

Notes on Lattice Graphics¹

J H Maindonald

Lattice Graphics:

Lattice	Lattice is a flavour of trellis graphics (the S-PLUS flavour was the original)
Grid	<i>grid</i> is a low-level graphics system. It was used to build <i>lattice</i> . For <i>grid</i> , see Part II of Paul Murrell's <i>R Graphics</i>
Lattice vs base	Lattice is more structured, automated and stylized. Much is done automatically, without user intervention. Changes to the default style are harder than for base.
Lattice syntax	Lattice syntax is consistent and tightly regulated. For lattice, graphics formulae are mandatory.

Lattice (trellis) graphics functions allow the use of the layout on the page to reflect meaningful aspects of data structure. Different levels of a factor may appear in different panels. Or they may appear in the same panel, distinguished by color and/or symbol. If lines or smooth curves are added, there is a different line or curve for each different group.

Similar considerations apply when columns of data are plotted in parallel. Different columns may appear in different panels. Or they may appear in the same panel, distinguished by color and/or symbol.

To see some of the possibilities that lattice graphics offers, enter

```
library(lattice)
demo(lattice)
```

1 Lattice Graphics vs Base Graphics

Contrast the different ways that base and lattice graphics are designed to operate.

- In base graphics, a graphics page (possibly the first of a sequence of pages) opens when a device is opened. Plots then go to the page or pages. For a screen device, plots go to the screen. For a hardcopy device, plots usually go, in the first place, to a file.
- Lattice functions create trellis objects. Objects can be created even if no device is open. Such objects can be updated. Objects are plotted (by this time, a device must be open), either when output from a lattice function goes to the command line (thus implicitly invoking the `print()` command), or by the explicit use of `print()`.

The updating feature allows the graphics object to be built up in steps, or even modified. Additionally, abilities that were added relatively late in lattice's development make it possible to add new features to the "printed" page, after a style of use that is common in base graphics.

¹© J. H. Maindonald 2008. Permission is given to make copies for personal study and class use. (June 9, 2008)

The lattice package comes already installed with all the binary distributions that are supplied from CRAN (Comprehensive R Archive Network: <http://mirror.aarnet.edu.au/pub//CRAN/>). For use in an R session, it must first be attached.

```
library(lattice)
```

Now compare:

```
plot(ACT ~ year, data=austpop)      # Base graphics
xyplot(ACT ~ year, data=austpop)    # Lattice graphics
```

In both cases, if these are typed from the command line, a graph is plotted. The reason is different in the two cases:

- `plot()` gives a graph as a side effect of the command.
- `xyplot()` generates a graphics object. As this is output to the command line, the object is “printed”, i.e., a graph appears.

The following makes this clear: The following makes clear the difference between the two functions:

```
invisible(plot(ACT ~ year, data=austpop))    # A graph is plotted
invisible(xyplot(ACT ~ year, data=austpop))  # Graph does appear
```

The function `invisible()` suppresses the command line printing. Hence `invisible(xyplot(...))` does not yield a graph.

Inside a function, `xyplot(...)` prints a graph only if it appears as the return value from the function, i.e. usually, as the final line. In a file that is sourced, no graph will appear. Inside a function (except as mentioned), or in a file that is sourced, there must be an explicit `print()`, i.e.

```
print(xyplot(ACT ~ year, data=austpop))
```

2 Groups within Data, and/or Columns in Parallel

Here are selected lines from the data set `grog` (*DAAGxtras* package):

	Beer	Wine	Spirit	Country	Year
1	5.24	2.86	1.81	Australia	1998
2	5.15	2.87	1.77	Australia	1999
...					
9	4.57	3.11	2.15	Australia	2006
10	4.50	2.59	1.77	NewZealand	1998
11	4.28	2.65	1.64	NewZealand	1999
...					
18	3.96	3.09	2.20	NewZealand	2006

There are three drinks (liquor products), shown in different columns, and two countries, occupying rows that are indexