting in a reduced chi square of one. With the **noerrorscaling** option the estimated errors are the unscaled standard deviations of the fit parameters. If no weights are specified for the data, parameter errors are always scaled.

If the **prescale** option is turned on, parameters are prescaled by their initial values before being passed to the Marquardt-Levenberg routine. This helps tremendously if there are parameters that differ in size by many orders of magnitude. Fit parameters with an initial value of exactly zero are never prescaled.

The maximum number of iterations may be limited with the **maxiter** option. A value of 0 or **default** means that there is no limit.

The **limit** option can be used to change the default epsilon limit (1e-5) to detect convergence. When the sum of squared residuals changes by a factor less than this number (epsilon), the fit is considered to have 'converged'. The **limit_abs** option imposes an additional absolute limit in the change of the sum of squared residuals and defaults to zero.

If you need even more control about the algorithm, and know the Marquardt-Levenberg algorithm well, the following options can be used to influence it. The startup value of **lambda** is normally calculated automatically from the ML-matrix, but if you want to, you may provide your own using the **start_lambda**

option. Setting it to **default** will re-enable the automatic selection. The option **lambda_factor** sets the factor by which **lambda** is increased or decreased whenever the chi-squared target function increased or decreased significantly. Setting it to **default** re-enables the default factor of 10.0.

The **script** option may be used to specify a **gnuplot** command to be executed when a fit is interrupted — see **fit** (p. 62). This setting takes precedence over the default of **replot** and the environment variable **FIT_SCRIPT**.

If the **covariancevariables** option is turned on, the covariances between final parameters will be saved to user-defined variables. The variable name for a certain parameter combination is formed by prepending "FIT_COV_" to the name of the first parameter and combining the two parameter names by "_". For example given the parameters "a" and "b" the covariance variable is named "FIT_COV_a.b".

In version 5 the syntax of the fit command changed and it now defaults to unitweights if no 'error' keyword is given. The **v4** option restores the default behavior of gnuplot version 4, see also **fit** (p. 62).

Fontpath

The **fontpath** setting defines additional locations for font files searched when including font files. Currently only the postscript terminal supports **fontpath**. If a file cannot be found in the current directory, the directories in **fontpath** are tried. Further documentation concerning the supported file formats is included in the **terminal postscript** section of the documentation.

Syntax:

```
set fontpath {"pathlist1" {"pathlist2"...}}
show fontpath
```

Path names may be entered as single directory names, or as a list of path names separated by a platform-specific path separator, eg. colon (':') on Unix, semicolon (';') on DOS/Windows/OS/2 platforms. The **show fontpath**, **save** and **save set** commands replace the platform-specific separator with a space character (' ') for maximum portability. If a directory name ends with an exclamation mark ('!') also the subdirectories of this directory are searched for font files.

If the environmental variable **GNUPLOT_FONTPATH** is set, its contents are appended to **fontpath**. If it is not set, a system dependent default value is used. It is set by testing several directories for existence when using the fontpath the first time. Thus, the first call of **set fontpath**, **show fontpath**, **save fontpath**, **plot**, or **splot** with embedded font files takes a little more time. If you want to save this time you may set the environmental variable **GNUPLOT_FONTPATH** since probing is switched off, then. You can find out which is the default fontpath by using **show fontpath**.

show fontpath prints the contents of the user-defined fontpath and the system fontpath separately. However, the **save** and **save set** commands save only the user-specified parts of **fontpath**.

For terminal drivers that access fonts by filename via the gd library, the font search path is controlled by the environmental variable **GDFONTPATH**.

Format

The format of the tic-mark labels can be set with the **set format** command or with the **set tics format** or individual **set {axis}tics format** commands.

Syntax:

```
set format {<axes>} {"<format-string>"} {numeric|timedate|geographic}
show format
```

where <axes> is either **x**, **y**, **xy**, **x2**, **y2**, **z**, **cb** or nothing (which applies the format to all axes). The following two commands are equivalent:

```
set format y "%.2f"
set ytics format "%.2f"
```

The length of the string is restricted to 100 characters. The default format is "% h", "\$%h\$" for LaTeX terminals. Other formats such as "%.2f" or "%3.0em" are often desirable. "set format" with no following string will restore the default.

If the empty string "" is given, tics will have no labels, although the tic mark will still be plotted. To eliminate the tic marks, use **unset xtics** or **set tics scale 0**.

Newline (\n) and enhanced text markup is accepted in the format string. Use double-quotes rather than single-quotes in this case. See also **syntax** (p. 41). Characters not preceded by "%" are printed verbatim. Thus you can include spaces and labels in your format string, such as "%g m", which will put " m" after each number. If you want "%" itself, double it: "%g %%".

See also **set xtics** (p. 168) for more information about tic labels, and **set decimalsign** (p. 111) for how to use non-default decimal separators in numbers printed this way. See also

[electron demo \(electron.dem\)](#).

Gprintf

The string function `gprintf("format",x)` uses gnuplot's own format specifiers, as do the gnuplot commands **set format**, **set timestamp**, and others. These format specifiers are not the same as those used by the standard C-language routine `sprintf()`. `gprintf()` accepts only a single variable to be formatted. Gnuplot also provides an `sprintf("format",x1,x2,...)` routine if you prefer. For a list of gnuplot's format options, see **format specifiers** (p. 117).

Format specifiers

The acceptable formats (if not in time/date mode) are:

Tic-mark label numerical format specifiers	
Format	Explanation
%f	floating point notation
%e or %E	exponential notation; an "e" or "E" before the power
%g or %G	the shorter of %e (or %E) and %f
%h or %H	like %g with "x10^{%S}" or "*10^{%S}" instead of "e%S"
%x or %X	hex
%o or %O	octal
%t	mantissa to base 10
%l	mantissa to base of current logscale
%s	mantissa to base of current logscale; scientific power
%T	power to base 10
%L	power to base of current logscale
%S	scientific power
%c	character replacement for scientific power
%b	mantissa of ISO/IEC 80000 notation (ki, Mi, Gi, Ti, Pi, Ei, Zi, Yi)
%B	prefix of ISO/IEC 80000 notation (ki, Mi, Gi, Ti, Pi, Ei, Zi, Yi)
%P	multiple of pi

A 'scientific' power is one such that the exponent is a multiple of three. Character replacement of scientific powers ("%c") has been implemented for powers in the range -18 to +18. For numbers outside of this range the format reverts to exponential.

Other acceptable modifiers (which come after the "%" but before the format specifier) are "-", which left-justifies the number; "+", which forces all numbers to be explicitly signed; " " (a space), which makes positive numbers have a space in front of them where negative numbers have "-"; "#", which places a decimal point after floats that have only zeroes following the decimal point; a positive integer, which defines the field width; "0" (the digit, not the letter) immediately preceding the field width, which indicates that leading zeroes are to be used instead of leading blanks; and a decimal point followed by a non-negative integer, which defines the precision (the minimum number of digits of an integer, or the number of digits following the decimal point of a float).

Some systems may not support all of these modifiers but may also support others; in case of doubt, check the appropriate documentation and then experiment.

Examples:

```
set format y "%t"; set ytics (5,10)          # "5.0" and "1.0"
set format y "%s"; set ytics (500,1000)      # "500" and "1.0"
set format y "%+-12.3f"; set ytics(12345)    # "+12345.000 "
set format y "%.2t*10^%+03T"; set ytic(12345)# "1.23*10^+04"
set format y "%s*10^{%S}"; set ytic(12345)   # "12.345*10^{3}"
set format y "%s %cg"; set ytic(12345)       # "12.345 kg"
set format y "%.0P pi"; set ytic(6.283185)   # "2 pi"
set format y "%.0f%%"; set ytic(50)          # "50%"

set log y 2; set format y '%l'; set ytics (1,2,3)
#displays "1.0", "1.0" and "1.5" (since 3 is 1.5 * 2^1)
```

There are some problem cases that arise when numbers like 9.999 are printed with a format that requires both rounding and a power.

If the data type for the axis is time/date, the format string must contain valid codes for the 'strftime' function (outside of **gnuplot**, type "man strftime"). See **set timefmt** (p. 161) for a list of the allowed input format codes.

Time/date specifiers

In time/date mode, the acceptable formats are:

Date Specifiers	
Format	Explanation
%a	abbreviated name of day of the week
%A	full name of day of the week
%b or %h	abbreviated name of the month
%B	full name of the month
%d	day of the month, 01–31
%D	shorthand for "%m/%d/%y" (only output)
%F	shorthand for "%Y-%m-%d" (only output)
%k	hour, 0–23 (one or two digits)
%H	hour, 00–23 (always two digits)
%l	hour, 1–12 (one or two digits)
%I	hour, 01–12 (always two digits)
%j	day of the year, 001–366
%m	month, 01–12
%M	minute, 00–60
%p	"am" or "pm"
%r	shorthand for "%I:%M:%S %p" (only output)
%R	shorthand for "%H:%M" (only output)
%S	second, integer 00–60 on output, (double) on input
%s	number of seconds since start of year 1970
%T	shorthand for "%H:%M:%S" (only output)
%U	week of the year (week starts on Sunday)
%w	day of the week, 0–6 (Sunday = 0)
%W	week of the year (week starts on Monday)
%y	year, 0–99 in range 1969–2068
%Y	year, 4-digit

Numerical formats may be preceded by a "0" ("zero") to pad the field with leading zeroes, and preceded by a positive digit to define the minimum field width. The %S, and %t formats also accept a precision specifier so that fractional hours/minutes/seconds can be written.

Time Specifiers	
Format	Explanation
%tH	+/- hours relative to time=0 (does not wrap at 24)
%tM	+/- minutes relative to time=0
%tS	+/- seconds associated with previous tH or tM field

Examples Examples of date format:

Suppose the x value in seconds corresponds a time slightly before midnight on 25 Dec 1976. The text printed for a tic label at this position would be

```
set format x          # defaults to "12/25/76 \n 23:11"
set format x "%A, %d %b %Y" # "Saturday, 25 Dec 1976"
set format x "%r %D"      # "11:11:11 pm 12/25/76"
```

Examples of time format:

The date format specifiers encode a time in seconds as a clock time on a particular day. So hours run only from 0-23, minutes from 0-59, and negative values correspond to dates prior to the epoch (1-Jan-1970). In order to report a time value in seconds as some number of hours/minutes/seconds relative to a time 0, use time formats %tH %tM %tS. To report a value of -3672.50 seconds

```
set format x          # default date format "12/31/69 \n 22:58"
set format x "%tH:%tM:%tS" # "-01:01:12"
set format x "%.2tH hours" # "-1.02 hours"
set format x "%tM:%.2tS"   # "-61:12.50"
```

Function style

This form of the command is deprecated. Please see **set style function** (p. 155).

Functions

The **show functions** command lists all user-defined functions and their definitions.

Syntax:

```
show functions
```

For information about the definition and usage of functions in **gnuplot**, please see **expressions** (p. 24). See also

[splines as user defined functions \(spline.dem\)](#)

and

[use of functions and complex variables for airfoils \(airfoil.dem\).](#)

Grid

The **set grid** command allows grid lines to be drawn on the plot.

Syntax:

```
set grid {{no}{m}xtics} {{no}{m}ytics} {{no}{m}ztics}
        {{no}{m}x2tics} {{no}{m}y2tics}
        {{no}{m}cbtics}
        {polar {<angle>}}
        {layerdefault | front | back}
        { {linestyle <major_linestyle>}
          | {linetype | lt <major_linetype>}
          | {linewidth | lw <major_linewidth>}}
```



```

        { , {linestyle | ls <minor_linestyle>}
          | {linetype | lt <minor_linetype>}
          {linewidth | lw <minor_linewidth>}} } }

unset grid
show grid

```

The grid can be enabled and disabled for the major and/or minor tic marks on any axis, and the linetype and linewidth can be specified for major and minor grid lines, also via a predefined linestyle, as far as the active terminal driver supports this.

Additionally, a polar grid can be selected for 2D plots — circles are drawn to intersect the selected tics, and radial lines are drawn at definable intervals. (The interval is given in degrees or radians, depending on the **set angles** setting.) Note that a polar grid is no longer automatically generated in polar mode.

The pertinent tics must be enabled before **set grid** can draw them; **gnuplot** will quietly ignore instructions to draw grid lines at non-existent tics, but they will appear if the tics are subsequently enabled.

If no linetype is specified for the minor gridlines, the same linetype as the major gridlines is used. The default polar angle is 30 degrees.

If **front** is given, the grid is drawn on top of the graphed data. If **back** is given, the grid is drawn underneath the graphed data. Using **front** will prevent the grid from being obscured by dense data. The default setup, **layerdefault**, is equivalent to **back** for 2D plots. In 3D plots the default is to split up the grid and the graph box into two layers: one behind, the other in front of the plotted data and functions. Since **hidden3d** mode does its own sorting, it ignores all grid drawing order options and passes the grid lines through the hidden line removal machinery instead. These options actually affect not only the grid, but also the lines output by **set border** and the various ticmarks (see **set xtics** (p. 168)).

Z grid lines are drawn on the bottom of the plot. This looks better if a partial box is drawn around the plot — see **set border** (p. 102).

Hidden3d

The **set hidden3d** command enables hidden line removal for surface plotting (see **splot** (p. 176)). Some optional features of the underlying algorithm can also be controlled using this command.

Syntax:

```

set hidden3d {defaults} |
    { {front|back}
      {{offset <offset>} | {nooffset}}
      {trianglepattern <bitpattern>}
      {{undefined <level>} | {noundefined}}
      {{no}altldiagonal}
      {{no}bentover} }
unset hidden3d
show hidden3d

```

In contrast to the usual display in gnuplot, hidden line removal actually treats the given function or data grids as real surfaces that can't be seen through, so plot elements behind the surface will be hidden by it. For this to work, the surface needs to have 'grid structure' (see **splot datafile** (p. 176) about this), and it has to be drawn **with lines** or **with linespoints**.

When **hidden3d** is set, both the hidden portion of the surface and possibly its contours drawn on the base (see **set contour** (p. 108)) as well as the grid will be hidden. Each surface has its hidden parts removed with respect to itself and to other surfaces, if more than one surface is plotted. Contours drawn on the surface (**set contour surface**) don't work.

As of gnuplot version 4.6, **hidden3d** also affects 3D plotting styles **points**, **labels**, **vectors**, and **impulses** even if no surface is present in the graph. Unobscured portions of each vector are drawn as line segments (no arrowheads). Individual plots within the graph may be explicitly excluded from this processing by appending the extra option **nohidden3d** to the **with** specifier.

Hidden3d does not affect solid surfaces drawn using the pm3d mode. To achieve a similar effect purely for pm3d surfaces, use instead **set pm3d depthorder**. To mix pm3d surfaces with normal **hidden3d** processing, use the option **set hidden3d front** to force all elements included in hidden3d processing to be drawn after any remaining plot elements, including the pm3d surface.

Functions are evaluated at isoline intersections. The algorithm interpolates linearly between function points or data points when determining the visible line segments. This means that the appearance of a function may be different when plotted with **hidden3d** than when plotted with **nohidden3d** because in the latter case functions are evaluated at each sample. Please see **set samples** (p. 151) and **set isosamples** (p. 122) for discussion of the difference.

The algorithm used to remove the hidden parts of the surfaces has some additional features controllable by this command. Specifying **defaults** will set them all to their default settings, as detailed below. If **defaults** is not given, only explicitly specified options will be influenced: all others will keep their previous values, so you can turn on/off hidden line removal via **set {no}hidden3d**, without modifying the set of options you chose.

The first option, **offset**, influences the linetype used for lines on the 'back' side. Normally, they are drawn in a linetype one index number higher than the one used for the front, to make the two sides of the surface distinguishable. You can specify a different linetype offset to add instead of the default 1, by **offset <offset>**. Option **nooffset** stands for **offset 0**, making the two sides of the surface use the same linetype.

Next comes the option **trianglepattern <bitpattern>**. <bitpattern> must be a number between 0 and 7, interpreted as a bit pattern. Each bit determines the visibility of one edge of the triangles each surface is split up into. Bit 0 is for the 'horizontal' edges of the grid, Bit 1 for the 'vertical' ones, and Bit 2 for the diagonals that split each cell of the original grid into two triangles. The default pattern is 3, making all horizontal and vertical lines visible, but not the diagonals. You may want to choose 7 to see those diagonals as well.

The **undefined <level>** option lets you decide what the algorithm is to do with data points that are undefined (missing data, or undefined function values), or exceed the given x-, y- or z-ranges. Such points can either be plotted nevertheless, or taken out of the input data set. All surface elements touching a point that is taken out will be taken out as well, thus creating a hole in the surface. If <level> = 3, equivalent to option **noundefined**, no points will be thrown away at all. This may produce all kinds of problems elsewhere, so you should avoid this. <level> = 2 will throw away undefined points, but keep the out-of-range ones. <level> = 1, the default, will get rid of out-of-range points as well.

By specifying **noaltdiagonal**, you can override the default handling of a special case can occur if **undefined** is active (i.e. <level> is not 3). Each cell of the grid-structured input surface will be divided in two triangles along one of its diagonals. Normally, all these diagonals have the same orientation relative to the grid. If exactly one of the four cell corners is excluded by the **undefined** handler, and this is on the usual diagonal, both triangles will be excluded. However if the default setting of **altdiagonal** is active, the other diagonal will be chosen for this cell instead, minimizing the size of the hole in the surface.

The **bentover** option controls what happens to another special case, this time in conjunction with the **trianglepattern**. For rather crumply surfaces, it can happen that the two triangles a surface cell is divided into are seen from opposite sides (i.e. the original quadrangle is 'bent over'), as illustrated in the following ASCII art:

original quadrangle:	A--B	displayed quadrangle:	C----B
("set view 0,0")	/	("set view 75,75" perhaps)	\
	/		\
	C--D		\
			A D

If the diagonal edges of the surface cells aren't generally made visible by bit 2 of the <bitpattern> there, the edge CB above wouldn't be drawn at all, normally, making the resulting display hard to understand. Therefore, the default option of **bentover** will turn it visible in this case. If you don't want that, you may choose **nobentover** instead. See also

[hidden line removal demo \(hidden.dem\)](#)

and

complex hidden line demo (singulr.dem).

Historysize

(Deprecated). **set historysize** *N* is equivalent to **set history size** *N*. **unset historysize** is equivalent to **set history size -1**.

History

Syntax:

```
set history {size <N>} {quiet|numbers} {full|trim} {default}
```

When leaving gnuplot the value of history size limits the number of lines saved to the history file. **set history size -1** allows an unlimited number of lines to be written to the history file.

By default the **history** command prints a line number in front of each command. **history quiet** suppresses the number for this command only. **set history quiet** suppresses numbers for all future **history** commands.

The **trim** option reduces the number of duplicate lines in the history list by removing earlier instances of the current command. This was the default behavior prior to gnuplot version 5.

Default settings: **set history size 500 numbers trim**.

Isosamples

The isoline density (grid) for plotting functions as surfaces may be changed by the **set isosamples** command.

Syntax:

```
set isosamples <iso_1> {,<iso_2>}
show isosamples
```

Each function surface plot will have <iso_1> iso-u lines and <iso_2> iso-v lines. If you only specify <iso_1>, <iso_2> will be set to the same value as <iso_1>. By default, sampling is set to 10 isolines per u or v axis. A higher sampling rate will produce more accurate plots, but will take longer. These parameters have no effect on data file plotting.

An isoline is a curve parameterized by one of the surface parameters while the other surface parameter is fixed. Isolines provide a simple means to display a surface. By fixing the u parameter of surface $s(u,v)$, the iso-u lines of the form $c(v) = s(u_0,v)$ are produced, and by fixing the v parameter, the iso-v lines of the form $c(u) = s(u,v_0)$ are produced.

When a function surface plot is being done without the removal of hidden lines, **set samples** controls the number of points sampled along each isoline; see **set samples** (p. 151) and **set hidden3d** (p. 120). The contour algorithm assumes that a function sample occurs at each isoline intersection, so change in **samples** as well as **isosamples** may be desired when changing the resolution of a function surface/contour.

Key

The **set key** command enables a key (or legend) containing a title and a sample (line, point, box) for each plot in the graph. The key may be turned off by requesting **set key off** or **unset key**. Individual key entries may be turned off by using the **notitle** keyword in the corresponding plot command. The text of the titles is controlled by the **set key autotitle** option or by the **title** keyword of individual **plot** and **splot** commands. See **plot title** (p. 92) for more information.

Syntax:

```
set key {on|off} {default}
    {{inside | outside} | {lmargin | rmargin | tmargin | bmargin}
     | {at <position>}}
    {left | right | center} {top | bottom | center}
```

```

{vertical | horizontal} {Left | Right}
{{no}opaque}
{{no}reverse} {{no}invert}
{samplen <sample_length>} {spacing <vertical_spacing>}
{width <width_increment>} {height <height_increment>}
{{no}autotitle {columnheader}}
{title "<text>"} {{no}enhanced}
{font "<face>,<size>"} {textcolor <colourspec>}
{{no}box {linestyle <style> | linetype <type> | linewidth <width>}}
{maxcols {<max no. of columns> | auto}}
{maxrows {<max no. of rows> | auto}}

unset key
show key

```

Elements within the key are stacked according to **vertical** or **horizontal**. In the case of **vertical**, the key occupies as few columns as possible. That is, elements are aligned in a column until running out of vertical space at which point a new column is started. The vertical space may be limited using 'maxrows'. In the case of **horizontal**, the key occupies as few rows as possible. The horizontal space may be limited using 'maxcols'.

By default the key is placed in the upper right inside corner of the graph. The keywords **left**, **right**, **top**, **bottom**, **center**, **inside**, **outside**, **lmargin**, **rmargin**, **tmargin**, **bmargin** (, **above**, **over**, **below** and **under**) may be used to automatically place the key in other positions of the graph. Also an **at <position>** may be given to indicate precisely where the plot should be placed. In this case, the keywords **left**, **right**, **top**, **bottom** and **center** serve an analogous purpose for alignment. For more information, see **key placement** (p. 124).

Justification of the plot titles within the key is controlled by **Left** or **Right** (default). The text and sample can be reversed (**reverse**) and a box can be drawn around the key (**box {...}**) in a specified **linetype** and **linewidth**, or a user-defined **linestyle**.

By default the key is built up one plot at a time. That is, the key symbol and title are drawn at the same time as the corresponding plot. That means newer plots may sometimes place elements on top of the key. **set key opaque** causes the key to be generated after all the plots. In this case the key area is filled with background color and then the key symbols and titles are written. Therefore the key itself may obscure portions of some plot elements. The default can be restored by **set key noopaque**.

By default the first plot label is at the top of the key and successive labels are entered below it. The **invert** option causes the first label to be placed at the bottom of the key, with successive labels entered above it. This option is useful to force the vertical ordering of labels in the key to match the order of box types in a stacked histogram.

The **<height_increment>** is a number of character heights to be added to or subtracted from the height of the key box. This is useful mainly when you are putting a box around the key and want larger borders around the key entries.

All plotted curves of **plots** and **splots** are titled according to the default option **autotitles**. The automatic generation of titles can be suppressed by **noautotitles**; then only those titles explicitly defined by **(s)plot ... title ...** will be drawn.

The command **set key autotitle columnheader** causes the first entry in each column of input data to be interpreted as a text string and used as a title for the corresponding plot. If the quantity being plotted is a function of data from several columns, gnuplot may be confused as to which column to draw the title from. In this case it is necessary to specify the column explicitly in the plot command, e.g.

```
plot "datafile" using (($2+$3)/$4) title columnhead(3) with lines
```

An overall title can be put on the key (**title "<text>"**) — see also **syntax** (p. 41) for the distinction between text in single- or double-quotes.

The defaults for **set key** are **on**, **right**, **top**, **vertical**, **Right**, **noreverse**, **noinvert**, **samplen 4**, **spacing 1.25**, **title ""**, and **nobox**. The default **<linetype>** is the same as that used for the plot borders. Entering **set key default** returns the key to its default configuration.

The key is drawn as a sequence of lines, with one plot described on each line. On the right-hand side (or the left-hand side, if **reverse** is selected) of each line is a representation that attempts to mimic the way the curve is plotted. On the other side of each line is the text description (the line title), obtained from the **plot** command. The lines are vertically arranged so that an imaginary straight line divides the left- and right-hand sides of the key. It is the coordinates of the top of this line that are specified with the **set key** command. In a **plot**, only the x and y coordinates are used to specify the line position. For a **splot**, x, y and z are all used as a 3D location mapped using the same mapping as the graph itself to form the required 2D screen position of the imaginary line.

When using the TeX or other terminals where formatting information is embedded in the string, **gnuplot** can only estimate the correctly exact width of the string for key positioning. If the key is to be positioned at the left, it may be convenient to use the combination **set key left Left reverse**.

If **splot** is being used to draw contours, by default the contour labels will be listed in the key. You can adjust this display using **set cntrlabel format**.

Examples:

This places the key at the default location:

```
set key default
```

This disables the key:

```
unset key
```

This places a key at coordinates 2,3.5,2 in the default (first) coordinate system:

```
set key at 2,3.5,2
```

This places the key below the graph:

```
set key below
```

This places the key in the bottom left corner, left-justifies the text, gives it a title, and draws a box around it in linetype 3:

```
set key left bottom Left title 'Legend' box 3
```

Key placement

To understand positioning, the best concept is to think of a region, i.e., inside/outside, or one of the margins. Along with the region, keywords **left/center/right** (l/c/r) and **top/center/bottom** (t/c/b) control where within the particular region the key should be placed.

When in **inside** mode, the keywords **left** (l), **right** (r), **top** (t), **bottom** (b), and **center** (c) push the key out toward the plot boundary as illustrated:

```
t/l   t/c   t/r
c/l    c    c/r
b/l   b/c   b/r
```

When in **outside** mode, automatic placement is similar to the above illustration, but with respect to the view, rather than the graph boundary. That is, a border is moved inward to make room for the key outside of the plotting area, although this may interfere with other labels and may cause an error on some devices. The particular plot border that is moved depends upon the position described above and the stacking direction. For options centered in one of the dimensions, there is no ambiguity about which border to move. For the corners, when the stack direction is **vertical**, the left or right border is moved inward appropriately. When the stack direction is **horizontal**, the top or bottom border is moved inward appropriately.

The margin syntax allows automatic placement of key regardless of stack direction. When one of the margins **lmargin** (lm), **rmargin** (rm), **tmargin** (tm), and **bmargin** (bm) is combined with a single, non-conflicting direction keyword, the following illustrated positions may contain the key:

```
l/tm  c/tm  r/tm
```

t/lm	t/rm
c/lm	c/rm
b/lm	b/rm
l/bm	c/bm
r/bm	

Keywords **above** and **over** are synonymous with **tmargin**. For version compatibility, **above** or **over** without an additional l/c/r or stack direction keyword uses **center** and **horizontal**. Keywords **below** and **under** are synonymous with **bmargin**. For compatibility, **below** or **under** without an additional l/c/r or stack direction keyword uses **center** and **horizontal**. A further compatibility issue is that **outside** appearing without an additional t/b/c or stack direction keyword uses **top**, **right** and **vertical** (i.e., the same as t/rm above).

The <position> can be a simple x,y,z as in previous versions, but these can be preceded by one of five keywords (**first**, **second**, **graph**, **screen**, **character**) which selects the coordinate system in which the position of the first sample line is specified. See **coordinates** (p. 21) for more details. The effect of **left**, **right**, **top**, **bottom**, and **center** when <position> is given is to align the key as though it were text positioned using the label command, i.e., **left** means left align with key to the right of <position>, etc.

Key samples

By default, each plot on the graph generates a corresponding entry in the key. This entry contains a plot title and a sample line/point/box of the same color and fill properties as used in the plot itself. The font and textcolor properties control the appearance of the individual plot titles that appear in the key. Setting the textcolor to "variable" causes the text for each key entry to be the same color as the line or fill color for that plot. This was the default in some earlier versions of gnuplot.

The length of the sample line can be controlled by **samplelen**. The sample length is computed as the sum of the tic length and <sample.length> times the character width. **samplelen** also affects the positions of point samples in the key since these are drawn at the midpoint of the sample line, even if the sample line itself is not drawn.

The vertical spacing between lines is controlled by **spacing**. The spacing is set equal to the product of the pointsize, the vertical tic size, and <vertical.spacing>. The program will guarantee that the vertical spacing is no smaller than the character height.

The <width.increment> is a number of character widths to be added to or subtracted from the length of the string. This is useful only when you are putting a box around the key and you are using control characters in the text. **gnuplot** simply counts the number of characters in the string when computing the box width; this allows you to correct it.

Label

Arbitrary labels can be placed on the plot using the **set label** command.

Syntax:

```
set label {<tag>} {"<label text>"} {at <position>}
        {left | center | right}
        {norotate | rotate {by <degrees>}}
        {font "<name>{,<size>}" }
        {noenhanced}
        {front | back}
        {textcolor <colourspec>}
        {point <pointstyle> | nopoint}
        {offset <offset>}
        {boxed}
        {hypertext}
unset label {<tag>}
show label
```

The `<position>` is specified by either `x,y` or `x,y,z`, and may be preceded by **first**, **second**, **graph**, **screen**, or **character** to select the coordinate system. See **coordinates** (p. 21) for details.

The tag is an integer that is used to identify the label. If no `<tag>` is given, the lowest unused tag value is assigned automatically. The tag can be used to delete or modify a specific label. To change any attribute of an existing label, use the **set label** command with the appropriate tag, and specify the parts of the label to be changed.

The `<label text>` can be a string constant, a string variable, or a string-valued expression. See **strings** (p. 39), **sprintf** (p. 26), and **gprintf** (p. 117).

By default, the text is placed flush left against the point `x,y,z`. To adjust the way the label is positioned with respect to the point `x,y,z`, add the justification parameter, which may be **left**, **right** or **center**, indicating that the point is to be at the left, right or center of the text. Labels outside the plotted boundaries are permitted but may interfere with axis labels or other text.

Some terminals support enclosing the label in a box. See **set style textbox** (p. 158). Note: Currently the boxed enclosure is limited to unrotated text.

If **rotate** is given, the label is written vertically (if the terminal can do so, of course). If **rotate by <degrees>** is given, conforming terminals will try to write the text at the specified angle; non-conforming terminals will treat this as vertical text.

Font and its size can be chosen explicitly by **font "<name>{,<size>}"** if the terminal supports font settings. Otherwise the default font of the terminal will be used.

Normally the enhanced text mode string interpretation, if enabled for the current terminal, is applied to all text strings including label text. The **noenhanced** property can be used to exempt a specific label from the enhanced text mode processing. This can be useful if the label contains underscores, for example. See **enhanced text** (p. 22).

If **front** is given, the label is written on top of the graphed data. If **back** is given (the default), the label is written underneath the graphed data. Using **front** will prevent a label from being obscured by dense data.

textcolor <colourspec> changes the color of the label text. `<colourspec>` can be a linetype, an rgb color, or a palette mapping. See help for **colourspec** (p. 33) and **palette** (p. 144). **textcolor** may be abbreviated **tc**.

```
'tc default' resets the text color to its default state.
'tc lt <n>' sets the text color to that of line type <n>.
'tc ls <n>' sets the text color to that of line style <n>.
'tc palette z' selects a palette color corresponding to the label z position.
'tc palette cb <val>' selects a color corresponding to <val> on the colorbar.
'tc palette fraction <val>', with 0<=val<=1, selects a color corresponding to
the mapping [0:1] to grays/colors of the 'palette'.
'tc rgb "#RRGGBB" or 'tc rgb "0xRRGGBB" sets an arbitrary 24-bit RGB color.
'tc rgb 0xRRGGBB' As above; a hexadecimal constant does not require quotes.
```

If a `<pointstyle>` is given, using keywords **lt**, **pt** and **ps**, see **style** (p. 93), a point with the given style and color of the given line type is plotted at the label position and the text of the label is displaced slightly. This option is used by default for placing labels in **mouse** enhanced terminals. Use **nopoint** to turn off the drawing of a point near the label (this is the default).

The displacement defaults to 1,1 in **pointsize** units if a `<pointstyle>` is given, 0,0 if no `<pointstyle>` is given. The displacement can be controlled by the optional **offset <offset>** where `<offset>` is specified by either `x,y` or `x,y,z`, and may be preceded by **first**, **second**, **graph**, **screen**, or **character** to select the coordinate system. See **coordinates** (p. 21) for details.

If one (or more) axis is timeseries, the appropriate coordinate should be given as a quoted time string according to the **timefmt** format string. See **set xdata** (p. 165) and **set timefmt** (p. 161).

The options available for **set label** are also available for the **labels** plot style. See **labels** (p. 54).

Examples

Examples:

To set a label at (1,2) to "y=x", use:

```
set label "y=x" at 1,2
```

To set a Sigma of size 24, from the Symbol font set, at the center of the graph, use:

```
set label "S" at graph 0.5,0.5 center font "Symbol,24"
```

To set a label "y=x^2" with the right of the text at (2,3,4), and tag the label as number 3, use:

```
set label 3 "y=x^2" at 2,3,4 right
```

To change the preceding label to center justification, use:

```
set label 3 center
```

To delete label number 2, use:

```
unset label 2
```

To delete all labels, use:

```
unset label
```

To show all labels (in tag order), use:

```
show label
```

To set a label on a graph with a timeseries on the x axis, use, for example:

```
set timefmt "%d/%m/%y,%H:%M"
set label "Harvest" at "25/8/93",1
```

To display a freshly fitted parameter on the plot with the data and the fitted function, do this after the **fit**, but before the **plot**:

```
set label sprintf("a = %3.5g",par_a) at 30,15
bfit = gprintf("b = %s*10^%S",par_b)
set label bfit at 30,20
```

To display a function definition along with its fitted parameters, use:

```
f(x)=a+b*x
fit f(x) 'datafile' via a,b
set label GPFUN_f at graph .05,.95
set label sprintf("a = %g", a) at graph .05,.90
set label sprintf("b = %g", b) at graph .05,.85
```

To set a label displaced a little bit from a small point:

```
set label 'origin' at 0,0 point lt 1 pt 2 ps 3 offset 1,-1
```

To set a label whose color matches the z value (in this case 5.5) of some point on a 3D splot colored using pm3d:

```
set label 'text' at 0,0,5.5 tc palette z
```

Hypertext

Some terminals (wxt, qt, svg, canvas, win) allow you to attach hypertext to specific points on the graph or elsewhere on the canvas. When the mouse hovers over the anchor point, a pop-up box containing the text is displayed. Terminals that do not support hypertext will display nothing. You must enable the **point** attribute of the label in order for the hypertext to be anchored. Examples:

```
set label at 0,0 "Plot origin" hypertext point pt 1
plot 'data' using 1:2:0 with labels hypertext point pt 7 \
title 'mouse over point to see its order in data set'
```


Linetype

The **set linetype** command allows you to redefine the basic linetypes used for plots. The command options are identical to those for "set style line". Unlike line styles, redefinitions by **set linetype** are persistent; they are not affected by **reset**.

For example, linetypes one and two default to red and green. If you redefine them like this:

```
set linetype 1 lw 2 lc rgb "blue" pointtype 6
set linetype 2 lw 2 lc rgb "forest-green" pointtype 8
```

everywhere that uses lt 1 will now get a thick blue line rather than a thin red line (the previous default meaning of lt 1). This includes uses such as the definition of a temporary linestyle derived from the base linetype 1.

Note: This command was introduced in gnuplot version 4.6. It supersedes an older rather cryptic command "set style increment user". The older command is now deprecated.

This mechanism can be used to define a set of personal preferences for the sequence of lines used in gnuplot. The recommended way to do this is to add to the run-time initialization file `~/.gnuplot` a sequence of commands like

```
set linetype 1 lc rgb "dark-violet" lw 2 pt 0
set linetype 2 lc rgb "sea-green"    lw 2 pt 7
set linetype 3 lc rgb "cyan"        lw 2 pt 6 pi -1
set linetype 4 lc rgb "dark-red"    lw 2 pt 5 pi -1
set linetype 5 lc rgb "blue"        lw 2 pt 8
set linetype 6 lc rgb "dark-orange" lw 2 pt 3
set linetype 7 lc rgb "black"       lw 2 pt 11
set linetype 8 lc rgb "goldenrod"   lw 2
set linetype cycle 8
```

Every time you run gnuplot the line types will be initialized to these values. You may initialize as many linetypes as you like. If you do not redefine, say, linetype 3 then it will continue to have the default properties (in this case blue, pt 3, lw 1, etc). The first few lines of the example script insure that the commands will be skipped by older versions of gnuplot.

Similar script files can be used to define theme-based color choices, or sets of colors optimized for a particular plot type or output device.

The command **set linetype cycle 8** tells gnuplot to re-use these definitions for the color and linewidth of higher-numbered linetypes. That is, linetypes 9-16, 17-24, and so on will use this same sequence of colors and widths. The point properties (pointtype, pointsize, pointinterval) are not affected by this command. **unset linetype cycle** disables this feature. If the line properties of a higher numbered linetype are explicitly defined, this takes precedence over the recycled low-number linetype properties.

Link

Syntax:

```
set link {x2 | y2} {via <expression1> inverse <expression2>}
unset link
```

The **set link** command establishes a mapping between the x and x2 axes, or the y and y2 axes. `<expression1>` maps primary axis coordinates onto the secondary axis. `<expression2>` maps secondary axis coordinates onto the primary axis.

Examples:

```
set link x
```

This is the simplest form of the command. It forces the x2 axis to have identically the same range, scale, and direction as the x axis. Commands **set xrange**, **set auto x**, etc will affect both the x and x2 axes. Commands **set x2range**, etc, will be ignored while the linkage is in effect.


```
set link x2 via x**2 inverse sqrt(x)
plot "sqrt_data" using 1:2 axes x2y1, "linear_data" using 1:2 axes x1y1
```

This command establishes forward and reverse mapping between the x and $x2$ axes. The forward mapping is used to generate $x2$ tic labels and $x2$ mouse coordinate. The reverse mapping is used to plot coordinates given in the $x2$ coordinate system. Note that the mapping as given is valid only for x non-negative. When mapping to the $y2$ axis, both $\langle\text{expression1}\rangle$ and $\langle\text{expression2}\rangle$ must use y as dummy variable.

Lmargin

The command **set lmargin** sets the size of the left margin. Please see **set margin** (p. 131) for details.

Loadpath

The **loadpath** setting defines additional locations for data and command files searched by the **call**, **load**, **plot** and **splot** commands. If a file cannot be found in the current directory, the directories in **loadpath** are tried.

Syntax:

```
set loadpath {"pathlist1" {"pathlist2"...}}
show loadpath
```

Path names may be entered as single directory names, or as a list of path names separated by a platform-specific path separator, eg. colon (':') on Unix, semicolon (';') on DOS/Windows/OS/2 platforms. The **show loadpath**, **save** and **save set** commands replace the platform-specific separator with a space character (' ').

If the environment variable `GNUPLOT_LIB` is set, its contents are appended to **loadpath**. However, **show loadpath** prints the contents of **set loadpath** and `GNUPLOT_LIB` separately. Also, the **save** and **save set** commands ignore the contents of `GNUPLOT_LIB`.

Locale

The **locale** setting determines the language with which **{x,y,z}{d,m}tics** will write the days and months.

Syntax:

```
set locale {"<locale>"}
```

$\langle\text{locale}\rangle$ may be any language designation acceptable to your installation. See your system documentation for the available options. The command **set locale ""** will try to determine the locale from the `LC_TIME`, `LC_ALL`, or `LANG` environment variables.

To change the decimal point locale, see **set decimalsign** (p. 111). To change the character encoding to the current locale, see **set encoding** (p. 114).

Logscale

Syntax:

```
set logscale <axes> {<base>}
unset logscale <axes>
show logscale
```

where $\langle\text{axes}\rangle$ may be any combinations of **x**, **x2**, **y**, **y2**, **z**, **cb**, and **r** in any order. $\langle\text{base}\rangle$ is the base of the log scaling (default is base 10). If no axes are specified, the command affects all axes except **r**. The command **unset logscale** turns off log scaling for all axes. Note that the ticmarks generated for logscaled axes are not uniformly spaced. See **set xtics** (p. 168).

Examples:

To enable log scaling in both x and z axes:

```
set logscale xz
```

To enable scaling log base 2 of the y axis:

```
set logscale y 2
```

To enable z and color log axes for a pm3d plot:

```
set logscale zcb
```

To disable z axis log scaling:

```
unset logscale z
```

Macros

If command line macro substitution is enabled, then tokens in the command line of the form @<stringvariablename> will be replaced by the text string contained in <stringvariablename>. See **substitution** (p. 40).

Syntax:

```
set macros
```

Mapping

If data are provided to **splot** in spherical or cylindrical coordinates, the **set mapping** command should be used to instruct **gnuplot** how to interpret them.

Syntax:

```
set mapping {cartesian | spherical | cylindrical}
```

A cartesian coordinate system is used by default.

For a spherical coordinate system, the data occupy two or three columns (or **using** entries). The first two are interpreted as the azimuthal and polar angles theta and phi (or "longitude" and "latitude"), in the units specified by **set angles**. The radius r is taken from the third column if there is one, or is set to unity if there is no third column. The mapping is:

```
x = r * cos(theta) * cos(phi)
y = r * sin(theta) * cos(phi)
z = r * sin(phi)
```

Note that this is a "geographic" spherical system, rather than a "polar" one (that is, phi is measured from the equator, rather than the pole).

For a cylindrical coordinate system, the data again occupy two or three columns. The first two are interpreted as theta (in the units specified by **set angles**) and z. The radius is either taken from the third column or set to unity, as in the spherical case. The mapping is:

```
x = r * cos(theta)
y = r * sin(theta)
z = z
```

The effects of **mapping** can be duplicated with the **using** filter on the **splot** command, but **mapping** may be more convenient if many data files are to be processed. However even if **mapping** is used, **using** may still be necessary if the data in the file are not in the required order.

mapping has no effect on **plot**.

[world.dem](#): [mapping demos](#).

Margin

The **margin** is the distance between the plot border and the outer edge of the canvas. The size of the margin is chosen automatically, but can be overridden by the **set margin** commands. **show margin** shows the current settings. To alter the distance between the inside of the plot border and the data in the plot itself, see **set offsets** (p. 139).

Syntax:

```
set lmargin {{at screen} <margin>}
set rmargin {{at screen} <margin>}
set tmargin {{at screen} <margin>}
set bmargin {{at screen} <margin>}
set margins <left>, <right>, <bottom>, <top>
show margin
```

The default units of <margin> are character heights or widths, as appropriate. A positive value defines the absolute size of the margin. A negative value (or none) causes **gnuplot** to revert to the computed value. For 3D plots, only the left margin can be set using character units.

The keywords **at screen** indicates that the margin is specified as a fraction of the full drawing area. This can be used to precisely line up the corners of individual 2D and 3D graphs in a multiplot. This placement ignores the current values of **set origin** and **set size**, and is intended as an alternative method for positioning graphs within a multiplot.

Normally the margins of a plot are automatically calculated based on tics, tic labels, axis labels, the plot title, the timestamp and the size of the key if it is outside the borders. If, however, tics are attached to the axes (**set xtics axis**, for example), neither the tics themselves nor their labels will be included in either the margin calculation or the calculation of the positions of other text to be written in the margin. This can lead to tic labels overwriting other text if the axis is very close to the border.

Monochrome

Syntax:

```
set monochrome {linetype N <linetype properties>}
```

The **set monochrome** command selects an alternative set of linetypes that differ by dot/dash pattern or line width rather than by color. This command replaces the monochrome option offered by certain terminal types in earlier versions of gnuplot. For backward compatibility these terminal types now implicitly invoke "set monochrome" if their own "mono" option is present. For example,

```
set terminal pdf mono
```

is equivalent to

```
set terminal pdf
set mono
```

Selecting monochrome mode does not prevent you from explicitly drawing lines using RGB or palette colors, but see also **set palette gray** (p. 147). Six monochrome linetypes are defined by default. You can change their properties or add additional monochrome linetypes by using the full form of the command. Changes made to the monochrome linetypes do not affect the color linetypes and vice versa. To restore the usual set of color linetypes, use either **unset monochrome** or **set color**.

Mouse

The command **set mouse** enables mouse actions for the current interactive terminal. It is usually enabled by default in interactive mode, but disabled by default if commands are being read from a file.

There are two mouse modes. The 2D mode works for **plot** commands and for **splot** maps (i.e. **set view** with z-rotation 0, 90, 180, 270 or 360 degrees, including **set view map**). In this mode the mouse position is

tracked and you can pan or zoom using the mouse buttons or arrow keys. Some terminals support toggling individual plots on/off by clicking on the corresponding key title or on a separate widget.

For 3D graphs **splot**, the view and scaling of the graph can be changed with mouse buttons 1 and 2, respectively. A vertical motion of Button 2 with the shift key held down changes the **xyplane**. If additionally to these buttons the modifier <ctrl> is held down, the coordinate axes are displayed but the data are suppressed. This is useful for large data sets.

Mousing is not available inside multiplot mode. When multiplot is completed using **unset multiplot**, then the mouse will be turned on again but acts only on the most recent plot within the multiplot (like replot does).

Syntax:

```
set mouse {doubleclick <ms>} {nodoubleclick}
        {{no}zoomcoordinates}
        {zoomfactors <xmultiplier>, <ymultiplier>}
        {noruler | ruler {at x,y}}
        {polardistance{deg|tan} | nopolardistance}
        {format <string>}
        {mouseformat <int>/<string>}
        {{no}labels {"labeloptions"}}
        {{no}zoomjump} {{no}verbose}

unset mouse
```

The options **noruler** and **ruler** switch the ruler off and on, the latter optionally setting the origin at the given coordinates. While the ruler is on, the distance in user units from the ruler origin to the mouse is displayed continuously. By default, toggling the ruler has the key binding 'r'.

The option **polardistance** determines if the distance between the mouse cursor and the ruler is also shown in polar coordinates (distance and angle in degrees or tangent (slope)). This corresponds to the default key binding '5'.

Choose the option **labels** to define persistent gnuplot labels using Button 2. The default is **no**labels, which makes Button 2 draw only a temporary label at the mouse position. Labels are drawn with the current setting of **mouseformat**. The **labeloptions** string is passed to the **set label** command. The default is "point pointstyle 1" which will plot a small plus at the label position. Temporary labels will disappear at the next **replot** or mouse zoom operation. Persistent labels can be removed by holding the Ctrl-Key down while clicking Button 2 on the label's point. The threshold for how close you must be to the label is also determined by the **pointsize**.

If the option **verbose** is turned on the communication commands are shown during execution. This option can also be toggled by hitting **6** in the driver's window. **verbose** is off by default.

Press 'h' in the driver's window for a short summary of the mouse and key bindings. This will also display user defined bindings or **hotkeys** which can be defined using the **bind** command, see help for **bind** (p. 36). Note, that user defined **hotkeys** may override the default bindings. See also help for **bind** (p. 36) and **label** (p. 125).

Doubleclick

The doubleclick resolution is given in milliseconds and used for Button 1, which copies the current mouse position to the **clipboard** on some terminals. The default value is 300 ms. Setting the value to 0 ms triggers the copy on a single click.

Mouseformat

The **set mouse format** command specifies a format string for `sprintf()` which determines how the mouse cursor [x,y] coordinates are printed to the plot window and to the clipboard. The default is "% #g".

set mouse mouseformat is used for formatting the text on Button1 and Button2 actions – copying the coordinates to the clipboard and temporarily annotating the mouse position. An integer argument selects

one of the format options in the table below. A string argument is used as a format for `sprintf()` in option 6 and should contain two float specifiers. Example:

```
'set mouse mouseformat "mouse x,y = %5.2g, %10.3f"'
```

Use **set mouse mouseformat ""** to turn this string off again.

The following formats are available:

0	default (same as 1)	
1	axis coordinates	1.23, 2.45
2	graph coordinates (from 0 to 1)	/0.00, 1.00/
3	x = timefmt y = axis	[(as set by 'set timefmt'), 2.45]
4	x = date y = axis	[31. 12. 1999, 2.45]
5	x = time y = axis	[23:59, 2.45]
6	x = date time y = axis	[31. 12. 1999 23:59, 2.45]
7	format from 'set mouse mouseformat', e.g. "mouse x,y = 1.23,	2.450"

Scrolling

X and Y axis scaling in both 2D and 3D graphs can be adjusted using the mouse wheel. `<wheel-up>` scrolls up (increases both YMIN and YMAX by ten percent of the Y range, and increases both Y2MIN and Y2MAX likewise), and `<wheel down>` scrolls down. `<shift-wheel-up>` scrolls left (decreases both XMIN and XMAX, and both X2MIN and X2MAX), and `<shift-wheel-down>` scrolls right. `<control-wheel-up>` zooms in toward the center of the plot, and `<control-wheel-down>` zooms out. `<shift-control-wheel-up>` zooms in along the X and X2 axes only, and `<shift-control-wheel-down>` zooms out along the X and X2 axes only.

X11 mouse

If multiple X11 plot windows have been opened using the **set term x11 <n>** terminal option, then only the current plot window supports the entire range of mouse commands and hotkeys. The other windows will, however, continue to display mouse coordinates at the lower left.

Zoom

Zooming is usually accomplished by holding down the left mouse button and dragging the mouse to delineate a zoom region. Some platforms may require using a different mouse button. The original plot can be restored by typing the 'u' hotkey in the plot window. The hotkeys 'p' and 'n' step back and forth through a history of zoom operations.

The option **zoomcoordinates** determines if the coordinates of the zoom box are drawn at the edges while zooming. This is on by default.

If the option **zoomjump** is on, the mouse pointer will be automatically offset a small distance after starting a zoom region with button 3. This can be useful to avoid a tiny (or even empty) zoom region. **zoomjump** is off by default.

Multiplot

The command **set multiplot** places **gnuplot** in the multiplot mode, in which several plots are placed on the same page, window, or screen.

Syntax:

```
set multiplot
{ title <page title> {font <fontspec>} {enhanced|noenhanced} }
{ layout <rows>,<cols>
  {rowsfirst|columnsfirst} {downwards|upwards}
  {scale <xscale>{,<yscale>}} {offset <xoff>{,<yoff>}}
```

```

    {margins <left>,<right>,<bottom>,<top>}
    {spacing <xspacing>{,<yspacing>}}
}
set multiplot {next|previous}
unset multiplot

```

For some terminals, no plot is displayed until the command **unset multiplot** is given, which causes the entire page to be drawn and then returns gnuplot to its normal single-plot mode. For other terminals, each separate **plot** command produces an updated display.

The **clear** command is used to erase the rectangular area of the page that will be used for the next plot. This is typically needed for "inset" plots.

Any labels or arrows that have been defined will be drawn for each plot according to the current size and origin (unless their coordinates are defined in the **screen** system). Just about everything else that can be **set** is applied to each plot, too. If you want something to appear only once on the page, for instance a single time stamp, you'll need to put a **set time/unset time** pair around one of the **plot**, **splot** or **replot** commands within the **set multiplot/unset multiplot** block.

The multiplot title is separate from the individual plot titles, if any. Space is reserved for it at the top of the page, spanning the full width of the canvas.

The commands **set origin** and **set size** must be used to correctly position each plot if no layout is specified or if fine tuning is desired. See **set origin** (p. 139) and **set size** (p. 151) for details of their usage.

Example:

```

set multiplot
set size 0.4,0.4
set origin 0.1,0.1
plot sin(x)
set size 0.2,0.2
set origin 0.5,0.5
plot cos(x)
unset multiplot

```

This displays a plot of $\cos(x)$ stacked above a plot of $\sin(x)$.

set size and **set origin** refer to the entire plotting area used for each plot. Please also see **set term size** (p. 20). If you want to have the axes themselves line up, you can guarantee that the margins are the same size with the **set margin** commands. See **set margin** (p. 131) for their use. Note that the margin settings are absolute, in character units, so the appearance of the graph in the remaining space will depend on the screen size of the display device, e.g., perhaps quite different on a video display and a printer.

With the **layout** option you can generate simple multiplots without having to give the **set size** and **set origin** commands before each plot: Those are generated automatically, but can be overridden at any time. With **layout** the display will be divided by a grid with **<rows>** rows and **<cols>** columns. This grid is filled rows first or columns first depending on whether the corresponding option is given in the multiplot command. The stack of plots can grow **downwards** or **upwards**. Default is **rowsfirst** and **downwards**. The commands **set multiplot next** and **set multiplot previous** are relevant only in the context of using the layout option. **next** skips the next position in the grid, leaving a blank space. **prev** returns to the grid position immediately preceding the most recently plotted position.

Each plot can be scaled by **scale** and shifted with **offset**; if the y-values for scale or offset are omitted, the x-value will be used. **unset multiplot** will turn off the automatic layout and restore the values of **set size** and **set origin** as they were before **set multiplot layout**.

Example:

```

set size 1,1
set origin 0,0
set multiplot layout 3,2 columnsfirst scale 1.1,0.9
[ up to 6 plot commands here ]
unset multiplot

```

The above example will produce 6 plots in 2 columns filled top to bottom, left to right. Each plot will have a horizontal size of $1.1/2$ and a vertical size of $0.9/3$.

Another possibility is to set uniform margins for all plots in the layout with options **layout margins** and **spacing**, which must be used together. With **margins** you set the outer margins of the whole multiplot grid.

spacing gives the gap size between two adjacent subplots, and can also be given in **character** or **screen** units. If a single value is given, it is used for both x and y direction, otherwise two different values can be selected.

If one value has no unit, the one of the preceding margin setting is used.

Example:

```
set multiplot layout 2,2 margins 0.1, 0.9, 0.1, 0.9 spacing 0.0
```

In this case the two left-most subplots will have left boundaries at screen coordinate 0.1, the two right-most subplots will have right boundaries at screen coordinate 0.9, and so on. Because the spacing between subplots is given as 0, their inner boundaries will superimpose.

Example:

```
set multiplot layout 2,2 margins char 5,1,1,2 spacing screen 0, char 2
```

This produces a layout in which the boundary of both left subplots is 5 character widths from the left edge of the canvas, the right boundary of the right subplots is 1 character width from the canvas edge. The overall bottom margin is one character height and the overall top margin is 2 character heights. There is no horizontal gap between the two columns of subplots. The vertical gap between subplots is equal to 2 character heights.

Example:

```
set multiplot layout 2,2 columnsfirst margins 0.1,0.9,0.1,0.9 spacing 0.1
set ylabel 'ylabel'
plot sin(x)
set xlabel 'xlabel'
plot cos(x)
unset ylabel
unset xlabel
plot sin(2*x)
set xlabel 'xlabel'
plot cos(2*x)
unset multiplot
```

See also

[multiplot demo \(multiplt.dem\)](#)

Mx2tics

Minor tic marks along the x2 (top) axis are controlled by **set mx2tics**. Please see **set mxtics** (p. 135).

Mxtics

Minor tic marks along the x axis are controlled by **set mxtics**. They can be turned off with **unset mxtics**. Similar commands control minor tics along the other axes.

Syntax:

```
set mxtics {<freq> | default}
unset mxtics
show mxtics
```

The same syntax applies to **mytics**, **mztics**, **mx2tics**, **my2tics**, **mrtics** and **mc2tics**.

<freq> is the number of sub-intervals (NOT the number of minor tics) between major tics (the default for a linear axis is either two or five depending on the major tics, so there are one or four minor tics between major tics). Selecting **default** will return the number of minor ticks to its default value.

If the axis is logarithmic, the number of sub-intervals will be set to a reasonable number by default (based upon the length of a decade). This will be overridden if <freq> is given. However the usual minor tics (2, 3, ..., 8, 9 between 1 and 10, for example) are obtained by setting <freq> to 10, even though there are but nine sub-intervals.

To set minor tics at arbitrary positions, use the ("**<label>**" **<pos>** **<level>**, ...) form of **set {x|x2|y|y2|z}tics** with **<label>** empty and **<level>** set to 1.

The **set m{x|x2|y|y2|z}tics** commands work only when there are uniformly spaced major tics. If all major tics were placed explicitly by **set {x|x2|y|y2|z}tics**, then minor tic commands are ignored. Implicit major tics and explicit minor tics can be combined using **set {x|x2|y|y2|z}tics** and **set {x|x2|y|y2|z}tics add**.

Examples:

```
set xtics 0, 5, 10
set xtics add (7.5)
set mxtics 5
```

Major tics at 0,5,7.5,10, minor tics at 1,2,3,4,6,7,8,9

```
set logscale y
set ytics format ""
set ytics 1e-6, 10, 1
set ytics add ("1" 1, ".1" 0.1, ".01" 0.01, "10^-3" 0.001, \
               "10^-4" 0.0001)
set mytics 10
```

Major tics with special formatting, minor tics at log positions

By default, minor tics are off for linear axes and on for logarithmic axes. They inherit the settings for **axis|border** and **{no}mirror** specified for the major tics. Please see **set xtics** (p. 168) for information about these.

My2tics

Minor tic marks along the y2 (right-hand) axis are controlled by **set my2tics**. Please see **set mxtics** (p. 135).

Mytics

Minor tic marks along the y axis are controlled by **set mytics**. Please see **set mxtics** (p. 135).

Mztics

Minor tic marks along the z axis are controlled by **set mztics**. Please see **set mxtics** (p. 135).

Object

The **set object** command defines a single object which will appear in all subsequent 2D plots. You may define as many objects as you like. Currently the supported object types are **rectangle**, **circle**, **ellipse**, and **polygon**. Rectangles inherit a default set of style properties (fill, color, border) from those set by the command **set style rectangle**, but each object can also be given individual style properties. Circles, ellipses, and polygons inherit the fill style from **set style fill**.

Syntax:

```
set object <index>
```



```

<object-type> <object-properties>
{front|back|behind} {clip|noclip}
{fc|fillcolor <colorespec>} {fs <fillstyle>}
{default} {lw|linewidth <width>} {dt|dashtype <dashtype>}
unset object <index>

```

<object-type> is either **rectangle**, **ellipse**, **circle**, or **polygon**. Each object type has its own set of characteristic properties.

Setting **front** will draw the object in front of all plot elements, but behind any labels that are also marked **front**. Setting **back** will place the object behind all plot curves and labels. Setting **behind** will place the object behind everything including the axes and **back** rectangles, thus

```
set object rectangle from screen 0,0 to screen 1,1 behind
```

can be used to provide a colored background for the entire graph or page.

By default, objects are clipped to the graph boundary unless one or more vertices are given in screen coordinates. Setting **noclip** will disable clipping to the graph boundary, but will still clip against the screen size.

The fill color of the object is taken from the <colorespec>. **fillcolor** may be abbreviated **fc**. The fill style is taken from <fillstyle>. See **colorespec** (p. 33) and **fillstyle** (p. 155). If the keyword **default** is given, these properties are inherited from the default settings at the time a plot is drawn. See **set style rectangle** (p. 158).

Rectangle

Syntax:

```

set object <index> rectangle
{from <position> {to|rto} <position> |
 center <position> size <w>,<h> |
 at <position> size <w>,<h>}}

```

The position of the rectangle may be specified by giving the position of two diagonal corners (bottom left and top right) or by giving the position of the center followed by the width and the height. In either case the positions may be given in axis, graph, or screen coordinates. See **coordinates** (p. 21). The options **at** and **center** are synonyms.

Examples:

```

# Force the entire area enclosed by the axes to have background color cyan
set object 1 rect from graph 0, graph 0 to graph 1, graph 1 back
set object 1 rect fc rgb "cyan" fillstyle solid 1.0

# Position a red square with lower left at 0,0 and upper right at 2,3
set object 2 rect from 0,0 to 2,3 fc lt 1

# Position an empty rectangle (no fill) with a blue border
set object 3 rect from 0,0 to 2,3 fs empty border rgb "blue"

# Return fill and color to the default style but leave vertices unchanged
set object 2 rect default

```

Rectangle corners specified in screen coordinates may extend beyond the edge of the current graph. Otherwise the rectangle is clipped to fit in the graph.

Ellipse

Syntax:

```
set object <index> ellipse {at|center} <position> size <w>,<h>
    {angle <orientation>} {units xy|xx|yy}
    {<other-object-properties>}
```

The position of the ellipse is specified by giving the center followed by the width and the height (actually the major and minor axes). The keywords **at** and **center** are synonyms. The center position may be given in axis, graph, or screen coordinates. See **coordinates** (p. 21). The major and minor axis lengths must be given in axis coordinates. The orientation of the ellipse is specified by the angle between the horizontal axis and the major diameter of the ellipse. If no angle is given, the default ellipse orientation will be used instead (see **set style ellipse** (p. 158)). The **units** keyword controls the scaling of the axes of the ellipse. **units xy** means that the major axis is interpreted in terms of units along the x axis, while the minor axis is that of the y axis. **units xx** means that both axes of the ellipses are scaled in the units of the x axis, while **units yy** means that both axes are in units of the y axis. The default is **xy** or whatever **set style ellipse units** was set to.

NB: If the x and y axis scales are not equal, (e.g. **units xy** is in effect) then the major/minor axis ratio will no longer be correct after rotation.

Note that **set object ellipse size <2r>,<2r>** does not in general produce the same result as **set object circle <r>**. The circle radius is always interpreted in terms of units along the x axis, and will always produce a circle even if the x and y axis scales are different and even if the aspect ratio of your plot is not 1. If **units** is set to **xy**, then 'set object ellipse' interprets the first <2r> in terms of x axis units and the second <2r> in terms of y axis units. This will only produce a circle if the x and y axis scales are identical and the plot aspect ratio is 1. On the other hand, if **units** is set to **xx** or **yy**, then the diameters specified in the 'set object' command will be interpreted in the same units, so the ellipse will have the correct aspect ratio, and it will maintain its aspect ratio even if the plot is resized.

Circle

Syntax:

```
set object <index> circle {at|center} <position> size <radius>
    {arc [<begin>:<end>]}
    {<other-object-properties>}
```

The position of the circle is specified by giving the position of the center followed by the radius. The keywords **at** and **center** are synonyms. The position and radius may be given in x-axis, graph, or canvas coordinates. See **coordinates** (p. 21). In all cases the radius is calculated relative to the horizontal scale of the axis, graph, or canvas. Any disparity between the horizontal and vertical scaling will be corrected for so that the result is always a circle. If you want to draw a circle in plot coordinates (such that it will appear as an ellipse if the horizontal and vertical scales are different), use **set object ellipse** instead.

By default a full circle is drawn. The optional qualifier **arc** specifies a starting angle and ending angle, in degrees, for one arc of the circle. The arc is always drawn counterclockwise.

See also **set object ellipse** (p. 137).

Polygon

Syntax:

```
set object <index> polygon
    from <position> to <position> ... {to <position>}
```

or

```
    from <position> rto <position> ... {rto <position>}
```

The position of the polygon may be specified by giving the position of a sequence of vertices. These may be given in axis, graph, or screen coordinates. If relative coordinates are used (rto) then the coordinate type must match that of the previous vertex. See **coordinates** (p. 21).

Example:

```
set object 1 polygon from 0,0 to 1,1 to 2,0
set object 1 fc rgb "cyan" fillstyle solid 1.0 border lt -1
```

Offsets

Offsets provide a mechanism to put an empty boundary around the data inside an autoscaled graph. The offsets only affect the x1 and y1 axes, and only in 2D **plot** commands.

Syntax:

```
set offsets <left>, <right>, <top>, <bottom>
unset offsets
show offsets
```

Each offset may be a constant or an expression. Each defaults to 0. By default, the left and right offsets are given in units of the first x axis, the top and bottom offsets in units of the first y axis. Alternatively, you may specify the offsets as a fraction of the total axis range by using the keyword "graph". A positive offset expands the axis range in the specified direction, e.g., a positive bottom offset makes ymin more negative. Negative offsets, while permitted, can have unexpected interactions with autoscaling and clipping. To prevent the auto-scaling from further adjusting your axis ranges, it is useful to also specify "set auto fix".

Example:

```
set auto fix
set offsets graph 0.05, 0, 2, 2
plot sin(x)
```

This graph of $\sin(x)$ will have a y range $[-3:3]$ because the function will be autoscaled to $[-1:1]$ and the vertical offsets are each two. The x range will be $[-11:10]$ because the default is $[-10:10]$ and it has been expanded to the left by 0.05 of that total range.

Origin

The **set origin** command is used to specify the origin of a plotting surface (i.e., the graph and its margins) on the screen. The coordinates are given in the **screen** coordinate system (see **coordinates** (p. 21) for information about this system).

Syntax:

```
set origin <x-origin>,<y-origin>
```

Output

By default, screens are displayed to the standard output. The **set output** command redirects the display to the specified file or device.

Syntax:

```
set output {"<filename>"}
show output
```

The filename must be enclosed in quotes. If the filename is omitted, any output file opened by a previous invocation of **set output** will be closed and new output will be sent to STDOUT. (If you give the command **set output "STDOUT"**, your output may be sent to a file named "STDOUT"! ["May be", not "will be", because some terminals, like **x11** or **wxt**, ignore **set output**.])

When both **set terminal** and **set output** are used together, it is safest to give **set terminal** first, because some terminals set a flag which is needed in some operating systems. This would be the case, for example, if the operating system needs a separate open command for binary files.

On platforms that support pipes, it may be useful to pipe terminal output. For instance,

```
set output "|lpr -Plaser filename"
set term png; set output "|display png:-"
```

On MSDOS machines, **set output "PRN"** will direct the output to the default printer. On VMS, output can be sent directly to any spooled device.

Parametric

The **set parametric** command changes the meaning of **plot** (**splot**) from normal functions to parametric functions. The command **unset parametric** restores the plotting style to normal, single-valued expression plotting.

Syntax:

```
set parametric
unset parametric
show parametric
```

For 2D plotting, a parametric function is determined by a pair of parametric functions operating on a parameter. An example of a 2D parametric function would be **plot sin(t),cos(t)**, which draws a circle (if the aspect ratio is set correctly — see **set size** (p. 151)). **gnuplot** will display an error message if both functions are not provided for a parametric **plot**.

For 3D plotting, the surface is described as $x=f(u,v)$, $y=g(u,v)$, $z=h(u,v)$. Therefore a triplet of functions is required. An example of a 3D parametric function would be **cos(u)*cos(v),cos(u)*sin(v),sin(u)**, which draws a sphere. **gnuplot** will display an error message if all three functions are not provided for a parametric **splot**.

The total set of possible plots is a superset of the simple $f(x)$ style plots, since the two functions can describe the x and y values to be computed separately. In fact, plots of the type $t,f(t)$ are equivalent to those produced with $f(x)$ because the x values are computed using the identity function. Similarly, 3D plots of the type $u,v,f(u,v)$ are equivalent to $f(x,y)$.

Note that the order the parametric functions are specified is x function, y function (and z function) and that each operates over the common parametric domain.

Also, the **set parametric** function implies a new range of values. Whereas the normal $f(x)$ and $f(x,y)$ style plotting assume an x range and y range (and z range), the parametric mode additionally specifies a $trange$, $urange$, and $vrangle$. These ranges may be set directly with **set trange**, **set urange**, and **set vrangle**, or by specifying the range on the **plot** or **splot** commands. Currently the default range for these parametric variables is [-5:5]. Setting the ranges to something more meaningful is expected.

Paxis

Syntax:

```
set paxis <axisno> {range <range-options> | tics <tic-options>}
show paxis <axisno> {range | tics}
```

The **set paxis** command is equivalent to the **set xrange** and **set xtics** commands except that it acts on one of the axes $p1$, $p2$, ... used in parallel axis plots. See **parallelaxes** (p. 55), **set xrange** (p. 167), and **set xtics** (p. 168). The normal options to the range and tics commands are accepted although not all options make sense for parallel axis plots.

Plot

The **show plot** command shows the current plotting command as it results from the last **plot** and/or **splot** and possible subsequent **replot** commands.

In addition, the **show plot add2history** command adds this current plot command into the **history**. It is useful if you have used **replot** to add more curves to the current plot and you want to edit the whole command now.

Pm3d

pm3d is an **splot** style for drawing palette-mapped 3d and 4d data as color/gray maps and surfaces. It uses an algorithm that allows plotting gridded as well as non-gridded data without preprocessing, even when the data scans do not have the same number of points.

Syntax (the options can be given in any order):

```
set pm3d {
    { at <position> }
    { interpolate <steps/points in scan, between scans> }
    { scansautomatic | scansforward | scansbackward | depthorder }
    { flush { begin | center | end } }
    { ftriangles | noftriangles }
    { clip1in | clip4in }
    { corners2color
      { mean|geomean|harmean|rms|median|min|max|c1|c2|c3|c4 }
    }
    { hidden3d {<linestyle>} | nohidden3d }
    { implicit | explicit }
    { map }
}
show pm3d
unset pm3d
```

A pm3d color surface is drawn if the splot command specifies **with pm3d**, if the data or function **style** is set to pm3d globally, or if the pm3d mode is **set pm3d implicit**. In the latter two cases, the pm3d surface is drawn in addition to the mesh produced by the style specified in the plot command. E.g.

```
splot 'fred.dat' with lines, 'lola.dat' with lines
```

would draw both a mesh of lines and a pm3d surface for each data set. If the option **explicit** is on (or **implicit** is off) only plots specified by the **with pm3d** attribute are plotted with a pm3d surface, e.g.:

```
splot 'fred.dat' with lines, 'lola.dat' with pm3d
```

would plot 'fred.dat' with lines (only) and 'lola.dat' with a pm3d surface.

On gnuplot start-up, the mode is **explicit**. For historical and compatibility reasons, the commands **set pm3d**; (i.e. no options) and **set pm3d at X ...** (i.e. **at** is the first option) change the mode to **implicit**. The command **set pm3d**; sets other options to their default state.

If you set the default data or function style to **pm3d**, e.g.:

```
set style data pm3d
```

then the options **implicit** and **explicit** have no effect.

Note that when plotting several plots, they are plotted in the order given on the command line. This can be of interest especially for filled surfaces which can overwrite and therefore hide part of earlier plots.

The pm3d coloring can be drawn at any or all of three different positions, **top**, **bottom**, or **surface**. See **pm3d position** (p. 142). The following command draws three color surfaces at different altitudes:

```
set border 4095
set pm3d at s
splot 10*x with pm3d at b, x*x-y*y, x*x+y*y with pm3d at t
```

See also help for **set palette** (p. 144), **set cbrange** (p. 175), **set colorbox** (p. 107), and definitely the demo file **demo/pm3d.dem**.

Algorithm

Let us first describe how a map/surface is drawn. The input data come from an evaluated function or from an **splot data file**. Each surface consists of a sequence of separate scans (isolines). The pm3d algorithm

fills the region between two neighbouring points in one scan with another two points in the next scan by a gray (or color) according to z-values (or according to an additional 'color' column, see help for **using** (p. 85)) of these 4 corners; by default the 4 corner values are averaged, but this can be changed by the option **corners2color**. In order to get a reasonable surface, the neighbouring scans should not cross and the number of points in the neighbouring scans should not differ too much; of course, the best plot is with scans having same number of points. There are no other requirements (e.g. the data need not be gridded). Another advantage is that the pm3d algorithm does not draw anything outside of the input (measured or calculated) region.

Surface coloring works with the following input data:

1. plot of function or of data file with one or three data columns: The gray/color scale is obtained by mapping the averaged (or **corners2color**) z-coordinate of the four corners of the above-specified quadrangle into the range [min_color_z,max_color_z] of **zrange** or **cbrange** providing a gray value in the range [0:1]. This value can be used directly as the gray for gray maps. The normalized gray value can be further mapped into a color — see **set palette** (p. 144) for the complete description.
2. plot of data file with two or four data columns: The gray/color value is obtained by using the last-column coordinate instead of the z-value, thus allowing the color and the z-coordinate be mutually independent. This can be used for 4d data drawing.

Other notes:

1. The term 'scan' referenced above is used more among physicists than the term 'iso_curve' referenced in gnuplot documentation and sources. You measure maps recorded one scan after another scan, that's why.
2. The 'gray' or 'color' scale is a linear mapping of a continuous variable onto a smoothly varying palette of colors. The mapping is shown in a rectangle next to the main plot. This documentation refers to this as a "colorbox", and refers to the indexing variable as lying on the colorbox axis. See **set colorbox** (p. 107), **set cbrange** (p. 175).

Position

Color surface can be drawn at the base or top (then it is a gray/color planar map) or at z-coordinates of surface points (gray/color surface). This is defined by the **at** option with a string of up to 6 combinations of **b**, **t** and **s**. For instance, **at b** plots at bottom only, **at st** plots firstly surface and then top map, while **at bstbst** will never be seriously used.

Colored quadrangles are plotted one after another. When plotting surfaces (**at s**), the later quadrangles overlap (overdraw) the previous ones. (Gnuplot is not virtual reality tool to calculate intersections of filled polygon meshes.) You may try to switch between **scansforward** and **scansbackward** to force the first scan of the data to be plotted first or last. The default is **scansautomatic** where gnuplot makes a guess about scans order. On the other hand, the **depthorder** option completely reorders the quadrangles. The rendering is performed after a depth sorting, which allows to visualize even complicated surfaces; see **pm3d depthorder** (p. 142) for more details.

Scanorder

By default the quadrangles making up a pm3d solid surface are rendered in the order they are encountered along the surface grid points. This order may be controlled by the options **scansautomatic|scansforward|scansbackward**. These scan options are not in general compatible with hidden-surface removal.

If two successive scans do not have same number of points, then it has to be decided whether to start taking points for quadrangles from the beginning of both scans (**flush begin**), from their ends (**flush end**) or to center them (**flush center**). Note, that **flush (center|end)** are incompatible with **scansautomatic**: if you specify **flush center** or **flush end** and **scansautomatic** is set, it is silently switched to **scansforward**.

If two subsequent scans do not have the same number of points, the option **ftriangles** specifies whether color triangles are drawn at the scan tail(s) where there are not enough points in either of the scans. This can be used to draw a smooth map boundary.

Gnuplot does not do true hidden surface removal for solid surfaces, but often it is sufficient to render the

component quadrangles in order from furthest to closest. This mode may be selected using the options

```
set pm3d depthorder hidden3d
```

The **depthorder** option orders the solid quadrangles; the **hidden3d** option similarly orders the bounding lines (if drawn). Note that the global option **set hidden3d** does not affect pm3d surfaces.

Clipping

Clipping with respect to x, y coordinates of quadrangles can be done in two ways. **clip1in**: all 4 points of each quadrangle must be defined and at least 1 point of the quadrangle must lie in the x and y ranges. **clip4in**: all 4 points of each quadrangle must lie in the x and y ranges.

Color assignment

3 columns of data (x,y,z):

The coloring setup as well as the color box drawing are determined by **set palette**. There can be only one palette for the current plot. Drawing of several surfaces with different palettes can be achieved by **multiplot** with fixed **origin** and **size**; don't forget to use **set palette maxcolors** when your terminal is running out of available colors.

There is a single gray/color value associated to each drawn pm3d quadrangle (it contains a solid color, not a gradient). The value is calculated from z-coordinates the four corners according to **corners2color** <option>. **4 columns of data (x,y,z,color)**:

If a fourth column of data is provided, it is normally interpreted as a separate palette-mapped gray value. The coloring of individual quadrangles works as above, except that the color value is distinct from the z value. As a separate coloring option, the fourth data column may provide instead an RGB color. See **rgbcolor variable** (p. 35). In this case the plotting command must be

```
splot ... using 1:2:3:4 with pm3d lc rgb variable
```

Another drawing algorithm, which would draw quadrangles around a given node by taking corners from averaged (x,y)-coordinates of its surrounding 4 nodes while using node's color, could be implemented in the future. This is already done for drawing images (2D grids) via **image** and **rgbimage** styles.

Notice that ranges of z-values and color-values for surfaces are adjustable independently by **set zrange**, **set cbrange**, as well as **set log** for z or cb. Maps can be adjusted by the cb-axis only; see also **set view map** (p. 164) and **set colorbox** (p. 107).

Corners2color

The color of each quadrangle in a pm3d surface is assigned based on the color values of its four bounding vertices. The options 'mean' (default), 'geomean', 'harmean', 'rms', and 'median' produce various kinds of surface color smoothing, while options 'min' and 'max' choose minimal or maximal value, respectively. This may not be desired for pixel images or for maps with sharp and intense peaks, in which case the options 'c1', 'c2', 'c3' or 'c4' can be used instead to assign the quadrangle color based on the z-coordinate of only one corner. Some experimentation may be needed to determine which corner corresponds to 'c1', as the orientation depends on the drawing direction. Because the pm3d algorithm does not extend the colored surface outside the range of the input data points, the 'c<j>' coloring options will result in pixels along two edges of the grid not contributing to the color of any quadrangle. For example, applying the pm3d algorithm to the 4x4 grid of data points in script **demo/pm3d.dem** (please have a look) produces only (4-1)x(4-1)=9 colored rectangles.

Border

The option **set pm3d border {line-properties}** draws bounding lines around each quadrangle as it is rendered. Normally this is used in conjunction with the **depthorder** option to approximate hidden line removal. Note that the global option **set hidden3d** has no effect on pm3d plots. Default line properties

(color, width) optionally follow the keyword **border**. These defaults can be overridden later in an `plot` command.

Example of recommended usage:

```
set pm3d at s depthorder border lw 0.2 lt black
unset hidden3d
unset surf
splot x*x+y*y linecolor rgb "blue"    # otherwise it would be black
```

NB: The deprecated option `set pm3d hidden3d N` is still accepted. It is equivalent to `set pm3d border ls N`.

Interpolate

The option **interpolate m,n** will interpolate between grid points to generate a finer mesh. For data files, this smooths the color surface and enhances the contrast of spikes in the surface. When working with functions, interpolation makes little sense. It would usually make more sense to increase **samples** and **isosamples**.

For positive *m* and *n*, each quadrangle or triangle is interpolated *m*-times and *n*-times in the respective direction. For negative *m* and *n*, the interpolation frequency is chosen so that there will be at least $|m|$ and $|n|$ points drawn; you can consider this as a special gridding function.

Note: **interpolate 0,0**, will automatically choose an optimal number of interpolated surface points.

Note: Currently color interpolation is always linear, even if `corners2color` is set to a nonlinear scheme such as the geometric mean.

Deprecated options

There used to be an option `{transparent|solid}` to this command. Now you get the same effect from `set grid {front|layerdefault}`, respectively.

The old command `set pm3d map` is equivalent to `set pm3d at b; set view map scale 1.0; set style data pm3d; set style func pm3d;`

Palette

Palette is a color storage for use by **pm3d**, filled color contours or polygons, color histograms, color gradient background, and whatever it is or it will be implemented... Here it stands for a palette of smooth "continuous" colors or grays, but let's call it just a palette.

Color palettes require terminal entries for filled color polygons and palettes of smooth colors, are currently available for terminals listed in help for `set pm3d`. The range of color values are adjustable independently by `set cbrange` and `set log cb`. The whole color palette is visualized in the **colorbox**.

Syntax:

```
set palette
set palette {
    { gray | color }
    { gamma <gamma> }
    { rgbformulae <r>,<g>,<b>
      | defined { ( <gray1> <color1> {, <grayN> <colorN>}... ) }
      | file '<filename>' {datafile-modifiers}
      | functions <R>,<G>,<B>
    }
    { cubehelix {start <val>} {cycles <val>} {saturation <val>} }
    { model { RGB | HSV | CMY | YIQ | XYZ } }
    { positive | negative }
    { nops_allcF | ps_allcF }
    { maxcolors <maxcolors> }
```



```

    }
show palette
show palette palette <n> {{float | int}}
show palette gradient
show palette fit2rgbformulae
show palette rgbformulae
show colnames

```

set palette (i.e. without options) sets up the default values. Otherwise, the options can be given in any order. **show palette** shows the current palette properties.

show palette gradient displays the gradient defining the palette (if appropriate). **show palette rgbformulae** prints the available fixed gray \rightarrow color transformation formulae. **show colnames** prints the known color names.

show palette palette <n> prints to the screen or to the file given by **set print** a table of RGB triplets calculated for the current palette settings and a palette having <n> discrete colors. The default wide table can be limited to 3 columns of r,g,b float values [0..1] or integer values [0..255] by options float or int, respectively. This way, the current gnuplot color palette can be loaded into other imaging applications, for example Octave. Alternatively, the **test palette** command will plot the R,G,B profiles for the current palette and leave the profile values in a datablock \$PALETTE.

The following options determine the coloring properties.

Figure using this palette can be **gray** or **color**. For instance, in **pm3d** color surfaces the gray of each small spot is obtained by mapping the averaged z-coordinate of the 4 corners of surface quadrangles into the range [min_z,max_z] providing range of grays [0:1]. This value can be used directly as the gray for gray maps. The color map requires a transformation gray \rightarrow (R,G,B), i.e. a mapping [0:1] \rightarrow ([0:1],[0:1],[0:1]).

Basically two different types of mappings can be used: Analytic formulae to convert gray to color, or discrete mapping tables which are interpolated. **palette rgbformulae** and **palette functions** use analytic formulae whereas **palette defined** and **palette file** use interpolated tables. **palette rgbformulae** reduces the size of postscript output to a minimum.

The command **show palette fit2rgbformulae** finds the best matching **set palette rgbformulae** for the current **set palette**. Naturally, it makes sense to use it for non-rgbformulae palettes. This command can be found useful mainly for external programs using the same rgbformulae definition of palettes as gnuplot, like zimg (

<http://zimg.sourceforge.net>

).

set palette gray switches to a gray only palette. **set palette rgbformulae**, **set palette defined**, **set palette file** and **set palette functions** switch to a color mapping. **set palette color** is an easy way to switch back from the gray palette to the last color mapping.

Automatic gamma correction via **set palette gamma <gamma>** can be done for gray maps (**set palette gray**) and for the **cubehelix** color palette schemes. Gamma = 1 produces a linear ramp of intensity. See **test palette** (p. 182).

Many terminals support only discrete number of colors (e.g. 256 colors in gif). After the default gnuplot linetype colors are allocated, the rest of the available colors are by default reserved for pm3d. Thus a multiplot using multiple palettes could fail because the first palette has used all the available color positions. You can mitigate this limitation by using **set palette maxcolors <N>** with a reasonably small value of N. This option causes N discrete colors to be selected from a continuous palette sampled at equally spaced intervals. If you want unequal spacing of N discrete colors, use **set palette defined** instead of a single continuous palette.

RGB color space might not be the most useful color space to work in. For that reason you may change the color space with **model** to one of **RGB**, **HSV**, **CMY**, **YIQ** and **XYZ**. Using color names for **set palette defined** tables and a color space other than RGB will result in funny colors. All explanation have been written for RGB color space, so please note, that **R** can be **H**, **C**, **Y**, or **X**, depending on the actual color space (**G** and **B** accordingly).

All values for all color spaces are limited to [0,1].

RGB stands for Red, Green and Blue; CMY stands for Cyan, Magenta and Yellow; HSV stands for Hue, Saturation, and Value; YIQ is the color model used by the U.S. Commercial Color Television Broadcasting, it is basically an RGB recoding with downward compatibility for black and white television; XYZ are the three primary colors of the color model defined by the 'Commission Internationale de l'Eclairage' (CIE). For more information on color models see:

http://en.wikipedia.org/wiki/Color_space

Rgbformulae

For **rgbformulae** three suitable mapping functions have to be chosen. This is done via **rgbformulae** `<r>,<g>,`. The available mapping functions are listed by **show palette rgbformulae**. Default is **7,5,15**, some other examples are **3,11,6**, **21,23,3** or **3,23,21**. Negative numbers, like **3,-11,-6**, mean inverted color (i.e. 1-gray passed into the formula, see also **positive** (p. 146) and **negative** (p. 146) options below).

Some nice schemes in RGB color space

```
7,5,15    ... traditional pm3d (black-blue-red-yellow)
3,11,6    ... green-red-violet
23,28,3   ... ocean (green-blue-white); try also all other permutations
21,22,23  ... hot (black-red-yellow-white)
30,31,32  ... color printable on gray (black-blue-violet-yellow-white)
33,13,10  ... rainbow (blue-green-yellow-red)
34,35,36  ... AFM hot (black-red-yellow-white)
```

A full color palette in HSV color space

```
3,2,2     ... red-yellow-green-cyan-blue-magenta-red
```

Please note that even if called **rgbformulae** the formulas might actually determine the `<H>,<S>,<V>` or `<X>,<Y>,<Z>` or ... color components as usual.

Use **positive** and **negative** to invert the figure colors.

Note that it is possible to find a set of the best matching rgbformulae for any other color scheme by the command

```
show palette fit2rgbformulae
```

Defined

Gray-to-rgb mapping can be manually set by use of **palette defined**: A color gradient is defined and used to give the rgb values. Such a gradient is a piecewise linear mapping from gray values in [0,1] to the RGB space [0,1]x[0,1]x[0,1]. You must specify the gray values and the corresponding RGB values between which linear interpolation will be done.

Syntax:

```
set palette defined { ( <gray1> <color1> {, <grayN> <colorN>}... ) }
```

`<grayX>` are gray values which are mapped to [0,1] and `<colorX>` are the corresponding rgb colors. The color can be specified in three different ways:

```
<color> := { <r> <g> <b> | '<color-name>' | '#rrggbb' }
```

Either by three numbers (each in [0,1]) for red, green and blue, separated by whitespace, or the name of the color in quotes or X style color specifiers also in quotes. You may freely mix the three types in a gradient definition, but the named color "red" will be something strange if RGB is not selected as color space. Use **show colornames** for a list of known color names.

Please note, that even if written as `<r>`, this might actually be the `<H>` component in HSV color space or `<X>` in CIE-XYZ space, or ... depending on the selected color model.

The `<gray>` values have to form an ascending sequence of real numbers; the sequence will be automatically rescaled to [0,1].

set palette defined (without a gradient definition in braces) switches to RGB color space and uses a preset full-spectrum color gradient. Use **show palette gradient** to display the gradient.

Examples:

To produce a gray palette (useless but instructive) use:

```
set palette model RGB
set palette defined ( 0 "black", 1 "white" )
```

To produce a blue yellow red palette use (all equivalent):

```
set palette defined ( 0 "blue", 1 "yellow", 2 "red" )
set palette defined ( 0 0 0 1, 1 1 1 0, 2 1 0 0 )
set palette defined ( 0 "#0000ff", 1 "#ffff00", 2 "#ff0000" )
```

To produce some rainbow-like palette use:

```
set palette defined ( 0 "blue", 3 "green", 6 "yellow", 10 "red" )
```

Full color spectrum within HSV color space:

```
set palette model HSV
set palette defined ( 0 0 1 1, 1 1 1 1 )
set palette defined ( 0 0 1 0, 1 0 1 1, 6 0.8333 1 1, 7 0.8333 0 1)
```

Approximate the default palette used by MATLAB:

```
set pal defined (1 '#00008f', 8 '#0000ff', 24 '#00ffff', \
                40 '#ffff00', 56 '#ff0000', 64 '#800000')
```

To produce a palette with only a few, equally-spaced colors:

```
set palette model RGB maxcolors 4
set palette defined ( 0 "yellow", 1 "red" )
```

'Traffic light' palette (non-smooth color jumps at gray = 1/3 and 2/3).

```
set palette model RGB
set palette defined (0 "dark-green", 1 "green", \
                    1 "yellow",      2 "dark-yellow", \
                    2 "red",          3 "dark-red" )
```

Functions

Use **set palette functions** **<Rexpr>**, **<Gexpr>**, **<Bexpr>** to define three formulae for the R(gray), G(gray) and B(gray) mapping. The three formulae may depend on the variable **gray** which will take values in [0,1] and should also produce values in [0,1]. Please note that **<Rexpr>** might be a formula for the H-value if HSV color space has been chosen (same for all other formulae and color spaces).

Examples:

To produce a full color palette use:

```
set palette model HSV functions gray, 1, 1
```

A nice black to gold palette:

```
set palette model XYZ functions gray**0.35, gray**0.5, gray**0.8
```

A gamma-corrected black and white palette

```
gamma = 2.2
color(gray) = gray**(1./gamma)
set palette model RGB functions color(gray), color(gray), color(gray)
```

Gray

set palette gray switches to a grayscale palette shading from 0.0 = black to 1.0 = white. **set palette color** is an easy way to switch back from the gray palette to the last color mapping.

Cubehelix

The "cubehelix" option defines a family of palettes in which color (hue) varies along the standard color wheel while at the same time the net intensity increases monotonically as the gray value goes from 0 to 1.

D A Green (2011) <http://arxiv.org/abs/1108.5083>

start defines the starting point along the color wheel in radians. **cycles** defines how many color wheel cycles span the palette range. Larger values of **saturation** produce more saturated color; saturation > 1 may lead to clipping of the individual RGB components and to intensity becoming non-monotonic. The palette is also affected by **set palette gamma**. The default values are

```
set palette cubehelix start 0.5 cycles -1.5 saturation 1
set palette gamma 1.5
```

File

set palette file is basically a **set palette defined** (<gradient>) where <gradient> is read from a datafile. Either 4 columns (gray,R,G,B) or just three columns (R,G,B) have to be selected via the **using** data file modifier. In the three column case, the line number will be used as gray. The gray range is automatically rescaled to [0,1]. The file is read as a normal data file, so all datafile modifiers can be used. Please note, that **R** might actually be e.g. **H** if HSV color space is selected.

As usual <filename> may be '-' which means that the data follow the command inline and are terminated by a single **e** on a line of its own.

Use **show palette gradient** to display the gradient.

Examples:

Read in a palette of RGB triples each in range [0,255]:

```
set palette file 'some-palette' using ($1/255):($2/255):($3/255)
```

Equidistant rainbow (blue-green-yellow-red) palette:

```
set palette model RGB file "-"
0 0 1
0 1 0
1 1 0
1 0 0
e
```

Binary palette files are supported as well, see **binary general** (p. 75). Example: put 64 triplets of R,G,B doubles into file palette.bin and load it by

```
set palette file "palette.bin" binary record=64 using 1:2:3
```

Gamma correction

For gray mappings gamma correction can be turned on by **set palette gamma** <gamma>. <gamma> defaults to 1.5 which is quite suitable for most terminals.

The gamma correction is applied to the cubehelix color palette family, but not to other palette coloring schemes. However, you may easily implement gamma correction for explicit color functions.

Example:

```
set palette model RGB
set palette functions gray**0.64, gray**0.67, gray**0.70
```

To use gamma correction with interpolated gradients specify intermediate gray values with appropriate colors. Instead of

```
set palette defined ( 0 0 0 0, 1 1 1 1 )
```

use e.g.

```
set palette defined ( 0 0 0 0, 0.5 .73 .73 .73, 1 1 1 1 )
```

or even more intermediate points until the linear interpolation fits the "gamma corrected" interpolation well enough.

Postscript

In order to reduce the size of postscript files, the gray value and not all three calculated r,g,b values are written to the file. Therefore the analytical formulae are coded directly in the postscript language as a header just before the pm3d drawing, see /g and /cF definitions. Usually, it makes sense to write therein definitions of only the 3 formulae used. But for multiplot or any other reason you may want to manually edit the transformations directly in the postscript file. This is the default option **nops_allcF**. Using the option **ps_allcF** writes postscript definitions of all formulae. This you may find interesting if you want to edit the postscript file in order to have different palettes for different surfaces in one graph. Well, you can achieve this functionality by **multiplot** with fixed **origin** and **size**.

If you are writing a pm3d surface to a postscript file, it may be possible to reduce the file size by up to 50% by the enclosed awk script **pm3dCompress.awk**. If the data lies on a rectangular grid, even greater compression may be possible using the script **pm3dConvertToImage.awk**. Usage:

```
awk -f pm3dCompress.awk thefile.ps >smallerfile.ps
awk -f pm3dConvertToImage.awk thefile.ps >smallerfile.ps
```

Pointintervalbox

The **pointinterval** property of line types is used in plot style **linespoints**. A negative value of pointinterval, e.g. -N, means that point symbols are drawn only for every Nth point, and that a box (actually circle) behind each point symbol is blanked out by filling with the background color. The command **set pointintervalbox** controls the radius of this blanked-out region. It is a multiplier for the default radius, which is equal to the point size.

Pointsize

The **set pointsize** command scales the size of the points used in plots.

Syntax:

```
set pointsize <multiplier>
show pointsize
```

The default is a multiplier of 1.0. Larger pointsizes may be useful to make points more visible in bitmapped graphics.

The pointsize of a single plot may be changed on the **plot** command. See **plot with** (p. 93) for details.

Please note that the pointsize setting is not supported by all terminal types.

Polar

The **set polar** command changes the meaning of the plot from rectangular coordinates to polar coordinates.

Syntax:

```
set polar
unset polar
show polar
```

In polar coordinates, the dummy variable (t) is an angle. The default range of t is [0:2*pi], or, if degree units have been selected, to [0:360] (see **set angles** (p. 98)).

The command **unset polar** changes the meaning of the plot back to the default rectangular coordinate system.

The **set polar** command is not supported for **splots**. See the **set mapping** (p. 130) command for similar functionality for **splot** (p. 176)s.

While in polar coordinates the meaning of an expression in t is really $r = f(t)$, where t is an angle of rotation. The **trange** controls the domain (the angle) of the function. The r , x and y ranges control the extent of the graph in the x and y directions. Each of these ranges, as well as the range, may be autoscaled or set explicitly. For details, see **set rrange** (p. 151) and **set xrange** (p. 167).

Example:

```
set polar
plot t*sin(t)
set trange [-2*pi:2*pi]
set rrange [0:3]
plot t*sin(t)
```

The first **plot** uses the default polar angular domain of 0 to 2π . The radius and the size of the graph are scaled automatically. The second **plot** expands the domain, and restricts the size of the graph to the area within 3 units of the origin. This has the effect of limiting x and y to $[-3:3]$.

You may want to **set size square** to have **gnuplot** try to make the aspect ratio equal to unity, so that circles look circular. See also

[polar demos \(polar.dem\)](#)

and

[polar data plot \(poldat.dem\)](#).

Print

The **set print** command redirects the output of the **print** command to a file.

Syntax:

```
set print
set print "-"
set print "<filename>" [append]
set print "|<shell_command>"
set print $datablock [append]
```

set print with no parameters restores output to $\langle \text{STDERR} \rangle$. The $\langle \text{filename} \rangle$ `"-"` means $\langle \text{STDOUT} \rangle$. The **append** flag causes the file to be opened in append mode. A $\langle \text{filename} \rangle$ starting with `"|"` is opened as a pipe to the $\langle \text{shell_command} \rangle$ on platforms that support piping.

The destination for **print** commands can also be a named data block. Data block names start with `'$'`, see also **inline data** (p. 81).

Psdire

The **set psdire** $\langle \text{directory} \rangle$ command controls the search path used by the postscript terminal to find prologue.ps and character encoding files. You can use this mechanism to switch between different sets of locally-customized prolog files. The search order is

- 1) The directory specified by 'set psdire', if any
- 2) The directory specified by environmental variable GNUPLT_PS_DIR
- 3) A built-in header or one from the default system directory
- 4) Directories set by 'set loadpath'

Raxis

The commands **set raxis** and **unset raxis** toggle whether the polar axis is drawn separately from grid lines and the x axis. If the minimum of the current rrange is non-zero (and not autoscaled), then a white circle is drawn at the center of the polar plot to indicate that the plot lines and axes do not reach 0. The axis line is drawn using the same line type as the plot border. See **polar** (p. 149), **rrange** (p. 151), **rtics** (p. 151), **set grid** (p. 119).

Rmargin

The command **set rmargin** sets the size of the right margin. Please see **set margin** (p. 131) for details.

Rrange

The **set rrange** command sets the range of the radial coordinate for a graph in polar mode. This has the effect of setting both xrange and yrange as well. The resulting xrange and yrange are both $[-(rmax-rmin) : +(rmax-rmin)]$. However if you later change the x or y range, for example by zooming, this does not change rrange, so data points continue to be clipped against rrange. Autoscaling of rmin always results in $rmin = 0$. Note: Setting a negative value for rmin may produce unexpected results.

Rtics

The **set rtics** command places ticks along the polar axis. These will only be shown in polar plot mode. The ticks and labels are drawn to the right of the origin. The **mirror** keyword causes them to be drawn also to the left of the origin. See **polar** (p. 149), **set xtics** (p. 168), and **set mxtics** (p. 135) for discussion of keywords.

Samples

The default sampling rate of functions, or for interpolating data, may be changed by the **set samples** command. To change the sampling range for a particular plot, see **plot sampling** (p. 91).

Syntax:

```
set samples <samples_1> {,<samples_2>}
show samples
```

By default, sampling is set to 100 points. A higher sampling rate will produce more accurate plots, but will take longer. This parameter has no effect on data file plotting unless one of the interpolation/approximation options is used. See **plot smooth** (p. 82) re 2D data and **set cntrparam** (p. 106) and **set dgrid3d** (p. 112) re 3D data.

When a 2D graph is being done, only the value of <samples.1> is relevant.

When a surface plot is being done without the removal of hidden lines, the value of samples specifies the number of samples that are to be evaluated for the isolines. Each iso-v line will have <sample.1> samples and each iso-u line will have <sample.2> samples. If you only specify <samples.1>, <samples.2> will be set to the same value as <samples.1>. See also **set isosamples** (p. 122).

Size

Syntax:

```
set size {{no}square | ratio <r> | noratio} {<xscale>,<yscale>}
show size
```

The <xscale> and <yscale> values are scale factors for the size of the plot, which includes the graph, labels, and margins.

Important note:

In earlier versions of **gnuplot**, some terminal types used the values from 'set size' to control also the size of the output canvas; others did not. Almost all terminals now follow the following convention:

set term <terminal.type> **size** <XX>, <YY> controls the size of the output file, or **canvas**. Please see individual terminal documentation for allowed values of the size parameters. By default, the plot will fill this canvas.

set size <XX>, <YY> scales the plot itself relative to the size of the canvas. Scale values less than 1 will cause the plot to not fill the entire canvas. Scale values larger than 1 will cause only a portion of the plot to fit on the canvas. Please be aware that setting scale values larger than 1 may cause problems on some terminal types.

ratio causes **gnuplot** to try to create a graph with an aspect ratio of <r> (the ratio of the y-axis length to the x-axis length) within the portion of the plot specified by <xscale> and <yscale>.

The meaning of a negative value for <r> is different. If <r>=-1, **gnuplot** tries to set the scales so that the unit has the same length on both the x and y axes. This is equivalent to **set view equal xy**. See **set view equal** (p. 164). If <r>=-2, the unit on y has twice the length of the unit on x, and so on.

The success of **gnuplot** in producing the requested aspect ratio depends on the terminal selected. The graph area will be the largest rectangle of aspect ratio <r> that will fit into the specified portion of the output (leaving adequate margins, of course).

set size square is a synonym for **set size ratio 1**.

Both **noratio** and **nosquare** return the graph to the default aspect ratio of the terminal, but do not return <xscale> or <yscale> to their default values (1.0).

ratio and **square** have no effect on 3D plots, but do affect 3D projections created using **set view map**. See also **set view equal** (p. 164), which forces the x and y axes of a 3D onto the same scale.

Examples:

To set the size so that the plot fills the available canvas:

```
set size 1,1
```

To make the graph half size and square use:

```
set size square 0.5,0.5
```

To make the graph twice as high as wide use:

```
set size ratio 2
```

Style

Default plotting styles are chosen with the **set style data** and **set style function** commands. See **plot with** (p. 93) for information about how to override the default plotting style for individual functions and data sets. See **plotting styles** (p. 44) for a complete list of styles.

Syntax:

```
set style function <style>
set style data <style>
show style function
show style data
```

Default styles for specific plotting elements may also be set.

Syntax:

```
set style arrow <n> <arrowstyle>
set style boxplot <boxplot style options>
set style circle radius <size> {clip|noclip}
set style ellipse size <size> units {xy|xx|yy} {clip|noclip}
set style fill <fillstyle>
```



```

set style histogram <histogram style options>
set style line <n> <linestyle>
set style rectangle <object options> <linestyle> <fillstyle>
set style textbox {opaque|transparent} {{no}border}

```

Set style arrow

Each terminal has a default set of arrow and point types, which can be seen by using the command **test**. **set style arrow** defines a set of arrow types and widths and point types and sizes so that you can refer to them later by an index instead of repeating all the information at each invocation.

Syntax:

```

set style arrow <index> default
set style arrow <index> {nohead | head | heads}
                        {size <length>,<angle>{,<backangle>}} {fixed}}
                        {filled | empty | nofilled | noborder}
                        {front | back}
                        { {linestyle | ls <line_style>}
                          | {linetype | lt <line_type>}
                          {linewidth | lw <line_width> }
unset style arrow
show style arrow

```

<index> is an integer that identifies the arrowstyle.

If **default** is given all arrow style parameters are set to their default values.

If the linestyle <index> already exists, only the given parameters are changed while all others are preserved. If not, all undefined values are set to the default values.

Specifying **nohead** produces arrows drawn without a head — a line segment. This gives you yet another way to draw a line segment on the plot. By default, arrows have one head. Specifying **heads** draws arrow heads on both ends of the line.

Head size can be controlled by **size <length>,<angle>** or **size <length>,<angle>,<backangle>**, where <length> defines length of each branch of the arrow head and <angle> the angle (in degrees) they make with the arrow. <Length> is in x-axis units; this can be changed by **first**, **second**, **graph**, **screen**, or **character** before the <length>; see **coordinates** (p. 21) for details.

By default the size of the arrow head is reduced for very short arrows. This can be disabled using the **fixed** keyword after the **size** command.

<backangle> is the angle (in degrees) the back branches make with the arrow (in the same direction as <angle>). It is ignored if the style is **nofilled**.

Specifying **filled** produces filled arrow heads with a border line around the arrow head. Specifying **noborder** produces filled arrow heads with no border. In this case the tip of the arrow head lies exactly on the endpoint of the vector and the arrow head is slightly smaller overall. Dashed arrows should always use **noborder**, since a dashed border is ugly. Not all terminals support filled arrow heads.

The line style may be selected from a user-defined list of line styles (see **set style line** (p. 156)) or may be defined here by providing values for <line_type> (an index from the default list of styles) and/or <line_width> (which is a multiplier for the default width).

Note, however, that if a user-defined line style has been selected, its properties (type and width) cannot be altered merely by issuing another **set style arrow** command with the appropriate index and **lt** or **lw**.

If **front** is given, the arrows are written on top of the graphed data. If **back** is given (the default), the arrow is written underneath the graphed data. Using **front** will prevent a arrow from being obscured by dense data.

Examples:

To draw an arrow without an arrow head and double width, use:

```
set style arrow 1 nohead lw 2
```

```
set arrow arrowstyle 1
```

See also **set arrow** (p. 99) for further examples.

Boxplot

The **set style boxplot** command allows you to change the layout of plots created using the **boxplot** plot style.

Syntax:

```
set style boxplot {range <r> | fraction <f>}
                  {{no}outliers} {pointtype <p>}
                  {candlesticks | financebars}
                  {separation <x>}
                  {labels off | auto | x | x2}
                  {sorted | unsorted}
```

The box in the boxplot always spans the range of values from the first quartile to the third quartile of the data points. The limit of the whiskers that extend from the box can be controlled in two different ways. By default the whiskers extend from each end of the box for a range equal to 1.5 times the interquartile range (i.e. the vertical height of the box proper). Each whisker is truncated back toward the median so that it terminates at a y value belonging to some point in the data set. Since there may be no point whose value is exactly 1.5 times the interquartile distance, the whisker may be shorter than its nominal range. This default corresponds to

```
set style boxplot range 1.5
```

Alternatively, you can specify the fraction of the total number of points that the whiskers should span. In this case the range is extended symmetrically from the median value until it encompasses the requested fraction of the data set. Here again each whisker is constrained to end at a point in the data set. To span 95% of the points in the set

```
set style boxplot fraction 0.95
```

Any points that lie outside the range of the whiskers are considered outliers. By default these are drawn as individual circles (pointtype 7). The option **nooutliers** disables this.

By default boxplots are drawn in a style similar to candlesticks, but you have the option of using instead a style similar to finance bars.

If the using specification for a boxplot contains a fourth column, the values in that column will be interpreted as the discrete levels of a factor variable. In this case more than one boxplots may be drawn, as many as the number of levels of the factor variable. These boxplots will be drawn next to each other, the distance between them is 1.0 by default (in x-axis units). This distance can be changed by the option **separation**.

The **labels** option governs how and where these boxplots (each representing a part of the dataset) are labeled. By default the value of the factor is put as a tick label on the horizontal axis – x or x2, depending on which one is used for the plot itself. This setting corresponds to option **labels auto**. The labels can be forced to use either of the x or x2 axes – options **labels x** and **labels x2**, respectively –, or they can be turned off altogether with the option **labels off**.

By default the boxplots corresponding to different levels of the factor variable are not sorted; they will be drawn in the same order the levels are encountered in the data file. This behavior corresponds to the **unsorted** option. If the **sorted** option is active, the levels are first sorted alphabetically, and the boxplots are drawn in the sorted order.

The **separation**, **labels**, **sorted** and **unsorted** option only have an effect if a fourth column is given the plot specification.

See **boxplot** (p. 45), **candlesticks** (p. 46), **financebars** (p. 49).

Set style data

The **set style data** command changes the default plotting style for data plots.

Syntax:

```
set style data <plotting-style>
show style data
```

See **plotting styles** (p. 44) for the choices. **show style data** shows the current default data plotting style.

Set style fill

The **set style fill** command is used to set the default style of the plot elements in plots with boxes, histograms, candlesticks and filledcurves. This default can be superseded by fillstyles attached to individual plots. See also 'set style rectangle'.

Syntax:

```
set style fill {empty
               | {transparent} solid {<density>}
               | {transparent} pattern {<n>}}
{border {lt} {lc <colourspec>} | noborder}
```

The default fillstyle is **empty**.

The **solid** option causes filling with a solid color, if the terminal supports that. The <density> parameter specifies the intensity of the fill color. At a <density> of 0.0, the box is empty, at <density> of 1.0, the inner area is of the same color as the current linetype. Some terminal types can vary the density continuously; others implement only a few levels of partial fill. If no <density> parameter is given, it defaults to 1.

The **pattern** option causes filling to be done with a fill pattern supplied by the terminal driver. The kind and number of available fill patterns depend on the terminal driver. If multiple datasets using filled boxes are plotted, the pattern cycles through all available pattern types, starting from pattern <n>, much as the line type cycles for multiple line plots.

The **empty** option causes filled boxes not to be filled. This is the default.

By default, **border**, the box is bounded by a solid line of the current linetype. **border <colourspec>** allows you to change the color of the border. **noborder** specifies that no bounding lines are drawn.

Set style fill transparent Some terminals support the attribute **transparent** for filled areas. In the case of transparent solid fill areas, the **density** parameter is interpreted as an alpha value; that is, density 0 is fully transparent, density 1 is fully opaque. In the case of transparent pattern fill, the background of the pattern is either fully transparent or fully opaque.

Terminal	solid	pattern	pm3d
gif	no	yes	no
jpeg	yes	no	yes
pdf	yes	yes	yes
png	TrueColor	index	yes
post	no	yes	no
svg	yes	no	yes
win	yes	yes	yes
wxt	yes	yes	yes
x11	no	yes	no

Note that there may be additional limitations on the creation or viewing of graphs containing transparent fill areas. For example, the png terminal can only use transparent fill if the "truecolor" option is set. Some pdf viewers may not correctly display the fill areas even if they are correctly described in the pdf file. Ghostscript/gv does not correctly display pattern-fill areas even though actual PostScript printers generally have no problem.

Set style function

The **set style function** command changes the default plotting style for function plots (e.g. lines, points, filledcurves). See **plotting styles** (p. 44).

Syntax:

```
set style function <plotting-style>
show style function
```

Set style increment

Note: This command has been deprecated. Instead please use the newer command **set linetype**, which redefines the linetypes themselves rather than searching for a suitable temporary line style to substitute. See **set linetype** (p. 128)

Syntax:

```
set style increment {default|userstyles}
show style increment
```

By default, successive plots within the same graph will use successive linetypes from the default set for the current terminal type. However, choosing **set style increment user** allows you to step through the user-defined line styles rather than through the default linetypes.

Example:

```
set style line 1 lw 2 lc rgb "gold"
set style line 2 lw 2 lc rgb "purple"
set style line 4 lw 1 lc rgb "sea-green"
set style increment user
```

```
plot f1(x), f2(x), f3(x), f4(x)
```

should plot functions f1, f2, f4 in your 3 newly defined line styles. If a user-defined line style is not found then the corresponding default linetype is used instead. E.g. in the example above, f3(x) will be plotted using the default linetype 3.

Set style line

Each terminal has a default set of line and point types, which can be seen by using the command **test**. **set style line** defines a set of line types and widths and point types and sizes so that you can refer to them later by an index instead of repeating all the information at each invocation.

Syntax:

```
set style line <index> default
set style line <index> {{linetype | lt} <line_type> | <colourspec>}
                        {{linecolor | lc} <colourspec>}
                        {{linewidth | lw} <line_width>}
                        {{pointtype | pt} <point_type>}
                        {{pointsize | ps} <point_size>}
                        {{pointinterval | pi} <interval>}
                        {palette}

unset style line
show style line
```

default sets all line style parameters to those of the linetype with that same index.

If the linestyle <index> already exists, only the given parameters are changed while all others are preserved. If not, all undefined values are set to the default values.

Line styles created by this mechanism do not replace the default linetype styles; both may be used. Line styles are temporary. They are lost whenever you execute a **reset** command. To redefine the linetype itself, please see **set linetype** (p. 128).

The line and point types default to the index value. The exact symbol that is drawn for that index value may vary from one terminal type to another.

The line width and point size are multipliers for the current terminal's default width and size (but note that `<point_size>` here is unaffected by the multiplier given by the command `set pointsize`).

The `pointinterval` controls the spacing between points in a plot drawn with style `linespoints`. The default is 0 (every point is drawn). For example, `set style line N pi 3` defines a linestyle that uses pointtype N, pointsize and linewidth equal to the current defaults for the terminal, and will draw every 3rd point in plots using `with linespoints`. A negative value for the interval is treated the same as a positive value, except that some terminals will try to interrupt the line where it passes through the point symbol.

Not all terminals support the `linewidth` and `pointsize` features; if not supported, the option will be ignored.

Terminal-independent colors may be assigned using either `linecolor <colourspec>` or `linetype <colourspec>`, abbreviated `lc` or `lt`. This requires giving a RGB color triple, a known palette color name, a fractional index into the current palette, or a constant value from the current mapping of the palette onto cbrange. See `colors` (p. 33), `colourspec` (p. 33), `set palette` (p. 144), `colnames` (p. 108), `cbrange` (p. 175).

`set style line <n> linetype <lt>` will set both a terminal-dependent dot/dash pattern and color. The command `set style line <n> linecolor <colourspec>` or `set style line <n> linetype <colourspec>` will set a new line color while leaving the existing dot-dash pattern unchanged.

In 3d mode (`splot` command), the special keyword `palette` is allowed as a shorthand for "linetype palette z". The color value corresponds to the z-value (elevation) of the splot, and varies smoothly along a line or surface.

Examples: Suppose that the default lines for indices 1, 2, and 3 are red, green, and blue, respectively, and the default point shapes for the same indices are a square, a cross, and a triangle, respectively. Then

```
set style line 1 lt 2 lw 2 pt 3 ps 0.5
```

defines a new linestyle that is green and twice the default width and a new pointstyle that is a half-sized triangle. The commands

```
set style function lines
plot f(x) lt 3, g(x) ls 1
```

will create a plot of f(x) using the default blue line and a plot of g(x) using the user-defined wide green line. Similarly the commands

```
set style function linespoints
plot p(x) lt 1 pt 3, q(x) ls 1
```

will create a plot of p(x) using the default triangles connected by a red line and q(x) using small triangles connected by a green line.

```
splot sin(sqrt(x*x+y*y))/sqrt(x*x+y*y) w l pal
```

creates a surface plot using smooth colors according to `palette`. Note, that this works only on some terminals. See also `set palette` (p. 144), `set pm3d` (p. 141).

```
set style line 10 linetype 1 linecolor rgb "cyan"
```

will assign linestyle 10 to be a solid cyan line on any terminal that supports rgb colors.

Set style circle

Syntax:

```
set style circle {radius {graph|screen} <R>}
                {{no}wedge}
                {clip|noclip}
```

This command sets the default radius used in plot style "with circles". It applies to data plots with only 2 columns of data (x,y) and to function plots. The default is "set style circle radius graph 0.02". **Nowedge** disables drawing of the two radii that connect the ends of an arc to the center. The default is **wedge**. This parameter has no effect on full circles. **Clip** clips the circle at the plot boundaries, **noclip** disables this. Default is **clip**.

Set style rectangle

Rectangles defined with the **set object** command can have individual styles. However, if the object is not assigned a private style then it inherits a default that is taken from the **set style rectangle** command.

Syntax:

```
set style rectangle {front|back} {lw|linewidth <lw>}
                  {fillcolor <colorspec>} {fs <fillstyle>}
```

See **colorspec** (p. 33) and **fillstyle** (p. 155). **fillcolor** may be abbreviated as **fc**.

Examples:

```
set style rectangle back fc rgb "white" fs solid 1.0 border lt -1
set style rectangle fc linestyle 3 fs pattern 2 noborder
```

The default values correspond to solid fill with the background color and a black border.

Set style ellipse

Syntax:

```
set style ellipse {units xx|xy|yy}
                  {size {graph|screen} <a>, {{graph|screen} <b>}}
                  {angle <angle>}
                  {clip|noclip}
```

This command governs whether the diameters of ellipses are interpreted in the same units or not. Default is **xy**, which means that the major diameter (first axis) of ellipses will be interpreted in the same units as the x (or x2) axis, while the minor (second) diameter in those of the y (or y2) axis. In this mode the ratio of the ellipse axes depends on the scales of the plot axes and aspect ratio of the plot. When set to **xx** or **yy**, both axes of all ellipses will be interpreted in the same units. This means that the ratio of the axes of the plotted ellipses will be correct even after rotation, but either their vertical or horizontal extent will not be correct.

This is a global setting that affects all ellipses, both those defined as objects and those generated with the **plot** command, however, the value of **units** can also be redefined on a per-plot and per-object basis.

It is also possible to set a default size for ellipses with the **size** keyword. This default size applies to data plots with only 2 columns of data (x,y) and to function plots. The two values are interpreted as the major and minor diameters (as opposed to semi-major and semi-minor axes) of the ellipse.

The default is "set style ellipse size graph 0.05,0.03".

Last, but not least it is possible to set the default orientation with the **angle** keyword. The orientation, which is defined as the angle between the major axis of the ellipse and the plot's x axis, must be given in degrees.

Clip clips the ellipse at the plot boundaries, **noclip** disables this. Default is **clip**.

For defining ellipse objects, see **set object ellipse** (p. 137); for the 2D plot style, see **ellipses** (p. 47).

Set style textbox

Syntax: set style textbox {opaque|transparent}{{no}border}

This command controls the appearance of labels with the attribute 'boxed'. Terminal types that do not support boxed text will ignore this style.

Surface

The **set surface** command is only relevant for 3D plots (**splot**).

Syntax:

```
set surface {implicit|explicit}
unset surface
show surface
```

unset surface will cause **splot** to not draw points or lines corresponding to any of the function or data file points. This is mainly useful for drawing only contour lines rather than the surface they were derived from. Contours may still be drawn on the surface, depending on the **set contour** option. To turn off the surface for an individual function or data file while leaving others active, use the **nosurface** keyword in the **splot** command. The combination **unset surface; set contour base** is useful for displaying contours on the grid base. See also **set contour** (p. 108).

If an 3D data set is recognizable as a mesh (grid) then by default the program implicitly treats the plot style **with lines** as requesting a gridded surface. See **grid.data** (p. 179). The command **set surface explicit** suppresses this expansion, plotting only the individual lines described by separate datablocks in the input file. A gridded surface can still be plotted by explicitly requesting **splot with surface**.

Table

When **table** mode is enabled, **plot** and **splot** commands print out a multicolumn text table of X Y {Z} R values rather than creating an actual plot on the current terminal. The character R takes on one of three values: "i" if the point is in the active range, "o" if it is out-of-range, or "u" if it is undefined. The data format is determined by the format of the axis tickmarks (see **set format** (p. 116)), and the columns are separated by single spaces. This can be useful if you want to generate contours and then save them for further use. The same method can be used to save interpolated data (see **set samples** (p. 151) and **set dgrid3d** (p. 112)).

Syntax:

```
set table {"outfile" | $datablock}
plot <whatever>
unset table
```

Tabular output is written to the named file, if any, otherwise it is written to the current value of **set output**. Alternatively, tabular output can be redirected to a named data block. Data block names start with '\$', see also **inline data** (p. 81). You must explicitly **unset table** in order to go back to normal plotting on the current terminal.

To avoid any style-dependent processing of the input data (smoothing, errorbar expansion, secondary range checking, etc), or to increase the number of columns that can be tabulated, you can use the keyword "table" instead of a normal plot style. For example

```
set table
plot <file> using 1:2:3:4:5:6:7:8:9:10 with table
```

Terminal

gnuplot supports many different graphics devices. Use **set terminal** to tell **gnuplot** what kind of output to generate. Use **set output** to redirect that output to a file or device.

Syntax:

```
set terminal {<terminal-type> | push | pop}
show terminal
```

If <terminal-type> is omitted, **gnuplot** will list the available terminal types. <terminal-type> may be abbreviated.

If both **set terminal** and **set output** are used together, it is safest to give **set terminal** first, because some terminals set a flag which is needed in some operating systems.

Some terminals have many additional options. The options used by a previous invocation **set term <term> <options>** of a given <term> are remembered, thus subsequent **set term <term>** does not reset them.

This helps in printing, for instance, when switching among different terminals — previous options don't have to be repeated.

The command **set term push** remembers the current terminal including its settings while **set term pop** restores it. This is equivalent to **save term** and **load term**, but without accessing the filesystem. Therefore they can be used to achieve platform independent restoring of the terminal after printing, for instance. After gnuplot's startup, the default terminal or that from **startup** file is pushed automatically. Therefore portable scripts can rely that **set term pop** restores the default terminal on a given platform unless another terminal has been pushed explicitly.

For more information, see the **complete list of terminals** (p. 184).

Termoption

The **set termoption** command allows you to change the behaviour of the current terminal without requiring a new **set terminal** command. Only one option can be changed per command, and only a small number of options can be changed this way. Currently the only options accepted are

```
set termoption {no}enhanced
set termoption font "<fontname>{,<fontsize>}"
set termoption fontscale <scale>
set termoption {solid|dashed}
set termoption {linewidth <lw>}{lw <lw>}
```

Tics

Control of the labeled tics on all axes at once is possible with the **set tics** command.

The tics may be turned off with the **unset tics** command, and may be turned on (the default state) with **set tics**. Fine control of tics on individual axes is possible using the alternative commands **set xtics**, **set ztics**, etc.

Syntax:

```
set tics {axis | border} {{no}mirror}
    {in | out} {front | back}
    {{no}rotate {by <ang>}} {offset <offset> | nooffset}
    {left | right | center | autojustify}
    {format "formatstring"} {font "name{,<size>}" } {{no}enhanced}
    { textcolor <colourspec> }
set tics scale {default | <major> {,<minor>}}
unset tics
show tics
```

The options can be applied to a single axis (x, y, z, x2, y2, cb), e.g.

```
set xtics rotate by -90
unset cbtics
```

Set tics **front** or **back** applies to all axes at once, but only for 2D plots (not splot). It controls whether the tics are placed behind or in front of the plot elements, in the case that there is overlap.

axis or **border** tells **gnuplot** to put the tics (both the tics themselves and the accompanying labels) along the axis or the border, respectively. If the axis is very close to the border, the **axis** option will move the tic labels to outside the border in case the border is printed (see **set border** (p. 102)). The relevant margin settings will usually be sized badly by the automatic layout algorithm in this case.

mirror tells **gnuplot** to put unlabeled tics at the same positions on the opposite border. **nomirror** does what you think it does.

in and **out** change the tic marks to be drawn inwards or outwards.

set tics scale controls the size of the tic marks. The first value <major> controls the auto-generated or user-specified major tics (level 0). The second value controls the auto-generated or user-specified minor tics

(level 1). `<major>` defaults to 1.0, `<minor>` defaults to `<major>/2`. Additional values control the size of user-specified tics with level 2, 3, ... Default tic sizes are restored by **set tics scale default**.

rotate asks **gnuplot** to rotate the text through 90 degrees, which will be done if the terminal driver in use supports text rotation. **norotate** cancels this. **rotate by <ang>** asks for rotation by `<ang>` degrees, supported by some terminal types.

The defaults are **border mirror norotate** for tics on the x and y axes, and **border nomirror norotate** for tics on the x2 and y2 axes. For the z axis, the default is **nomirror**.

The `<offset>` is specified by either x,y or x,y,z, and may be preceded by **first**, **second**, **graph**, **screen**, or **character** to select the coordinate system. `<offset>` is the offset of the tics texts from their default positions, while the default coordinate system is **character**. See **coordinates** (p. 21) for details. **nooffset** switches off the offset.

By default, tic labels are justified automatically depending on the axis and rotation angle to produce aesthetically pleasing results. If this is not desired, justification can be overridden with an explicit **left**, **right** or **center** keyword. **autojustify** restores the default behavior.

set tics with no options restores to place tics inwards. Every other options are retained.

See also **set xtics** (p. 168) for more control of major (labeled) tic marks and **set mxtics** for control of minor tic marks. These commands provide control of each axis independently.

Ticslevel

Deprecated. See **set xyplane** (p. 172).

Ticscale

The **set ticscale** command is deprecated, use **set tics scale** instead.

Timestamp

The command **set timestamp** places the time and date of the plot in the left margin.

Syntax:

```
set timestamp {"<format>"} {top|bottom} [{no}rotate]
               {offset <xoff>{,<yoff>}} {font "<fontspec>"}
               {textcolor <colourspec>}

unset timestamp
show timestamp
```

The format string allows you to choose the format used to write the date and time. Its default value is what `asctime()` uses: `"%a %b %d %H:%M:%S %Y"` (weekday, month name, day of the month, hours, minutes, seconds, four-digit year). With **top** or **bottom** you can place the timestamp at the top or bottom of the left margin (default: bottom). **rotate** lets you write the timestamp vertically, if your terminal supports vertical text. The constants `<xoff>` and `<yoff>` are offsets that let you adjust the position more finely. `` is used to specify the font with which the time is to be written.

The abbreviation **time** may be used in place of **timestamp**.

Example:

```
set timestamp "%d/%m/%y %H:%M" offset 80,-2 font "Helvetica"
```

See **set timefmt** (p. 161) for more information about time format strings.

Timefmt

This command applies to timeseries where data are composed of dates/times. It has no meaning unless the command **set *data time** is given also.

Syntax:

```
set timefmt "<format string>"
show timefmt
```

The string argument tells **gnuplot** how to read timedata from the datafile. The valid formats are:

Time Series timedata Format Specifiers	
Format	Explanation
%d	day of the month, 1–31
%m	month of the year, 1–12
%y	year, 0–99
%Y	year, 4-digit
%j	day of the year, 1–365
%H	hour, 0–24
%M	minute, 0–60
%s	seconds since the Unix epoch (1970-01-01 00:00 UTC)
%S	second, integer 0–60 on output, (double) on input
%b	three-character abbreviation of the name of the month
%B	name of the month

Any character is allowed in the string, but must match exactly. `\t` (tab) is recognized. Backslash-octals (`\nnn`) are converted to char. If there is no separating character between the time/date elements, then `%d`, `%m`, `%y`, `%H`, `%M` and `%S` read two digits each. If a decimal point immediately follows the field read by `%S`, the decimal and any following digits are interpreted as a fractional second. `%Y` reads four digits. `%j` reads three digits. `%b` requires three characters, and `%B` requires as many as it needs.

Spaces are treated slightly differently. A space in the string stands for zero or more whitespace characters in the file. That is, `"%H %M"` can be used to read `"1220"` and `"12 20"` as well as `"12 20"`.

Each set of non-blank characters in the timedata counts as one column in the **using n:n** specification. Thus **11:11 25/12/76 21.0** consists of three columns. To avoid confusion, **gnuplot** requires that you provide a complete **using** specification if your file contains timedata.

If the date format includes the day or month in words, the format string must exclude this text. But it can still be printed with the `"%a"`, `"%A"`, `"%b"`, or `"%B"` specifier. **gnuplot** will determine the proper month and weekday from the numerical values. See **set format** (p. 116) for more details about these and other options for printing time data.

When reading two-digit years with `%y`, values 69–99 refer to the 20th century, while values 00–68 refer to the 21st century. NB: This is in accordance with the UNIX98 spec, but conventions vary widely and two-digit year values are inherently ambiguous.

See also **set xdata** (p. 165) and **time/date** (p. 42) for more information.

Example:

```
set timefmt "%d/%m/%Y\t%H:%M"
```

tells **gnuplot** to read date and time separated by tab. (But look closely at your data — what began as a tab may have been converted to spaces somewhere along the line; the format string must match what is actually in the file.) See also

[time data demo](#).

Title

The **set title** command produces a plot title that is centered at the top of the plot. **set title** is a special case of **set label**.

Syntax:

```
set title {"<title-text>"} {offset <offset>} {font "<font>{,<size>}" }
      {{textcolor | tc} {<colorespec> | default}} {{no}enhanced}
show title
```

If `<offset>` is specified by either `x,y` or `x,y,z` the title is moved by the given offset. It may be preceded by **first**, **second**, **graph**, **screen**, or **character** to select the coordinate system. See **coordinates** (p. 21) for details. By default, the **character** coordinate system is used. For example, "**set title offset 0,-1**" will change only the y offset of the title, moving the title down by roughly the height of one character. The size of a character depends on both the font and the terminal.

`` is used to specify the font with which the title is to be written; the units of the font `<size>` depend upon which terminal is used.

textcolor `<colorspec>` changes the color of the text. `<colorspec>` can be a linetype, an rgb color, or a palette mapping. See help for **colorspec** (p. 33) and **palette** (p. 144).

noenhanced requests that the title not be processed by the enhanced text mode parser, even if enhanced text mode is currently active.

set title with no parameters clears the title.

See **syntax** (p. 41) for details about the processing of backslash sequences and the distinction between single- and double-quotes.

Tmargin

The command **set tmargin** sets the size of the top margin. Please see **set margin** (p. 131) for details.

Trange

The **set trange** command sets the parametric range used to compute x and y values when in parametric or polar modes. Please see **set xrange** (p. 167) for details.

Urange

The **set urange** and **set vrange** commands set the parametric ranges used to compute x, y, and z values when in **splot** parametric mode. Please see **set xrange** (p. 167) for details.

Variables

The **show variables** command lists the current value of user-defined and internal variables. Gnuplot internally defines variables whose names begin with `GPVAL_`, `MOUSE_`, `FIT_`, and `TERM_`.

Syntax:

```
show variables          # show variables that do not begin with GPVAL_
show variables all      # show all variables including those beginning GPVAL_
show variables NAME     # show only variables beginning with NAME
```

Version

The **show version** command lists the version of gnuplot being run, its last modification date, the copyright holders, and email addresses for the FAQ, the gnuplot-info mailing list, and reporting bugs—in short, the information listed on the screen when the program is invoked interactively.

Syntax:

```
show version {long}
```

When the **long** option is given, it also lists the operating system, the compilation options used when **gnuplot** was installed, the location of the help file, and (again) the useful email addresses.

View

The **set view** command sets the viewing angle for **splots**. It controls how the 3D coordinates of the plot are mapped into the 2D screen space. It provides controls for both rotation and scaling of the plotted data, but supports orthographic projections only. It supports both 3D projection or orthogonal 2D projection into a 2D plot-like map.

Syntax:

```
set view <rot_x>{,{<rot_z>}{,{<scale>}{,<scale_z>}}}  
set view map {scale <scale>}  
set view {no}equal {xy|xyz}  
show view
```

where <rot_x> and <rot_z> control the rotation angles (in degrees) in a virtual 3D coordinate system aligned with the screen such that initially (that is, before the rotations are performed) the screen horizontal axis is x, screen vertical axis is y, and the axis perpendicular to the screen is z. The first rotation applied is <rot_x> around the x axis. The second rotation applied is <rot_z> around the new z axis.

Command **set view map** is used to represent the drawing as a map. It is useful for **contour** plots or 2D heatmaps using pm3d mode rather than **with image**. In the latter case, take care that you properly use **zrange** and **cbrange** for input data point filtering and color range scaling, respectively.

<rot_x> is bounded to the [0:180] range with a default of 60 degrees, while <rot_z> is bounded to the [0:360] range with a default of 30 degrees. <scale> controls the scaling of the entire **splot**, while <scale_z> scales the z axis only. Both scales default to 1.0.

Examples:

```
set view 60, 30, 1, 1  
set view , ,0.5
```

The first sets all the four default values. The second changes only scale, to 0.5.

Equal axes

The command **set view equal xy** forces the unit length of the x and y axes to be on the same scale, and chooses that scale so that the plot will fit on the page. The command **set view equal xyz** additionally sets the z axis scale to match the x and y axes; however there is no guarantee that the current z axis range will fit within the plot boundary. By default all three axes are scaled independently to fill the available area.

See also **set xyplane** (p. 172).

Vrange

The **set urange** and **set vrange** commands set the parametric ranges used to compute x, y, and z values when in **splot** parametric mode. Please see **set xrange** (p. 167) for details.

X2data

The **set x2data** command sets data on the x2 (top) axis to timeseries (dates/times). Please see **set xdata** (p. 165).

X2dtics

The **set x2dtics** command changes ticks on the x2 (top) axis to days of the week. Please see **set xdtics** (p. 166) for details.

X2label

The **set x2label** command sets the label for the x2 (top) axis. Please see **set xlabel** (p. 166).

X2mtics

The **set x2mtics** command changes tics on the x2 (top) axis to months of the year. Please see **set xmtics** (p. 167) for details.

X2range

The **set x2range** command sets the horizontal range that will be displayed on the x2 (top) axis. See **set xrange** (p. 167) for the full set of command options. This command is ignored if the x2 axis range is explicitly linked to the x axis. See **set link** (p. 128).

X2tics

The **set x2tics** command controls major (labeled) tics on the x2 (top) axis. Please see **set xtics** (p. 168) for details.

X2zeroaxis

The **set x2zeroaxis** command draws a line at the origin of the x2 (top) axis ($y_2 = 0$). For details, please see **set zeroaxis** (p. 174).

Xdata

This command controls interpretation of data on the x axis. An analogous command acts on each of the other axes.

Syntax:

```
set xdata time
show xdata
```

The same syntax applies to **ydata**, **zdata**, **x2data**, **y2data** and **cbdata**.

The **time** option signals that data represents a time/date in seconds. The current version of gnuplot stores time to a millisecond precision.

If no option is specified, the data interpretation reverts to normal.

Time

set xdata time indicates that the x coordinate is to be interpreted as a date or time to millisecond precision. Equivalent commands exist for **set ydata time** and so on.

Two separate formats control interpretation of time data. Input data is read from a file using the **timefmt** for that axis. See **set timefmt** (p. 161). Use quoted strings in this same **timefmt** to specify the axis range.

Example:

```
set xdata time
set timefmt x "%d-%b-%Y"
set xrange ["01-Jan-2013" : "31-Dec-2014"]
```

For output, i.e. tick labels along that axis or coordinates output by mousing, the function 'strftime' (type "man strftime" on unix to look it up) is used to convert from the internal time in seconds to a string representation of a date. **gnuplot** tries to figure out a reasonable format for this. You can customize the format using either **set format x** or **set xtics format**. See **time_specifiers** (p. 118) for a special set of time format specifiers. See also **time/date** (p. 42) for more information.

Xdtics

The **set xdtics** commands converts the x-axis tic marks to days of the week where 0=Sun and 6=Sat. Overflows are converted modulo 7 to dates. **set noxdtics** returns the labels to their default values. Similar commands do the same things for the other axes.

Syntax:

```
set xdtics
unset xdtics
show xdtics
```

The same syntax applies to **ydtics**, **zdtics**, **x2dtics**, **y2dtics** and **cbdtics**.

See also the **set format** (p. 116) command.

Xlabel

The **set xlabel** command sets the x axis label. Similar commands set labels on the other axes.

Syntax:

```
set xlabel {"<label>"} {offset <offset>} {font "<font>{,<size>}" }
      {textcolor <colourspec>} {{no}enhanced}
      {rotate by <degrees> | rotate parallel | norotate}
show xlabel
```

The same syntax applies to **x2label**, **ylabel**, **y2label**, **zlabel** and **cblabel**.

If <offset> is specified by either x,y or x,y,z the label is moved by the given offset. It may be preceded by **first**, **second**, **graph**, **screen**, or **character** to select the coordinate system. See **coordinates** (p. 21) for details. By default, the **character** coordinate system is used. For example, "**set xlabel offset -1,0**" will change only the x offset of the title, moving the label roughly one character width to the left. The size of a character depends on both the font and the terminal.

 is used to specify the font in which the label is written; the units of the font <size> depend upon which terminal is used.

noenhanced requests that the label text not be processed by the enhanced text mode parser, even if enhanced text mode is currently active.

To clear a label, put no options on the command line, e.g., "**set y2label**".

The default positions of the axis labels are as follows:

xlabel: The x-axis label is centered below the bottom of the plot.

ylabel: The y-axis label is centered to the left of the plot, defaulting to either horizontal or vertical orientation depending on the terminal type.

zlabel: The z-axis label is centered along the z axis and placed in the space above the grid level.

cblabel: The color box axis label is centered along the box and placed below or to the right according to horizontal or vertical color box gradient.

y2label: The y2-axis label is placed to the right of the y2 axis. The position is terminal-dependent in the same manner as is the y-axis label.

x2label: The x2-axis label is placed above the plot but below the title. It is also possible to create an x2-axis label by using new-line characters to make a multi-line plot title, e.g.,

```
set title "This is the title\n\nThis is the x2label"
```

Note that double quotes must be used. The same font will be used for both lines, of course.

The orientation (rotation angle) of the x, x2, y and y2 axis labels in 2D plots can be changed by specifying **rotate by <degrees>**. The orientation of the x and y axis labels in 3D plots defaults to horizontal but can be changed to run parallel to the axis by specifying **rotate parallel**.

If you are not satisfied with the default position of an axis label, use **set label** instead—that command gives you much more control over where text is placed.

Please see **syntax** (p. 41) for further information about backslash processing and the difference between single- and double-quoted strings.

Xmtics

The **set xmtics** command converts the x-axis tic marks to months of the year where 1=Jan and 12=Dec. Overflows are converted modulo 12 to months. The tics are returned to their default labels by **unset xmtics**. Similar commands perform the same duties for the other axes.

Syntax:

```
set xmtics
unset xmtics
show xmtics
```

The same syntax applies to **x2mtics**, **ymtics**, **y2mtics**, **zmtics** and **cbmtics**.

See also the **set format** (p. 116) command.

Xrange

The **set xrange** command sets the horizontal range that will be displayed. A similar command exists for each of the other axes, as well as for the polar radius *r* and the parametric variables *t*, *u*, and *v*.

Syntax:

```
set xrange [{<min>}:{<max>}] [{no}reverse] [{no}writeback] [{no}extend]
| restore
show xrange
```

where *<min>* and *<max>* terms are constants, expressions or an asterisk to set autoscaling. If the data are time/date, you must give the range as a quoted string according to the **set timefmt** format. If *<min>* or *<max>* is omitted the current value will not be changed. See below for full autoscaling syntax. See also **noextend** (p. 100).

The same syntax applies to **yrange**, **zrange**, **x2range**, **y2range**, **cbrange**, **rrange**, **trange**, **urange** and **vrange**.

See **set link** (p. 128) for options that link the ranges of *x* and *x2*, or *y* and *y2*.

The **reverse** option reverses the direction of an autoscaled axis. For example, if the data values range from 10 to 100, it will autoscale to the equivalent of **set xrange [100:10]**. The **reverse** flag has no effect if the axis is not autoscaled. NB: This is a change introduced in version 4.7.

Autoscaling: If *<min>* (the same applies for correspondingly to *<max>*) is an asterisk "*" autoscaling is turned on. The range in which autoscaling is being performed may be limited by a lower bound *<lb>* or an upper bound *<ub>* or both. The syntax is

```
{ <lb> < > } * { < <ub> > }
```

For example,

```
0 < * < 200
```

sets *<lb>* = 0 and *<ub>* = 200. With such a setting *<min>* would be autoscaled, but its final value will be between 0 and 200 (both inclusive despite the '*<*' sign). If no lower or upper bound is specified, the '*<*' to also be omitted. If *<ub>* is lower than *<lb>* the constraints will be turned off and full autoscaling will happen. This feature is useful to plot measured data with autoscaling but providing a limit on the range, to clip outliers, or to guarantee a minimum range that will be displayed even if the data would not need such a big range.

The **writeback** option essentially saves the range found by **autoscale** in the buffers that would be filled by **set xrange**. This is useful if you wish to plot several functions together but have the range determined

by only some of them. The **writeback** operation is performed during the **plot** execution, so it must be specified before that command. To restore, the last saved horizontal range use **set xrange restore**. For example,

```
set xrange [-10:10]
set yrange [] writeback
plot sin(x)
set yrange restore
replot x/2
```

results in a yrange of [-1:1] as found only from the range of sin(x); the [-5:5] range of x/2 is ignored. Executing **show yrange** after each command in the above example should help you understand what is going on.

In 2D, **xrange** and **yrange** determine the extent of the axes, **trange** determines the range of the parametric variable in parametric mode or the range of the angle in polar mode. Similarly in parametric 3D, **xrange**, **yrange**, and **zrange** govern the axes and **urange** and **vrangle** govern the parametric variables.

In polar mode, **rrange** determines the radial range plotted. $\langle rmin \rangle$ acts as an additive constant to the radius, whereas $\langle rmax \rangle$ acts as a clip to the radius — no point with radius greater than $\langle rmax \rangle$ will be plotted. **xrange** and **yrange** are affected — the ranges can be set as if the graph was of $r(t)-rmin$, with $rmin$ added to all the labels.

Any range may be partially or totally autoscaled, although it may not make sense to autoscale a parametric variable unless it is plotted with data.

Ranges may also be specified on the **plot** command line. A range given on the plot line will be used for that single **plot** command; a range given by a **set** command will be used for all subsequent plots that do not specify their own ranges. The same holds true for **splot**.

Examples:

To set the xrange to the default:

```
set xrange [-10:10]
```

To set the yrange to increase downwards:

```
set yrange [10:-10]
```

To change zmax to 10 without affecting zmin (which may still be autoscaled):

```
set zrange [:10]
```

To autoscale xmin while leaving xmax unchanged:

```
set xrange [*:]
```

To autoscale xmin but keeping xmin positive:

```
set xrange [0<*:]
```

To autoscale x but keep minimum range of 10 to 50 (actual might be larger):

```
set xrange [*<10:50<*]
```

Autoscaling but limit maximum xrange to -1000 to 1000, i.e. autoscaling within [-1000:1000]

```
set xrange [-1000<*:*<1000]
```

Make sure xmin is somewhere between -200 and 100:

```
set xrange [-200<*<100:]
```

Xtics

Fine control of the major (labeled) ticks on the x axis is possible with the **set xtics** command. The ticks may be turned off with the **unset xtics** command, and may be turned on (the default state) with **set xtics**. Similar commands control the major ticks on the y, z, x2 and y2 axes.

Syntax:


```

set xtics {axis | border} {{no}mirror}
    {in | out} {scale {default | <major> {,<minor>}}}
    {{no}rotate {by <ang>}} {offset <offset> | nooffset}
    {left | right | center | autojustify}
    {add}
    {
        autofreq
        | <incr>
        | <start>, <incr> {,<end>}
        | ({"<label>" <pos> {<level>} {,{("<label>")...} }
    {format "formatstring"} {font "name{,<size>"} } {{no}enhanced}
    { numeric | timdate | geographic }
    { rangelimited }
    { textcolor <colorspec> }
unset xtics
show xtics

```

The same syntax applies to **ytics**, **ztics**, **x2tics**, **y2tics** and **cbtics**.

axis or **border** tells **gnuplot** to put the tics (both the tics themselves and the accompanying labels) along the axis or the border, respectively. If the axis is very close to the border, the **axis** option will move the tic labels to outside the border. The relevant margin settings will usually be sized badly by the automatic layout algorithm in this case.

mirror tells **gnuplot** to put unlabeled tics at the same positions on the opposite border. **nomirror** does what you think it does.

in and **out** change the tic marks to be drawn inwards or outwards.

With **scale**, the size of the tic marks can be adjusted. If **<minor>** is not specified, it is $0.5 * \text{<major>}$. The default size 1.0 for major tics and 0.5 for minor tics is requested by **scale default**.

rotate asks **gnuplot** to rotate the text through 90 degrees, which will be done if the terminal driver in use supports text rotation. **norotate** cancels this. **rotate by <ang>** asks for rotation by **<ang>** degrees, supported by some terminal types.

The defaults are **border mirror norotate** for tics on the x and y axes, and **border nomirror norotate** for tics on the x2 and y2 axes. For the z axis, the **{axis | border}** option is not available and the default is **nomirror**. If you do want to mirror the z-axis tics, you might want to create a bit more room for them with **set border**.

The **<offset>** is specified by either x,y or x,y,z, and may be preceded by **first**, **second**, **graph**, **screen**, or **character** to select the coordinate system. **<offset>** is the offset of the tics texts from their default positions, while the default coordinate system is **character**. See **coordinates** (p. 21) for details. **nooffset** switches off the offset.

Example:

Move xtics more closely to the plot.

```
set xtics offset 0,graph 0.05
```

By default, tic labels are justified automatically depending on the axis and rotation angle to produce aesthetically pleasing results. If this is not desired, justification can be overridden with an explicit **left**, **right** or **center** keyword. **autojustify** restores the default behavior.

set xtics with no options restores the default border or axis if xtics are being displayed; otherwise it has no effect. Any previously specified tic frequency or position {and labels} are retained.

Positions of the tics are calculated automatically by default or if the **autofreq** option is given; otherwise they may be specified in either of two forms:

The implicit **<start>**, **<incr>**, **<end>** form specifies that a series of tics will be plotted on the axis between the values **<start>** and **<end>** with an increment of **<incr>**. If **<end>** is not given, it is assumed to be infinity. The increment may be negative. If neither **<start>** nor **<end>** is given, **<start>** is assumed to be negative infinity, **<end>** is assumed to be positive infinity, and the tics will be drawn at integral multiples of **<incr>**. If the axis is logarithmic, the increment will be used as a multiplicative factor.

If you specify to a negative <start> or <incr> after a numerical value (e.g., **rotate by** <angle> or **offset** <offset>), the parser fails because it subtracts <start> or <incr> from that value. As a workaround, specify **0-<start>** resp. **0-<incr>** in that case.

Example:

```
set xtics border offset 0,0.5 -5,1,5
```

Fails with 'invalid expression' at the last comma.

```
set xtics border offset 0,0.5 0-5,1,5
```

or

```
set xtics offset 0,0.5 border -5,1,5
```

Sets tics at the border, tics text with an offset of 0,0.5 characters, and sets the start, increment, and end to -5, 1, and 5, as requested.

The **set grid** options 'front', 'back' and 'layerdefault' affect the drawing order of the xtics, too.

Examples:

Make tics at 0, 0.5, 1, 1.5, ..., 9.5, 10.

```
set xtics 0,.5,10
```

Make tics at ..., -10, -5, 0, 5, 10, ...

```
set xtics 5
```

Make tics at 1, 100, 1e4, 1e6, 1e8.

```
set logscale x; set xtics 1,100,1e8
```

The explicit ("**<label>**" **<pos>** **<level>**, ...) form allows arbitrary tic positions or non-numeric tic labels. In this form, the tics do not need to be listed in numerical order. Each tic has a position, optionally with a label. Note that the label is a string enclosed by quotes. It may be a constant string, such as "hello", may contain formatting information for converting the position into its label, such as "%3f clients", or may be empty, "". See **set format** (p. 116) for more information. If no string is given, the default label (numerical) is used.

An explicit tic mark has a third parameter, the level. The default is level 0, a major tic. Level 1 generates a minor tic. Labels are never printed for minor tics. Major and minor tics may be auto-generated by the program or specified explicitly by the user. Tics with level 2 and higher must be explicitly specified by the user, and take priority over auto-generated tics. The size of tics marks at each level is controlled by the command **set tics scale**.

Examples:

```
set xtics ("low" 0, "medium" 50, "high" 100)
set xtics (1,2,4,8,16,32,64,128,256,512,1024)
set ytics ("bottom" 0, "" 10, "top" 20)
set ytics ("bottom" 0, "" 10 1, "top" 20)
```

In the second example, all tics are labeled. In the third, only the end tics are labeled. In the fourth, the unlabeled tic is a minor tic.

Normally if explicit tics are given, they are used instead of auto-generated tics. Conversely if you specify **set xtics auto** or the like it will erase any previously specified explicit tics. You can mix explicit and auto-generated tics by using the keyword **add**, which must appear before the tic style being added.

Example:

```
set xtics 0,.5,10
set xtics add ("Pi" 3.14159)
```

This will automatically generate tic marks every 0.5 along x, but will also add an explicit labeled tic mark at pi.

However they are specified, tics will only be plotted when in range.

Format (or omission) of the tic labels is controlled by **set format**, unless the explicit text of a label is included in the **set xtics** ("**<label>**") form.

Minor (unlabeled) tics can be added automatically by the **set mxtics** command, or at explicit positions by the **set xtics** (" " **<pos>** 1, ...) form.

Xtics timedata

Times and dates are stored internally as a number of seconds.

Input: Non-numeric time and date values are converted to seconds on input using the format specifier in **timefmt**. Axis positions and range limits also may be given as quoted dates or times interpreted using **timefmt**. If the **<start>**, **<incr>**, **<end>** form is used, **<incr>** must be in seconds. Use of **timefmt** to interpret input data, range, and tic positions is triggered by **set xdata time**.

Output: Axis tic labels are generated using a separate format specified either by **set format** or **set xtics format**. By default the usual numeric format specifiers are expected (**set xtics numeric**). Other options are geographic coordinates (**set xtics geographic**), or times or dates (**set xtics time**).

Note: For backward compatibility with earlier gnuplot versions, the command **set xdata time** will implicitly also act do **set xtics time**, and **set xdata** or **unset xdata** will implicitly reset to **set xtics numeric**. However you can change this with a later call to **set xtics**.

Examples:

```
set xdata time           # controls interpretation of input data
set timefmt "%d/%m"      # format used to read input data
set xtics timedate       # controls interpretation of output format
set xtics format "%b %d" # format used for tic labels
set xrange ["01/12":"06/12"]
set xtics "01/12", 172800, "05/12"

set xdata time
set timefmt "%d/%m"
set xtics format "%b %d" time
set xrange ["01/12":"06/12"]
set xtics ("01/12", "" "03/12", "05/12")
```

Both of these will produce tics "Dec 1", "Dec 3", and "Dec 5", but in the second example the tic at "Dec 3" will be unlabeled.

Geographic

set xtics geographic indicates that x-axis values are to be interpreted as a geographic coordinate measured in degrees. Use **set xtics format** or **set format x** to specify the appearance of the axis tick labels. The format specifiers for geographic data are as follows:

%D	= integer degrees
%<width.precision>d	= floating point degrees
%M	= integer minutes
%<width.precision>m	= floating point minutes
%S	= integer seconds
%<width.precision>s	= floating point seconds
%E	= label with E/W instead of +/-
%N	= label with N/S instead of +/-

For example, the command **set format x "%Ddeg %5.2mmin %E"** will cause x coordinate -1.51 to be labeled as " 1deg 30.60min W".

If the xtics are left in the default state (**set xtics numeric**) the coordinate will be reported as a decimal number of degrees, and **format** will be assumed to contain normal numeric format specifiers rather than the special set above.

Xtics rangelimited

This option limits both the auto-generated axis tic labels and the corresponding plot border to the range of values actually present in the data that has been plotted. Note that this is independent of the current range limits for the plot. For example, suppose that the data in "file.dat" all lies in the range $2 < y < 4$. Then

the following commands will create a plot for which the left-hand plot border (y axis) is drawn for only this portion of the total y range, and only the axis ticks in this region are generated. I.e., the plot will be scaled to the full range on y, but there will be a gap between 0 and 2 on the left border and another gap between 4 and 10. This style is sometimes referred to as a **range-frame** graph.

```
set border 3
set yrange [0:10]
set ytics nomirror rangelimited
plot "file.dat"
```

Xyplane

The **set xyplane** command adjusts the position at which the xy plane is drawn in a 3D plot. The synonym "set ticslevel" is accepted for backwards compatibility.

Syntax:

```
set xyplane at <zvalue>
set xyplane relative <frac>
set ticslevel <frac>          # equivalent to set xyplane relative
show xyplane
```

The form **set xyplane relative <frac>** places the xy plane below the range in Z, where the distance from the xy plane to Zmin is given as a fraction of the total range in z. The default value is 0.5. Negative values are permitted, but tic labels on the three axes may overlap. The older, deprecated, form **set ticslevel** is retained for backwards compatibility.

To place the xy-plane at a position 'pos' on the z-axis, **ticslevel** may be set equal to $(\text{pos} - \text{zmin}) / (\text{zmin} - \text{zmax})$. However, this position will change if the z range is changed.

The alternative form **set xyplane at <zvalue>** fixes the placement of the xy plane at a specific Z value regardless of the current z range. Thus to force the x, y, and z axes to meet at a common origin one would specify **set xyplane at 0**.

See also **set view** (p. 164), and **set zeroaxis** (p. 174).

Xzeroaxis

The **set xzeroaxis** command draws a line at $y = 0$. For details, please see **set zeroaxis** (p. 174).

Y2data

The **set y2data** command sets y2 (right-hand) axis data to timeseries (dates/times). Please see **set xdata** (p. 165).

Y2dtics

The **set y2dtics** command changes ticks on the y2 (right-hand) axis to days of the week. Please see **set xdtics** (p. 166) for details.

Y2label

The **set y2label** command sets the label for the y2 (right-hand) axis. Please see **set xlabel** (p. 166).

Y2mtics

The **set y2mtics** command changes ticks on the y2 (right-hand) axis to months of the year. Please see **set xmtics** (p. 167) for details.

Y2range

The **set y2range** command sets the vertical range that will be displayed on the y2 (right) axis. See **set xrange** (p. 167) for the full set of command options. This command is ignored if the y2 axis range is explicitly linked to the y axis. See **set link** (p. 128).

Y2tics

The **set y2tics** command controls major (labeled) ticks on the y2 (right-hand) axis. Please see **set xtics** (p. 168) for details.

Y2zeroaxis

The **set y2zeroaxis** command draws a line at the origin of the y2 (right-hand) axis ($x_2 = 0$). For details, please see **set zeroaxis** (p. 174).

Ydata

The **set ydata** command sets y-axis data to timeseries (dates/times). Please see **set xdata** (p. 165).

Ydtics

The **set ydtics** command changes ticks on the y axis to days of the week. Please see **set xdtics** (p. 166) for details.

Ylabel

This command sets the label for the y axis. Please see **set xlabel** (p. 166).

Ymtics

The **set ymtics** command changes ticks on the y axis to months of the year. Please see **set xmtics** (p. 167) for details.

Yrange

The **set yrange** command sets the vertical range that will be displayed on the y axis. Please see **set xrange** (p. 167) for details.

Ytics

The **set ytics** command controls major (labeled) ticks on the y axis. Please see **set xtics** (p. 168) for details.

Yzeroaxis

The **set yzeroaxis** command draws a line at $x = 0$. For details, please see **set zeroaxis** (p. 174).

Zdata

The **set zdata** command sets zaxis data to timeseries (dates/times). Please see **set xdata** (p. 165).

Zdtics

The **set zdtics** command changes tics on the z axis to days of the week. Please see **set xdtics** (p. 166) for details.

Zzeroaxis

The **set zzeroaxis** command draws a line through (x=0,y=0). This has no effect on 2D plots, including **set view map**. For details, please see **set zeroaxis** (p. 174) and **set xyplane** (p. 172).

Cbdata

Set color box axis data to timeseries (dates/times). Please see **set xdata** (p. 165).

Cbdtics

The **set cbdtics** command changes tics on the color box axis to days of the week. Please see **set xdtics** (p. 166) for details.

Zero

The **zero** value is the default threshold for values approaching 0.0.

Syntax:

```
set zero <expression>
show zero
```

gnuplot will not plot a point if its imaginary part is greater in magnitude than the **zero** threshold. This threshold is also used in various other parts of **gnuplot** as a (crude) numerical-error threshold. The default **zero** value is 1e-8. **zero** values larger than 1e-3 (the reciprocal of the number of pixels in a typical bitmap display) should probably be avoided, but it is not unreasonable to set **zero** to 0.0.

Zeroaxis

The x axis may be drawn by **set xzeroaxis** and removed by **unset xzeroaxis**. Similar commands behave similarly for the y, x2, y2, and z axes. **set zeroaxis ...** (no prefix) acts on the x, y, and z axes jointly.

Syntax:

```
set {x|x2|y|y2|z}zeroaxis { {linestyle | ls <line_style>}
                           | { linetype | lt <line_type>}
                           { linewidth | lw <line_width>}}
unset {x|x2|y|y2|z}zeroaxis
show {x|y|z}zeroaxis
```

By default, these options are off. The selected zero axis is drawn with a line of type <line_type> and width <line_width> (if supported by the terminal driver currently in use), or a user-defined style <line_style>.

If no linetype is specified, any zero axes selected will be drawn using the axis linetype (linetype 0).

Examples:

To simply have the y=0 axis drawn visibly:

```
set xzeroaxis
```

If you want a thick line in a different color or pattern, instead:

```
set xzeroaxis linetype 3 linewidth 2.5
```

Zlabel

This command sets the label for the z axis. Please see **set xlabel** (p. 166).

Zmtics

The **set zmtics** command changes tics on the z axis to months of the year. Please see **set xmtics** (p. 167) for details.

Zrange

The **set zrange** command sets the range that will be displayed on the z axis. The zrange is used only by **splot** and is ignored by **plot**. Please see **set xrange** (p. 167) for details.

Ztics

The **set ztics** command controls major (labeled) tics on the z axis. Please see **set xtics** (p. 168) for details.

Cblabel

This command sets the label for the color box axis. Please see **set xlabel** (p. 166).

Cbmtics

The **set cbmtics** command changes tics on the color box axis to months of the year. Please see **set xmtics** (p. 167) for details.

Cbrange

The **set cbrange** command sets the range of values which are colored using the current **palette** by styles **with pm3d**, **with image** and **with palette**. Values outside of the color range use color of the nearest extreme.

If the cb-axis is autoscaled in **splot**, then the colorbox range is taken from **zrange**. Points drawn in **splot ... pm3d|palette** can be filtered by using different **zrange** and **cbrange**.

Please see **set xrange** (p. 167) for details on **set cbrange** (p. 175) syntax. See also **set palette** (p. 144) and **set colorbox** (p. 107).

Cbtics

The **set cbtics** command controls major (labeled) tics on the color box axis. Please see **set xtics** (p. 168) for details.

Shell

The **shell** command spawns an interactive shell. To return to **gnuplot**, type **logout** if using VMS, **exit** or the END-OF-FILE character if using Unix, or **exit** if using MS-DOS or OS/2.

There are two ways of spawning a shell command: using **system** command or via **!** (\$ if using VMS). The former command takes a string as a parameter and thus it can be used anywhere among other gnuplot commands, while the latter syntax requires to be the only command on the line. Control will return immediately to **gnuplot** after this command is executed. For example, in MS-DOS or OS/2,

```
! dir
```

or

```
system "dir"
```

prints a directory listing and then returns to **gnuplot**.

Other examples of the former syntax:

```
system "date"; set time; plot "a.dat"
print=1; if (print) replot; set out; system "lpr x.ps"
```

Splot

splot is the command for drawing 3D plots (well, actually projections on a 2D surface, but you knew that). It is the 3D equivalent of the **plot** command. **splot** provides only a single x, y, and z axis; there is no equivalent to the x2 and y2 secondary axes provided by **plot**.

See the **plot** (p. 73) command for many options available in both 2D and 3D plots.

Syntax:

```
splot {<ranges>}
      {<iteration>}
      <function> | {{<file name> | <data block name>} {datafile-modifiers}}
      {<title-spec>} {with <style>}
      {, {definitions{,}} <function> ...}
```

The **splot** command operates on a data generated by a function, read from a data file, or stored previously in a named data block. Data file names are usually provided as a quoted string. The function can be a mathematical expression, or a triple of mathematical expressions in parametric mode.

By default **splot** draws the xy plane completely below the plotted data. The offset between the lowest ztic and the xy plane can be changed by **set xyplane**. The orientation of a **splot** projection is controlled by **set view**. See **set view** (p. 164) and **set xyplane** (p. 172) for more information.

The syntax for setting ranges on the **splot** command is the same as for **plot**. In non-parametric mode, ranges must be given in the order

```
splot [<xrange>] [<yrange>] [<zrange>] ...
```

In parametric mode, the order is

```
splot [<urange>] [<vrangle>] [<xrange>] [<yrange>] [<zrange>] ...
```

The **title** option is the same as in **plot**. The operation of **with** is also the same as in **plot** except that not all 2D plotting styles are available.

The **datafile** options have more differences.

As an alternative to surfaces drawn using parametric or function mode, the pseudo-file '++' can be used to generate samples on a grid in the xy plane.

See also **show plot** (p. 140), **set view map** (p. 164), and **sampling** (p. 91).

Data-file

Splot, like **plot**, can display from a file.

Syntax:

```
splot '<file_name>' {binary <binary list>}
                        {{nonuniform} matrix}
                        {index <index list>}
                        {every <every list>}
                        {using <using list>}
```


The special filenames "" and "-" are permitted, as in **plot**. See **special-filenames** (p. 83).

In brief, **binary** and **matrix** indicate that the data are in a special form, **index** selects which data sets in a multi-data-set file are to be plotted, **every** specifies which datalines (subsets) within a single data set are to be plotted, and **using** determines how the columns within a single record are to be interpreted.

The options **index** and **every** behave the same way as with **plot**; **using** does so also, except that the **using** list must provide three entries instead of two.

The **plot** option **smooth** is not available for **splot**, but **cntrparam** and **dgrid3d** provide limited smoothing capabilities.

Data file organization is essentially the same as for **plot**, except that each point is an (x,y,z) triple. If only a single value is provided, it will be used for z, the datablock number will be used for y, and the index of the data point in the datablock will be used for x. If two or four values are provided, **gnuplot** uses the last value for calculating the color in pm3d plots. Three values are interpreted as an (x,y,z) triple. Additional values are generally used as errors, which can be used by **fit**.

Single blank records separate datablocks in a **splot** datafile; **splot** treats datablocks as the equivalent of function y-isolines. No line will join points separated by a blank record. If all datablocks contain the same number of points, **gnuplot** will draw cross-isolines between datablocks, connecting corresponding points. This is termed "grid data", and is required for drawing a surface, for contouring (**set contour**) and hidden-line removal (**set hidden3d**). See also **splot grid_data** (p. 179).

It is no longer necessary to specify **parametric** mode for three-column **splots**.

Matrix

Gnuplot can interpret matrix data input in two different ways.

The first of these assumes a uniform grid of x and y coordinates and assigns each value in the input matrix to one element M[i,j] of this uniform grid. The assigned x coordinates are the integers [0:NCOLS-1]. The assigned y coordinates are the integers [0:NROWS-1]. This is the default for text data input, but not for binary input. See **matrix uniform** (p. 177) for examples and additional keywords.

The second interpretation assumes a non-uniform grid with explicit x and y coordinates. The first row of input data contains the y coordinates; the first column of input data contains the x coordinates. For binary input data, the first element of the first row must contain the number of columns. This is the default for **binary matrix** input, but requires an additional keyword **nonuniform** for text input data. See **matrix nonuniform** (p. 178) for examples.

Uniform Example commands for plotting uniform matrix data:

```
splot 'file' matrix using 1:2:3      # text input
splot 'file' binary general using 1:2:3 # binary input
```

In a uniform grid matrix the z-values are read in a row at a time, i. e.,

```
z11 z12 z13 z14 ...
z21 z22 z23 z24 ...
z31 z32 z33 z34 ...
```

and so forth.

For text input, if the first row contains column labels rather than data, use the additional keyword **column-headers**. Similarly if the first field in each row contains a label rather than data, use the additional keyword **rowheaders**. Here is an example that uses both:

```
$DATA << EOD
xxx A   B   C   D
aa  z11 z12 z13 z14
bb  z21 z22 z23 z24
cc  z31 z32 z33 z34
EOD
plot $DATA matrix columnheaders rowheaders with image
```

For text input, a blank line or comment line ends the matrix, and starts a new surface mesh. You can select among the meshes inside a file by the **index** option to the **splot** command, as usual.

Nonuniform The first row of input data contains the y coordinates. The first column of input data contains the x coordinates. For binary input data, the first field of the first row must contain the number of columns. (This number is ignored for text input).

Example commands for plotting non-uniform matrix data:

```
splot 'file' nonuniform matrix using 1:2:3 # text input
splot 'file' binary matrix using 1:2:3     # binary input
```

Thus the data organization for non-uniform matrix input is

```
<N+1> <x0>  <x1>  <x2>  ... <xN>
<y0> <z0,0> <z0,1> <z0,2> ... <z0,N>
<y1> <z1,0> <z1,1> <z1,2> ... <z1,N>
:      :      :      :      ...   :
```

which is then converted into triplets:

```
<x0> <y0> <z0,0>
<x0> <y1> <z0,1>
<x0> <y2> <z0,2>
:      :      :
<x0> <yN> <z0,N>

<x1> <y0> <z1,0>
<x1> <y1> <z1,1>
:      :      :
```

These triplets are then converted into **gnuplot** iso-curves and then **gnuplot** proceeds in the usual manner to do the rest of the plotting.

Examples A collection of matrix and vector manipulation routines (in C) is provided in **binary.c**. The routine to write binary data is

```
int fwrite_matrix(file,m,nrl,nrl,ncl,nch,row_title,column_title)
```

An example of using these routines is provided in the file **bf_test.c**, which generates binary files for the demo file **demo/binary.dem**.

Usage in **plot**:

```
plot 'a.dat' matrix
plot 'a.dat' matrix using 1:3
plot 'a.gpbin' {matrix} binary using 1:3
```

will plot rows of the matrix, while using 2:3 will plot matrix columns, and using 1:2 the point coordinates (rather useless). Applying the **every** option you can specify explicit rows and columns.

Example – rescale axes of a matrix in a text file:

```
splot 'a.dat' matrix using (1+$1):(1+$2*10):3
```

Example – plot the 3rd row of a matrix in a text file:

```
plot 'a.dat' matrix using 1:3 every 1:999:1:2
```

(rows are enumerated from 0, thus 2 instead of 3).

Gnuplot can read matrix binary files by use of the option **binary** appearing without keyword qualifications unique to general binary, i.e., **array**, **record**, **format**, or **filetype**. Other general binary keywords for translation should also apply to matrix binary. (See **binary general** (p. 75) for more details.)

Example datafile

A simple example of plotting a 3D data file is

```
splot 'datafile.dat'
```

where the file "datafile.dat" might contain:

```
# The valley of the Gnu.
 0 0 10
 0 1 10
 0 2 10

 1 0 10
 1 1 5
 1 2 10

 2 0 10
 2 1 1
 2 2 10

 3 0 10
 3 1 0
 3 2 10
```

Note that "datafile.dat" defines a 4 by 3 grid (4 rows of 3 points each). Rows (datablocks) are separated by blank records.

Note also that the x value is held constant within each dataline. If you instead keep y constant, and plot with hidden-line removal enabled, you will find that the surface is drawn 'inside-out'.

Actually for grid data it is not necessary to keep the x values constant within a datablock, nor is it necessary to keep the same sequence of y values. **gnuplot** requires only that the number of points be the same for each datablock. However since the surface mesh, from which contours are derived, connects sequentially corresponding points, the effect of an irregular grid on a surface plot is unpredictable and should be examined on a case-by-case basis.

Grid data

The 3D routines are designed for points in a grid format, with one sample, datapoint, at each mesh intersection; the datapoints may originate from either evaluating a function, see **set isosamples** (p. 122), or reading a datafile, see **splot datafile** (p. 176). The term "isoline" is applied to the mesh lines for both functions and data. Note that the mesh need not be rectangular in x and y, as it may be parameterized in u and v, see **set isosamples** (p. 122).

However, **gnuplot** does not require that format. In the case of functions, 'samples' need not be equal to 'isosamples', i.e., not every x-isoline sample need intersect a y-isoline. In the case of data files, if there are an equal number of scattered data points in each datablock, then "isolines" will connect the points in a datablock, and "cross-isolines" will connect the corresponding points in each datablock to generate a "surface". In either case, contour and hidden3d modes may give different plots than if the points were in the intended format. Scattered data can be converted to a {different} grid format with **set dgrid3d**.

The contour code tests for z intensity along a line between a point on a y-isoline and the corresponding point in the next y-isoline. Thus a **splot** contour of a surface with samples on the x-isolines that do not coincide with a y-isoline intersection will ignore such samples. Try:

```
set xrange [-pi/2:pi/2]; set yrange [-pi/2:pi/2]
set style function lp
set contour
set isosamples 10,10; set samples 10,10;
splot cos(x)*cos(y)
set samples 4,10; replot
set samples 10,4; replot
```

Plot surfaces

splot can display a surface as a collection of points, or by connecting those points. As with **plot**, the points may be read from a data file or result from evaluation of a function at specified intervals, see **set isosamples** (p. 122). The surface may be approximated by connecting the points with straight line segments, see **set surface** (p. 158), in which case the surface can be made opaque with **set hidden3d**. The orientation from which the 3d surface is viewed can be changed with **set view**.

Additionally, for points in a grid format, **splot** can interpolate points having a common amplitude (see **set contour** (p. 108)) and can then connect those new points to display contour lines, either directly with straight-line segments or smoothed lines (see **set cntrparam** (p. 106)). Functions are already evaluated in a grid format, determined by **set isosamples** and **set samples**, while file data must either be in a grid format, as described in **data-file**, or be used to generate a grid (see **set dgrid3d** (p. 112)).

Contour lines may be displayed either on the surface or projected onto the base. The base projections of the contour lines may be written to a file, and then read with **plot**, to take advantage of **plot**'s additional formatting capabilities.

Stats (Statistical Summary)

Syntax:

```
stats {<ranges>} 'filename' {matrix | using N{:M}} {name 'prefix'} {{no}output}
```

This command prepares a statistical summary of the data in one or two columns of a file. The using specifier is interpreted in the same way as for plot commands. See **plot** (p. 73) for details on the **index** (p. 80), **every** (p. 79), and **using** (p. 85) directives. Data points are filtered against both xrange and yrange before analysis. See **set xrange** (p. 167). The summary is printed to the screen by default. Output can be redirected to a file by prior use of the command **set print**, or suppressed altogether using the **nooutput** option.

In addition to printed output, the program stores the individual statistics into three sets of variables. The first set of variables reports how the data is laid out in the file:

STATS_records	N	total number of in-range data records
STATS_outofrange		number of records filtered out by range limits
STATS_invalid		number of invalid/incomplete/missing records
STATS_blank		number of blank lines in the file
STATS_blocks		number of indexable datablocks in the file
STATS_columns		number of data columns in the first row of data

The second set reports properties of the in-range data from a single column. This column is treated as y. If the y axis is autoscaled then no range limits are applied. Otherwise only values in the range [ymin:ymax] are considered.

If two columns are analysed jointly by a single **stats** command, the suffix "_x" or "_y" is appended to each variable name. I.e. STATS_min_x is the minimum value found in the first column, while STATS_min_y is the minimum value found in the second column. In this case points are filtered by testing against both xrange and yrange.

STATS_min	$\min(y)$	minimum value of in-range data points
STATS_max	$\max(y)$	maximum value of in-range data points
STATS_index_min	$i \mid y_i = \min(y)$	index i for which data[i] == STATS_min
STATS_index_max	$i \mid y_i = \max(y)$	index i for which data[i] == STATS_max
STATS_mean	$\bar{y} = \frac{1}{N} \sum y$	mean value of the in-range data points
STATS_stddev	$\sigma_y = \sqrt{\frac{1}{N} \sum (y - \bar{y})^2}$	population standard deviation of the in-range data
STATS_ssd	$s_y = \sqrt{\frac{1}{N-1} \sum (y - \bar{y})^2}$	sample standard deviation of the in-range data
STATS_lo_quartile		value of the lower (1st) quartile boundary
STATS_median		median value
STATS_up_quartile		value of the upper (3rd) quartile boundary
STATS_sum	$\sum y$	sum
STATS_sumsq	$\sum y^2$	sum of squares
STATS_skewness	$\frac{1}{N\sigma^3} \sum (y - \bar{y})^3$	skewness of the in-range data points
STATS_kurtosis	$\frac{1}{N\sigma^4} \sum (y - \bar{y})^4$	kurtosis of the in-range data points
STATS_adev	$\frac{1}{N} \sum y - \bar{y} $	mean absolute deviation of the in-range data
STATS_mean_err	σ_y / \sqrt{N}	standard error of the mean value
STATS_stddev_err	$\sigma_y / \sqrt{2N}$	standard error of the standard deviation
STATS_skewness_err	$\sqrt{6/N}$	standard error of the skewness
STATS_kurtosis_err	$\sqrt{24/N}$	standard error of the kurtosis

The third set of variables is only relevant to analysis of two data columns.

STATS_correlation	sample correlation coefficient between x and y values
STATS_slope	A corresponding to a linear fit $y = Ax + B$
STATS_slope_err	uncertainty of A
STATS_intercept	B corresponding to a linear fit $y = Ax + B$
STATS_intercept_err	uncertainty of B
STATS_sumxy	sum of $x*y$
STATS_pos_min_y	x coordinate of a point with minimum y value
STATS_pos_max_y	x coordinate of a point with maximum y value

When **matrix** is specified, the **using** option is ignored and the "z"-values are treated as one-dimensional data set. The matrix dimensions are saved in the variables STATS.size.x and STATS.size.y.

It may be convenient to track the statistics from more than one file at the same time. The **name** option causes the default prefix "STATS" to be replaced by a user-specified string. For example, the mean value of column 2 data from two different files could be compared by

```
stats "file1.dat" using 2 name "A"
stats "file2.dat" using 2 name "B"
if (A_mean < B_mean) {...}
```

The index reported in STATS.index.xxx corresponds to the value of pseudo-column 0 (\$0) in plot commands. I.e. the first point has index 0, the last point has index N-1.

Data values are sorted to find the median and quartile boundaries. If the total number of points N is odd, then the median value is taken as the value of data point (N+1)/2. If N is even, then the median is reported as the mean value of points N/2 and (N+2)/2. Equivalent treatment is used for the quartile boundaries.

For an example of using the **stats** command to annotate a subsequent plot, see

[stats.dem.](#)

The current implementation does not allow analysis if either the X or Y axis is set to log-scaling. This restriction may be removed in a later version.

System

system "command" executes "command" using the standard shell. See **shell** (p. 175). If called as a function, **system("command")** returns the resulting character stream from stdout as a string. One

optional trailing newline is ignored.

This can be used to import external functions into gnuplot scripts:

```
f(x) = real(system(sprintf("somecommand %f", x)))
```

Test

This command graphically tests or presents terminal and palette capabilities.

Syntax:

```
test {terminal | palette}
```

test or **test terminal** creates a display of line and point styles and other useful things supported by the **terminal** you are currently using.

test palette plots profiles of $R(z)$, $G(z)$, $B(z)$, where $0 \leq z \leq 1$. These are the RGB components of the current color **palette**. It also plots the apparent net intensity as calculated using NTSC coefficients to map RGB onto a grayscale. The profile values are also loaded into a datablock named \$PALETTE.

Undefine

Clear one or more previously defined user variables. This is useful in order to reset the state of a script containing an initialization test.

A variable name can contain the wildcard character ***** as last character. If the wildcard character is found, all variables with names that begin with the prefix preceding the wildcard will be removed. This is useful to remove several variables sharing a common prefix. Note that the wildcard character is only allowed at the end of the variable name! Specifying the wildcard character as sole argument to **undefine** has no effect.

Example:

```
undefine foo foo1 foo2
if (!exists("foo")) load "initialize.gp"

bar = 1; bar1 = 2; bar2 = 3
undefine bar*           # removes all three variables
```

Unset

Options set using the **set** command may be returned to their default state by the corresponding **unset** command. The **unset** command may contain an optional iteration clause. See **plot for** (p. 91).

Examples:

```
set xtics mirror rotate by -45 0,10,100
...
unset xtics

# Unset labels numbered between 100 and 200
unset for [i=100:200] label i
```

Linetype

Syntax:

```
unset linetype N
```

Remove all characteristics previously associated with a single linetype. Subsequent use of this linetype will use whatever characteristics and color that is native to the current terminal type (i.e. the default linetypes properties available in gnuplot versions prior to 4.6).

Monochrome

Switches the active set of linetypes from monochrome to color. Equivalent to **set color**.

Output

Because some terminal types allow multiple plots to be written into a single output file, the output file is not automatically closed after plotting. In order to print or otherwise use the file safely, it should first be closed explicitly by using **unset output** or by using **set output** to close the previous file and then open a new one.

Terminal

The default terminal that is active at the time of program entry depends on the system platform, gnuplot build options, and the environmental variable GNUTERM. Whatever this default may be, gnuplot saves it in the internal variable GNUTERM. The **unset terminal** command restores this initial state. It is equivalent to **set terminal GNUTERM**.

Update

This command updates the current values of variables stored in the given file, which has to be formatted as an initial-value file (as described in the **fit** section).

If the file name does not exist, a new file is created containing all currently defined user variables. All variables not used in the last fit are marked as "#FIXED". This is useful for saving the current values of fit variables for later use or for restarting a converged or stopped fit.

Syntax:

```
update <filename> {<filename>}
```

If a second filename is supplied, the updated values are written to this file, and the original parameter file is left unmodified.

Otherwise, if the file already exists, **gnuplot** first renames it by appending **.old** and then opens a new file. That is, "**update 'fred'**" behaves the same as "**!rename fred fred.old; update 'fred.old' 'fred'**". If renaming is not possible because that file already exists, update aborts with an error message. [Renaming is not done at all on VMS systems, since they use file-versioning.]

Please see **fit** (p. 62) for more information.

While

Syntax:

```
while (<expr>) {  
    <commands>  
}
```

Execute a block of commands repeatedly so long as <expr> evaluates to a non-zero value. This command cannot be mixed with old-style (un-bracketed) if/else statements. See **if** (p. 70).

Part IV

Terminal types

Complete list of terminals

Gnuplot supports a large number of output formats. These are selected by choosing an appropriate terminal type, possibly with additional modifying options. See **set terminal** (p. 159).

This document may describe terminal types that are not available to you because they were not configured or installed on your system. To see a list of terminals available on a particular gnuplot installation, type 'set terminal' with no modifiers.

Aifm

NOTE: Outdated terminal, originally written for Adobe Illustrator 3.0+. Since Adobe Illustrator understands PostScript level 1 commands directly, you should use **set terminal post level1** instead.

Syntax:

```
set terminal aifm {color|monochrome} {"<fontname>"} {<fontsize>}
```

Aqua

This terminal relies on AquaTerm.app for display on Mac OS X.

Syntax:

```
set terminal aqua {<n>} {title "<wintitle>"} {size <x> <y>}
    {font "<fontname>{,<fontsize>}"}
    {<no>enhanced} {solid|dashed} {dl <dashlength>}}
```

where <n> is the number of the window to draw in (default is 0), <wintitle> is the name shown in the title bar (default "Figure <n>"), <x> <y> is the size of the plot (default is 846x594 pt = 11.75x8.25 in).

Use <fontname> to specify the font (default is "Times-Roman"), and <fontsize> to specify the font size (default is 14.0 pt).

The aqua terminal supports enhanced text mode (see **enhanced** (p. 22)), except for overprint. Font support is limited to the fonts available on the system. Character encoding can be selected by **set encoding** and currently supports iso_latin.1, iso_latin.2, cp1250, and UTF8 (default).

Lines can be drawn either solid or dashed, (default is solid) and the dash spacing can be modified by <dashlength> which is a multiplier > 0.

Be

The **be** terminal type is present if gnuplot is built for the **beos** operating system and for use with X servers. It is selected at program startup if the **DISPLAY** environment variable is set, if the **TERM** environment variable is set to **xterm**, or if the **-display** command line option is used.

Syntax:

```
set terminal be {reset} {<n>}
```

Multiple plot windows are supported: **set terminal be <n>** directs the output to plot window number n. If n>0, the terminal number will be appended to the window title and the icon will be labeled **gplt <n>**. The active window may distinguished by a change in cursor (from default to crosshair.)

Plot windows remain open even when the **gnuplot** driver is changed to a different device. A plot window can be closed by pressing the letter q while that window has input focus, or by choosing **close** from a

window manager menu. All plot windows can be closed by specifying **reset**, which actually terminates the subprocess which maintains the windows (unless **-persist** was specified).

Plot windows will automatically be closed at the end of the session unless the **-persist** option was given.

The size or aspect ratio of a plot may be changed by resizing the **gnuplot** window.

Linewidths and point sizes may be changed from within **gnuplot** with **set linestyle**.

For terminal type **be**, **gnuplot** accepts (when initialized) the standard X Toolkit options and resources such as geometry, font, and name from the command line arguments or a configuration file. See the X(1) man page (or its equivalent) for a description of such options.

A number of other **gnuplot** options are available for the **be** terminal. These may be specified either as command-line options when **gnuplot** is invoked or as resources in the configuration file ".Xdefaults". They are set upon initialization and cannot be altered during a **gnuplot** session.

Command-line options

In addition to the X Toolkit options, the following options may be specified on the command line when starting **gnuplot** or as resources in your ".Xdefaults" file:

'-mono'	forces monochrome rendering on color displays.
'-gray'	requests grayscale rendering on grayscale or color displays. (Grayscale displays receive monochrome rendering by default.)
'-clear'	requests that the window be cleared momentarily before a new plot is displayed.
'-raise'	raises plot window after each plot.
'-noraise'	does not raise plot window after each plot.
'-persist'	plots windows survive after main gnuplot program exits.

The options are shown above in their command-line syntax. When entered as resources in ".Xdefaults", they require a different syntax.

Example:

```
gnuplot*gray: on
```

gnuplot also provides a command line option (**-pointsize <v>**) and a resource, **gnuplot*pointsize: <v>**, to control the size of points plotted with the **points** plotting style. The value **v** is a real number (greater than 0 and less than or equal to ten) used as a scaling factor for point sizes. For example, **-pointsize 2** uses points twice the default size, and **-pointsize 0.5** uses points half the normal size.

Monochrome options

For monochrome displays, **gnuplot** does not honor foreground or background colors. The default is black-on-white. **-rv** or **gnuplot*reverseVideo: on** requests white-on-black.

Color resources

For color displays, **gnuplot** honors the following resources (shown here with their default values) or the greyscale resources. The values may be color names as listed in the BE rgb.txt file on your system, hexadecimal RGB color specifications (see BE documentation), or a color name followed by a comma and an **intensity** value from 0 to 1. For example, **blue, 0.5** means a half intensity blue.

```

gnuplot*background: white
gnuplot*textColor: black
gnuplot*borderColor: black
gnuplot*axisColor: black
gnuplot*line1Color: red
gnuplot*line2Color: green
gnuplot*line3Color: blue
gnuplot*line4Color: magenta
gnuplot*line5Color: cyan
gnuplot*line6Color: sienna
gnuplot*line7Color: orange
gnuplot*line8Color: coral

```

The command-line syntax for these is, for example,

Example:

```
gnuplot -background coral
```

Grayscale_resources

When **-gray** is selected, **gnuplot** honors the following resources for grayscale or color displays (shown here with their default values). Note that the default background is black.

```

gnuplot*background: black
gnuplot*textGray: white
gnuplot*borderGray: gray50
gnuplot*axisGray: gray50
gnuplot*line1Gray: gray100
gnuplot*line2Gray: gray60
gnuplot*line3Gray: gray80
gnuplot*line4Gray: gray40
gnuplot*line5Gray: gray90
gnuplot*line6Gray: gray50
gnuplot*line7Gray: gray70
gnuplot*line8Gray: gray30

```

Line_resources

gnuplot honors the following resources for setting the width (in pixels) of plot lines (shown here with their default values.) 0 or 1 means a minimal width line of 1 pixel width. A value of 2 or 3 may improve the appearance of some plots.

```

gnuplot*borderWidth: 2
gnuplot*axisWidth: 0
gnuplot*line1Width: 0
gnuplot*line2Width: 0
gnuplot*line3Width: 0
gnuplot*line4Width: 0
gnuplot*line5Width: 0
gnuplot*line6Width: 0
gnuplot*line7Width: 0
gnuplot*line8Width: 0

```

gnuplot honors the following resources for setting the dash style used for plotting lines. 0 means a solid line. A two-digit number **jk** (**j** and **k** are ≥ 1 and ≤ 9) means a dashed line with a repeated pattern of **j** pixels on followed by **k** pixels off. For example, '16' is a "dotted" line with one pixel on followed by six

pixels off. More elaborate on/off patterns can be specified with a four-digit value. For example, '4441' is four on, four off, four on, one off. The default values shown below are for monochrome displays or monochrome rendering on color or grayscale displays. For color displays, the default for each is 0 (solid line) except for **axisDashes** which defaults to a '16' dotted line.

```
gnuplot*borderDashes: 0
gnuplot*axisDashes: 16
gnuplot*line1Dashes: 0
gnuplot*line2Dashes: 42
gnuplot*line3Dashes: 13
gnuplot*line4Dashes: 44
gnuplot*line5Dashes: 15
gnuplot*line6Dashes: 4441
gnuplot*line7Dashes: 42
gnuplot*line8Dashes: 13
```

Cairolatex

The **cairolatex** terminal device generates encapsulated PostScript (*.eps) or PDF output using the cairo and pango support libraries and uses LaTeX for text output using the same routines as the **epslatex** terminal.

Syntax:

```
set terminal cairolatex
    {eps | pdf}
    {standalone | input}
    {blacktext | colortext | colourtext}
    {header <header> | noheader}
    {mono|color}
    {{no}transparent} {{no}crop} {background <rgbcolor>}
    {font <font>} {fontscale <scale>}
    {linewidth <lw>} {rounded|butt|square} {dashlength <dl>}
    {size <XX>{unit},<YY>{unit}}
```

The cairolatex terminal prints a plot like **terminal epscairo** or **terminal pdfcairo** but transfers the texts to LaTeX instead of including them in the graph. For reference of options not explained here see **pdfcairo** (p. ??).

eps and **pdf** select the type of graphics output. Use **eps** with latex/dvips and **pdf** for pdflatex.

blacktext forces all text to be written in black even in color mode;

The **cairolatex** driver offers a special way of controlling text positioning: (a) If any text string begins with '{', you also need to include a '}' at the end of the text, and the whole text will be centered both horizontally and vertically by LaTeX. (b) If the text string begins with '[', you need to continue it with: a position specification (up to two out of t,b,l,r,c), ']', the text itself, and finally, '}'. The text itself may be anything LaTeX can typeset as an LR-box. \rule{...}{...}'s may help for best positioning. See also the documentation for the **pslatex** (p. ??) terminal driver. To create multiline labels, use \shortstack, for example

```
set ylabel '[r]{\shortstack{first line \\ second line}}'
```

The **back** option of **set label** commands is handled slightly different than in other terminals. Labels using 'back' are printed behind all other elements of the plot while labels using 'front' are printed above everything else.

The driver produces two different files, one for the eps or pdf part of the figure and one for the LaTeX part. The name of the LaTeX file is taken from the **set output** command. The name of the eps/pdf file is derived by replacing the file extension (normally '.tex') with '.eps' or '.pdf' instead. There is no LaTeX output if no output file is given! Remember to close the **output file** before next plot unless in **multiplot** mode.

In your LaTeX documents use '\input{filename}' to include the figure. The '.eps' or '.pdf' file is included by the command \includegraphics{...}, so you must also include \usepackage{graphicx} in the LaTeX preamble.

If you want to use coloured text (option **colourtext**) you also have to include `\usepackage{color}` in the LaTeX preamble.

The behaviour concerning font selection depends on the header mode. In all cases, the given font size is used for the calculation of proper spacing. When not using the **standalone** mode the actual LaTeX font and font size at the point of inclusion is taken, so use LaTeX commands for changing fonts. If you use e.g. 12pt as font size for your LaTeX document, use `'', 12''` as options. The font name is ignored. If using **standalone** the given font and font size are used, see below for a detailed description.

If text is printed coloured is controlled by the TeX booleans `\ifGPcolor` and `\ifGPblacktext`. Only if `\ifGPcolor` is true and `\ifGPblacktext` is false, text is printed coloured. You may either change them in the generated TeX file or provide them globally in your TeX file, for example by using

```
\newif\ifGPblacktext
\GPblacktexttrue
```

in the preamble of your document. The local assignment is only done if no global value is given.

When using the `cairolatex` terminal give the name of the TeX file in the **set output** command including the file extension (normally `".tex"`). The graph filename is generated by replacing the extension.

If using the **standalone** mode a complete LaTeX header is added to the LaTeX file; and `"-inc"` is added to the filename of the gaph file. The **standalone** mode generates a TeX file that produces output with the correct size when using `dvips`, `pdfTeX`, or `VTex`. The default, **input**, generates a file that has to be included into a LaTeX document using the `\input` command.

If a font other than `""` or `"default"` is given it is interpreted as LaTeX font name. It contains up to three parts, separated by a comma: `'fontname,fontseries,fontshape'`. If the default `fontshape` or `fontseries` are requested, they can be omitted. Thus, the real syntax for the fontname is `'{fontname}{fontseries}{fontshape}'`. The naming convention for all parts is given by the LaTeX font scheme. The fontname is 3 to 4 characters long and is built as follows: One character for the font vendor, two characters for the name of the font, and optionally one additional character for special fonts, e.g., `'j'` for fonts with old-style numerals or `'x'` for expert fonts. The names of many fonts is described in

<http://www.tug.org/fontname/fontname.pdf>

For example, `'cmr'` stands for Computer Modern Roman, `'ptm'` for Times-Roman, and `'phv'` for Helvetica. The font series denotes the thickness of the glyphs, in most cases `'m'` for normal (`"medium"`) and `'bx'` or `'b'` for bold fonts. The font shape is `'n'` for upright, `'it'` for italics, `'sl'` for slanted, or `'sc'` for small caps, in general. Some fonts may provide different font series or shapes.

Examples:

Use Times-Roman boldface (with the same shape as in the surrounding text):

```
set terminal cairolatex font 'ptm,bx'
```

Use Helvetica, boldface, italics:

```
set terminal cairolatex font 'phv,bx,it'
```

Continue to use the surrounding font in slanted shape:

```
set terminal cairolatex font ',,sl'
```

Use small capitals:

```
set terminal cairolatex font ',,sc'
```

By this method, only text fonts are changed. If you also want to change the math fonts you have to use the `"gnuplot.cfg"` file or the **header** option, described below.

In **standalone** mode, the font size is taken from the given font size in the **set terminal** command. To be able to use a specified font size, a file `"size<size>.clo"` has to reside in the LaTeX search path. By default, 10pt, 11pt, and 12pt are supported. If the package `"extsizes"` is installed, 8pt, 9pt, 14pt, 17pt, and 20pt are added.

The **header** option takes a string as argument. This string is written into the generated LaTeX file. If using the **standalone** mode, it is written into the preamble, directly before the `\begin{document}` command. In

the **input** mode, it is placed directly after the `\begingroup` command to ensure that all settings are local to the plot.

Examples:

Use T1 fontencoding, change the text and math font to Times-Roman as well as the sans-serif font to Helvetica:

```
set terminal cairolatex standalone header \
"\usepackage[T1]{fontenc}\n\\usepackage{mathptmx}\n\\usepackage{helvet}"
```

Use a boldface font in the plot, not influencing the text outside the plot:

```
set terminal cairolatex input header "\\bfseries"
```

If the file "gnuplot.cfg" is found by LaTeX it is input in the preamble the LaTeX document, when using **standalone** mode. It can be used for further settings, e.g., changing the document font to Times-Roman, Helvetica, and Courier, including math fonts (handled by "mathptmx.sty"):

```
\usepackage{mathptmx}
\usepackage[scaled=0.92]{helvet}
\usepackage{courier}
```

The file "gnuplot.cfg" is loaded before the header information given by the **header** command. Thus, you can use **header** to overwrite some of settings performed using "gnuplot.cfg"

Canvas

The **canvas** terminal creates a set of javascript commands that draw onto the HTML5 canvas element. Syntax:

```
set terminal canvas {size <xsize>, <ysize>} {background <rgb_color>}
                    {font {<fontname>}{,<fontsize>}} | {fsize <fontsize>}
                    {{no}enhanced} {linewidth <lw>}
                    {rounded | butt | square}
                    {dashlength <dl>}
                    {standalone {mousing} | name '<funcname>'}
                    {jsdir 'URL/for/javascripts'}
                    {title '<some string>'}
```

where `<xsize>` and `<ysize>` set the size of the plot area in pixels. The default size in standalone mode is 600 by 400 pixels. The default font size is 10.

NB: Only one font is available, the ascii portion of Hershey simplex Roman provided in the file `canvastext.js`. You can replace this with the file `canvasmath.js`, which contains also UTF-8 encoded Hershey simplex Greek and math symbols. For consistency with other terminals, it is also possible to use **font "name,size"**. Currently the font **name** is ignored, but browser support for named fonts is likely to arrive eventually.

The default **standalone** mode creates an html page containing javascript code that renders the plot using the HTML 5 canvas element. The html page links to two required javascript files 'canvastext.js' and 'gnuplot_common.js'. An additional file 'gnuplot_dashedlines.js' is needed to support dashed lines. By default these point to local files, on unix-like systems usually in directory `/usr/local/share/gnuplot/<version>/js`. See installation notes for other platforms. You can change this by using the **jsdir** option to specify either a different local directory or a general URL. The latter is usually appropriate if the plot is exported for viewing on remote client machines.

All plots produced by the canvas terminal are mouseable. The additional keyword **mousing** causes the **standalone** mode to add a mouse-tracking box underneath the plot. It also adds a link to a javascript file 'gnuplot_mouse.js' and to a stylesheet for the mouse box 'gnuplot_mouse.css' in the same local or URL directory as 'canvastext.js'.

The **name** option creates a file containing only javascript. Both the javascript function it contains and the id of the canvas element that it draws onto are taken from the following string parameter. The commands

```
set term canvas name 'fishplot'
set output 'fishplot.js'
```

will create a file containing a javascript function `fishplot()` that will draw onto a canvas with `id=fishplot`. An html page that invokes this javascript function must also load the `canvastext.js` function as described above. A minimal html file to wrap the fishplot created above might be:

```
<html>
<head>
  <script src="canvastext.js"></script>
  <script src="gnuplot_common.js"></script>
</head>
<body onload="fishplot();">
  <script src="fishplot.js"></script>
  <canvas id="fishplot" width=600 height=400>
    <div id="err_msg">No support for HTML 5 canvas element</div>
  </canvas>
</body>
</html>
```

The individual plots drawn on this canvas will have names `fishplot_plot_1`, `fishplot_plot_2`, and so on. These can be referenced by external javascript routines, for example `gnuplot.toggle_visibility("fishplot_plot_2")`.

Cgm

The **cgm** terminal generates a Computer Graphics Metafile, Version 1. This file format is a subset of the ANSI X3.122-1986 standard entitled "Computer Graphics - Metafile for the Storage and Transfer of Picture Description Information".

Syntax:

```
set terminal cgm {color | monochrome} {solid | dashed} {{no}rotate}
                {<mode>} {width <plot_width>} {linewidth <line_width>}
                {font "<fontname>,<fontsize>"}
                {background <rgb_color>}
[deprecated]    {<color0> <color1> <color2> ...}
```

solid draws all curves with solid lines, overriding any dashed patterns; `<mode>` is **landscape**, **portrait**, or **default**; `<plot_width>` is the assumed width of the plot in points; `<line_width>` is the line width in points (default 1); `<fontname>` is the name of a font (see list of fonts below) `<fontsize>` is the size of the font in points (default 12).

The first six options can be in any order. Selecting **default** sets all options to their default values.

The mechanism of setting line colors in the **set term** command is deprecated. Instead you should set the background using a separate keyword and set the line colors using **set linetype**. The deprecated mechanism accepted colors of the form 'xrrggbb', where x is the literal character 'x' and 'rrggbb' are the red, green and blue components in hex. The first color was used for the background, subsequent colors are assigned to successive line types.

Examples:

```
set terminal cgm landscape color rotate dashed width 432 \
                linewidth 1 'Helvetica Bold' 12          # defaults
set terminal cgm linewidth 2 14 # wider lines & larger font
set terminal cgm portrait "Times Italic" 12
set terminal cgm color solid      # no pesky dashes!
```

Cgm font

The first part of a Computer Graphics Metafile, the metafile description, includes a font table. In the picture body, a font is designated by an index into this table. By default, this terminal generates a table with the following 35 fonts, plus six more with **italic** replaced by **oblique**, or vice-versa (since at least the Microsoft Office and Corel Draw CGM import filters treat **italic** and **oblique** as equivalent):

CGM fonts	
Helvetica	Hershey/Cartographic_Roman
Helvetica Bold	Hershey/Cartographic_Greek
Helvetica Oblique	Hershey/Simplex_Roman
Helvetica Bold Oblique	Hershey/Simplex_Greek
Times Roman	Hershey/Simplex_Script
Times Bold	Hershey/Complex_Roman
Times Italic	Hershey/Complex_Greek
Times Bold Italic	Hershey/Complex_Italic
Courier	Hershey/Complex_Cyrillic
Courier Bold	Hershey/Duplex_Roman
Courier Oblique	Hershey/Triplex_Roman
Courier Bold Oblique	Hershey/Triplex_Italic
Symbol	Hershey/Gothic_German
ZapfDingbats	Hershey/Gothic_English
Script	Hershey/Gothic_Italian
15	Hershey/Symbol_Set_1
	Hershey/Symbol_Set_2
	Hershey/Symbol_Math

The first thirteen of these fonts are required for WebCGM. The Microsoft Office CGM import filter implements the 13 standard fonts listed above, and also 'ZapfDingbats' and 'Script'. However, the script font may only be accessed under the name '15'. For more on Microsoft import filter font substitutions, check its help file which you may find here:

C:\Program Files\Microsoft Office\Office\Cgmimp32.hlp

and/or its configuration file, which you may find here:

C:\Program Files\Common Files\Microsoft Shared\Grphflt\Cgmimp32.cfg

In the **set term** command, you may specify a font name which does not appear in the default font table. In that case, a new font table is constructed with the specified font as its first entry. You must ensure that the spelling, capitalization, and spacing of the name are appropriate for the application that will read the CGM file. (Gnuplot and any MIL-D-28003A compliant application ignore case in font names.) If you need to add several new fonts, use several **set term** commands.

Example:

```
set terminal cgm 'Old English'
set terminal cgm 'Tengwar'
set terminal cgm 'Arabic'
set output 'myfile.cgm'
plot ...
set output
```

You cannot introduce a new font in a **set label** command.

Cgm fontsize

Fonts are scaled assuming the page is 6 inches wide. If the **size** command is used to change the aspect ratio of the page or the CGM file is converted to a different width, the resulting font sizes will be scaled up or down accordingly. To change the assumed width, use the **width** option.

Cgm linewidth

The **linewidth** option sets the width of lines in pt. The default width is 1 pt. Scaling is affected by the actual width of the page, as discussed under the **fontsize** and **width** options.

Cgm rotate

The **norotate** option may be used to disable text rotation. For example, the CGM input filter for Word for Windows 6.0c can accept rotated text, but the DRAW editor within Word cannot. If you edit a graph (for example, to label a curve), all rotated text is restored to horizontal. The Y axis label will then extend beyond the clip boundary. With **norotate**, the Y axis label starts in a less attractive location, but the page can be edited without damage. The **rotate** option confirms the default behavior.

Cgm solid

The **solid** option may be used to disable dashed line styles in the plots. This is useful when color is enabled and the dashing of the lines detracts from the appearance of the plot. The **dashed** option confirms the default behavior, which gives a different dash pattern to each line type.

Cgm size

Default size of a CGM plot is 32599 units wide and 23457 units high for landscape, or 23457 units wide by 32599 units high for portrait.

Cgm width

All distances in the CGM file are in abstract units. The application that reads the file determines the size of the final plot. By default, the width of the final plot is assumed to be 6 inches (15.24 cm). This distance is used to calculate the correct font size, and may be changed with the **width** option. The keyword should be followed by the width in points. (Here, a point is 1/72 inch, as in PostScript. This unit is known as a "big point" in TeX.) Gnuplot **expressions** can be used to convert from other units.

Example:

```
set terminal cgm width 432           # default
set terminal cgm width 6*72          # same as above
set terminal cgm width 10/2.54*72    # 10 cm wide
```

Cgm nofontlist

The default font table includes the fonts recommended for WebCGM, which are compatible with the Computer Graphics Metafile input filter for Microsoft Office and Corel Draw. Another application might use different fonts and/or different font names, which may not be documented. The **nofontlist** (synonym **winword6**) option deletes the font table from the CGM file. In this case, the reading application should use a default table. Gnuplot will still use its own default font table to select font indices. Thus, 'Helvetica' will give you an index of 1, which should get you the first entry in your application's default font table. 'Helvetica Bold' will give you its second entry, etc.

Part V

Bugs

Please e-mail bug reports to the gnuplot-bugs mailing list or upload the report to the gnuplot web site on SourceForge. Please give complete information on the version of gnuplot you are using and, if possible, a test script that demonstrates the bug. See **seeking-assistance** (p. 16).

Known limitations

It is not possible to use inline data (e.g. plot '-' ...) inside the curly brackets of a **do** or **while** loop.

Floating point exceptions (floating point number too large/small, divide by zero, etc.) may be generated by user defined functions. Some of the demos in particular may cause numbers to exceed the floating point range. Whether the system ignores such exceptions (in which case **gnuplot** labels the corresponding point as undefined) or aborts **gnuplot** depends on the compiler/runtime environment.

The gamma and bessel functions do not support complex arguments.

Coordinates specified as "time" wrap at 24 hours.

Parametric curves: The 'nohidden3d' option to exempt individual plots from the global setting 'set hidden3d' does not work for parametric curves. Iteration inside a **plot** command does not work for parametric curves.

X11 terminal: It is difficult to select UTF-8 fonts. Only one color palette at a time is active for any given x11 plot window. This means that multiplots whose constituent plots use different palettes will not display correctly in x11.

Qt terminal: 3D rotation of polygons and surfaces can be very slow; this is strongly affected by the Qt rendering mode (see Qt documentation).

External libraries

External library GD (used by PNG/JPEG/GIF drivers): Versions of libgd through 2.0.33 contain various bugs in mapping the characters of Adobe's Symbol font. Also it is possible to trigger a library segfault if an anti-aliased line crosses an upper corner of the canvas.

External library PDFlib (used by PDF driver): Gnuplot can be linked against libpdf versions 4, 5, or 6. However, these versions differ in their handling of piped I/O. Therefore gnuplot scripts using piped output to PDF may work only for some versions of PDFlib.

External library svgalib (used by linux and vga driver): Requires gnuplot to be suid root (bad!) and has many bugs that are specific to the video card or graphics driver used in X11.

Internationalization (locale settings): Gnuplot uses the C runtime library routine setlocale() to control locale-specific formatting of input and output number, times, and date strings. The locales available, and the level of support for locale features such as "thousands' grouping separator", depend on the internationalization support provided by your individual machine.

Part VI

Index

Index

- .gnuplot, 39
- 3D, 58

- abs, 24
- acos, 24
- acosh, 25
- acsplines, 82
- aifm, 184
- airy, 25
- all, 95
- angles, 98, 149
- Aqua, 184
- aqua, 184
- arg, 25
- arrow, 99, 154
- arrowstyle, 56, 99, 153
- asin, 25
- asinh, 25
- atan, 25
- atan2, 25
- atanh, 25
- automated, 52
- autoscale, 100
- avs, 76
- axes, 21, 32, 39, 74

- back, 36
- background, 34
- backquotes, 40
- bars, 46, 47, 49, 56–58, 88, 89, 102
- batch/interactive, 16, 19, 62, 72, 98
- BE, 184
- be, 184
- behind, 36
- besj0, 25
- besj1, 25
- bessel, 193
- besy0, 25
- besy1, 25
- bezier, 82
- bgnd, 34
- binary, 74, 76
- bind, 36, 97, 102, 132
- bitwise operators, 28
- black, 34
- bmarg, 102
- border, 102, 120, 160
- boxerrorbars, 44, 104
- boxes, 44, 46, 104
- boxplot, 45–47, 154
- boxwidth, 46, 47, 104
- boxxyerrorbars, 46
- branch, 68
- bugs, 193

- cairolatex, 187
- call, 59, 72
- candlesticks, 46, 49, 154
- canvas, 20
- cbdata, 174
- cbdtics, 174
- cblabel, 175
- cbmtics, 175
- cbrange, 33, 34, 141, 142, 157, 175
- cbtics, 175
- cd, 59
- cdawson, 25
- ceil, 25
- center, 53
- cerf, 25
- cgm, 190
- changes, 18
- circle, 47, 138
- circles, 47
- clabel, 105
- clear, 60
- clip, 56, 105
- cnormal, 83
- cntrlrlabel, 39, 105, 107, 108
- cntrparam, 39, 106, 108, 151, 180
- colorbox, 34, 107, 141–143, 175
- colornames, 33, 34, 108, 157
- colors, 33, 94, 104, 146, 157
- colorsequence, 104
- colourspec, 17, 33, 48, 126, 137, 157, 158, 163
- column, 26, 86
- columnhead, 26, 86
- columnheader, 22, 92, 123
- command line editing, 20
- command-line-editing, 96
- command-line-options, 38
- commands, 59
- comments, 16, 21, 72
- commentschars, 21, 111
- compatibility, 19
- contour, 39, 58, 107, 108, 120, 159, 180
- contours, 108
- coordinates, 21, 99, 125, 126, 137–139, 153, 161, 163, 166, 169
- copyright, 15
- corners2color, 143
- cos, 25
- cosh, 25
- csplines, 82
- cubehelix, 148
- cumulative, 83
- cycle, 128

- dashtype, 17, 33, 35, 109

- data, [17](#), [54](#), [73](#), [78](#), [109](#), [150](#), [159](#)
- data file, [78](#)
- datablocks, [32](#), [81](#)
- datafile, [39](#), [62](#), [78](#), [100](#), [108](#), [120](#), [179](#)
- datastrings, [21](#), [54](#), [92](#)
- date specifiers, [118](#)
- dawson, [25](#)
- decimalsign, [111](#), [114](#), [117](#), [129](#)
- degrees, [98](#)
- depthorder, [58](#), [142](#)
- dgrid3d, [112](#), [151](#), [159](#), [180](#)
- division, [24](#)
- do, [33](#), [61](#), [71](#)
- dots, [48](#)
- dummy, [101](#), [113](#)
- dx, [53](#), [77](#)
- dy, [53](#), [77](#)
- edf, [76](#)
- editing, [20](#)
- ehf, [76](#)
- ellipse, [48](#), [137](#), [138](#), [158](#)
- ellipses, [47](#), [158](#)
- elliptic integrals, [26](#)
- encoding, [22](#), [31](#), [114](#), [129](#)
- encodings, [114](#)
- enhanced, [22](#), [184](#)
- environment, [23](#), [67](#)
- eps, [32](#)
- equal, [152](#)
- equal axes, [164](#)
- erf, [25](#)
- erfc, [25](#)
- erfi, [26](#)
- error estimates, [64](#), [65](#)
- error state, [30](#), [97](#)
- errorbars, [78](#), [88](#)
- errorlines, [78](#), [88](#), [89](#)
- errors, [30](#)
- errorscaling, [67](#)
- evaluate, [61](#)
- every, [79](#), [180](#)
- example, [80](#)
- examples, [86](#)
- exists, [26](#), [40](#)
- exit, [62](#)
- exp, [25](#)
- expint, [25](#)
- exponentiation, [28](#)
- expressions, [24](#), [95](#), [119](#)
- factorial, [27](#)
- faddeeva, [25](#)
- failsafe, [53](#)
- FAQ, [16](#)
- faq, [16](#)
- file, [78](#)
- filetype, [76](#)
- fill, [44](#), [46](#), [47](#), [50](#), [51](#)
- filledcurves, [48](#)
- fillsteps, [50](#)
- fillstyle, [46–48](#), [93](#), [137](#), [155](#), [158](#)
- financebars, [46](#), [49](#), [154](#)
- fit, [23](#), [30](#), [62](#), [65](#), [67](#), [86](#), [116](#), [183](#)
- FIT LAMBDA FACTOR, [67](#)
- FIT LIMIT, [67](#)
- FIT LOG, [68](#)
- FIT MAXITER, [67](#)
- fit parameters, [64](#)
- FIT SCRIPT, [68](#)
- FIT START LAMBDA, [67](#)
- fitting, [64](#)
- flipx, [53](#)
- floating point exceptions, [109](#), [193](#)
- floor, [25](#)
- flush, [142](#)
- fontfile, [32](#)
- fontpath, [116](#)
- fonts, [31](#), [32](#)
- for, [33](#), [51](#), [52](#), [71](#), [98](#), [182](#)
- format, [116](#), [159](#), [162](#), [166](#), [167](#), [170](#)
- format specifiers, [117](#)
- fortran, [109](#)
- fpe trap, [109](#)
- frequency, [83](#)
- front, [36](#)
- fsteps, [49](#)
- ftriangles, [142](#)
- function, [89](#), [119](#)
- functions, [31](#), [73](#), [89](#)
- gamma, [25](#), [193](#)
- gamma correction, [148](#)
- gd, [31](#)
- general, [74](#), [111](#), [148](#), [178](#)
- geographic, [17](#), [171](#)
- geomean, [143](#)
- gif, [31](#)
- glossary, [32](#), [79](#)
- gnuplot, [15](#)
- gnuplot defined, [29](#)
- gprintf, [40](#), [117](#), [126](#)
- GPVAL, [29](#)
- gpval, [29](#)
- gray, [131](#)
- grid, [119](#), [151](#)
- grid data, [108](#), [112](#), [159](#), [177](#), [179](#)
- guidelines, [66](#)
- harmean, [143](#)
- heatmap, [53](#)
- help, [69](#)
- help desk, [16](#)
- hexadecimal, [24](#)

- hidden3d, 58, 120, 122
- histeps, 50
- histogram, 83
- histograms, 50
- history, 70, 96
- historysize, 122
- hotkey, 36
- hotkeys, 36
- hsv, 26, 33
- hsv2rgb, 26
- hypertext, 54, 127

- ibeta, 25
- if, 33, 61, 70, 183
- if old, 70
- if-old, 70
- igamma, 25
- imag, 25
- image, 53, 56
- import, 71
- impulses, 54
- index, 79, 80, 180
- initialization, 39, 97
- inline, 81
- int, 25
- internationalization, 193
- interval, 156
- introduction, 15
- inverf, 25
- invnorm, 25
- isosamples, 39, 89, 121, 122, 151, 179, 180
- iterate, 33
- iteration, 33, 61, 71, 91, 98, 182

- jpeg, 31

- kdensity, 18, 83, 112
- key, 92, 122

- label, 54, 125, 132
- labels, 22, 54, 79, 126, 132
- lambertw, 25
- layers, 36
- layout, 133
- lc, 33
- least squares, 62
- legend, 122
- lgamma, 25
- libgd, 193
- license, 15
- limit, 65
- line, 108, 153
- line editing, 20
- linecolor, 33, 44, 46, 49, 56
- lines, 54, 55
- linespoints, 55, 156
- linestyle, 55, 156
- linetype, 17, 33, 55, 94, 104, 128, 156
- linetypes, 33
- linewidth, 55, 156
- link, 17, 101, 128, 165, 167, 173
- lmargin, 129
- load, 72
- loadpath, 129
- locale, 31, 111, 114, 129, 193
- log, 25, 104
- log10, 25
- logscale, 129
- lower, 72
- lp, 55

- macros, 31, 40, 61
- map, 58, 143, 144, 176
- mapping, 39, 130, 150
- margin, 102, 129, 134, 151, 163
- margins, 131
- markup, 22
- Marquardt, 62
- matrix, 74, 75, 80, 111, 177
- max, 143
- maxiter, 65
- mcsplines, 82
- mean, 143
- median, 143
- min, 143
- missing, 18, 109
- mixing macros backquotes, 41
- modulo, 28
- modulus, 24
- monochrome, 33, 104, 131
- mouse, 36, 37, 131
- mouseformat, 132
- mousewheel, 133
- mousing, 131
- multi branch, 68
- multi-branch, 63, 68
- multiplot, 61, 133
- mx2tics, 135
- mxtics, 135, 136, 151
- my2tics, 136
- mytics, 136
- mztics, 136

- NaN, 24, 30, 87
- negation, 27
- negative, 146
- new features, 17
- newhistogram, 52
- noarrow, 99
- noautoscale, 100
- noborder, 102
- nocbdtics, 174
- nocbmtics, 175
- nocbtics, 175
- noclip, 105

- nocontour, 108
- nodgrid3d, 112
- noextend, 100, 167
- nofpe trap, 109
- nogrid, 119
- nohidden3d, 120, 193
- nokey, 122
- nolabel, 125
- nologscale, 129
- nomouse, 131
- nomultiplot, 133
- nomx2tics, 135
- nomxtics, 135
- nomy2tics, 136
- nomytics, 136
- nomztics, 136
- nonuniform, 177
- nooffsets, 139
- noparametric, 140
- nopolar, 149
- norm, 24, 25
- nosurface, 158
- notimestamp, 161
- nox2dtics, 164
- nox2mtics, 165
- nox2tics, 165
- nox2zeroaxis, 165
- noxdtics, 166
- noxmtics, 167
- noxtics, 168
- noxzeroaxis, 172
- noy2dtics, 172
- noy2mtics, 172
- noy2tics, 173
- noy2zeroaxis, 173
- noydtics, 173
- noymtics, 173
- noytics, 173
- noyzeroaxis, 173
- nozdtics, 174
- nozmtics, 175
- noztics, 175
- nozeroaxis, 174
- object, 136
- octal, 24
- offsets, 131, 139
- old-style, 59
- one's complement, 27
- operator precedence, 27
- operators, 27
- options, 16
- origin, 61, 134, 139
- output, 139
- parallelaxes, 55, 140
- parametric, 101, 140
- pause, 72
- paxis, 55
- pdf, 31, 193
- pdfcairo, 187
- persist, 38
- pi, 30
- pixels, 53
- placement, 123
- plot, 17, 73, 74, 96, 176, 180
- plotting, 38
- pm3d, 58, 108, 141, 157
- png, 31
- pointinterval, 55, 156
- pointintervalbox, 149
- points, 55
- pointsize, 93, 149
- polar, 39, 55, 149, 151
- polygon, 138
- pop, 159
- position, 141
- positive, 146
- postscript, 32
- practical guidelines, 66
- print, 95
- printerr, 95
- projection, 58
- prologue, 24
- psdir, 150
- pseudocolumns, 81, 86, 87
- pslatex, 187
- punctuation, 41
- push, 159
- pwd, 95
- quit, 95
- quotes, 16, 42
- raise, 95
- rand, 25, 26
- random, 26
- range frame, 171
- rangelimited, 171
- ranges, 62, 90
- ratio, 151
- raxis, 151
- real, 25
- rectangle, 137
- refresh, 85, 88, 96
- replot, 85, 96
- reread, 96
- reset, 97, 98
- restore, 167
- rgbalpha, 53
- rgbcolor, 33, 35
- rgbformulae, 146

- rgbimage, 53
- rmargin, 151
- rms, 143
- rotate, 53
- rrange, 55, 102, 150, 151
- rtics, 151

- sample, 91
- samples, 82, 89, 121, 122, 151, 159
- sampling, 84, 90, 91, 151, 176
- save, 97
- sbezier, 83
- scansautomatic, 142
- scansbackward, 142
- scansforward, 142
- scrolling, 133
- seeking assistance, 16
- seeking-assistance, 193
- separator, 86, 110
- session, 97
- set, 98
- sgn, 25
- shell, 175, 181
- show, 98
- sin, 25
- sinh, 25
- size, 61, 134, 140, 151
- SJIS, 114
- skip, 80, 81
- smooth, 82, 151
- space, 37
- special filenames, 83
- special-filenames, 40, 81, 177
- specifiers, 117
- specify, 41
- splot, 96, 120, 150, 176
- sprintf, 26, 40, 126
- sqrt, 25
- square, 55, 151, 152
- start, 39
- start up, 39
- starting values, 68
- startup, 23, 39
- statistical overview, 65
- statistics, 180
- stats, 180
- steps, 50, 56
- strcol, 26
- strftime, 26
- string operators, 28
- stringcolumn, 26
- strings, 39, 126
- strlen, 26
- strptime, 26
- strstrt, 26
- style, 86, 88, 89, 93, 126
- styles, 73, 93, 152, 155

- substitution, 40, 42, 130
- substr, 26
- substring, 26, 40
- summation, 29, 33
- surface, 58, 108, 158, 180
- svglib, 193
- syntax, 16, 41, 117, 123, 163, 167
- system, 26, 181

- table, 159
- tan, 25
- tanh, 25
- tc, 33
- term, 98, 184
- terminal, 184
- terminals, 160
- termoption, 160
- ternary, 28
- test, 34, 93, 182
- text, 17, 42, 126
- text markup, 22
- textbox, 126, 158
- textcolor, 33
- thru, 85
- tics, 160
- ticscale, 161
- ticslevel, 161
- time, 26
- time specifiers, 43, 118, 165
- time/date, 42, 162, 165
- timecolumn, 26
- timefmt, 21, 90, 118, 126, 161, 165, 193
- timestamp, 161
- tips, 69
- title, 22, 122, 162
- tm hour, 26
- tm mday, 26
- tm min, 26
- tm mon, 26
- tm sec, 26
- tm wday, 26
- tm yday, 26
- tm year, 26
- tmargin, 163
- trange, 163
- transparency, 53
- transparent, 155

- unary, 27
- undefine, 182
- uniform, 177
- unique, 82, 83
- unset, 182
- unwrap, 83
- update, 64, 183
- urange, 163
- user defined, 30

- user-defined, 89
- using, 22, 24, 29, 44, 85, 88, 89, 142, 180
- UTF 8, 114
- valid, 26
- value, 27, 31
- variable, 44, 46, 49, 54, 56, 81, 143
- variables, 27, 30, 31, 36, 38, 67, 73
- vectors, 56
- view, 164, 172, 176
- voigt, 25
- volatile, 88
- VP, 26
- vrange, 164
- while, 33, 183
- with, 88, 89, 93, 149, 152
- word, 26, 27
- words, 26, 27
- writeback, 167
- wxt, 31
- x11, 16
- x11 mouse, 133
- x2data, 164
- x2dtics, 164
- x2label, 164
- x2mtics, 165
- x2range, 165
- x2tics, 165
- x2zeroaxis, 165
- xdata, 21, 126, 162, 164, 165, 172–174
- xdtics, 164, 166, 172–174
- xerrorbars, 56
- xerrorlines, 57
- xlabel, 164, 166, 172, 173, 175
- xmtics, 165, 167, 172, 173, 175
- xrange, 101, 140, 150, 163–165, 167, 173, 175, 180
- xticlabels, 22, 87, 88
- xtics, 103, 117, 120, 129, 136, 140, 151, 161, 165, 168, 173, 175
- xyerrorbars, 57
- xyerrorlines, 57
- xyplane, 21, 161, 164, 172, 174, 176
- xzeroaxis, 172
- y2data, 172
- y2dtics, 172
- y2label, 172
- y2mtics, 172
- y2range, 173
- y2tics, 173
- y2zeroaxis, 173
- ydata, 173
- ydtics, 173
- yerrorbars, 57
- yerrorlines, 58
- ylabel, 173
- ymtics, 173
- yrange, 173
- ytics, 173
- yzeroaxis, 173
- zdata, 173
- zdtics, 174
- zero, 174
- zeroaxis, 165, 172–174
- zlabel, 175
- zmtics, 175
- zoom, 133
- zrange, 175
- ztics, 175
- zzeroaxis, 174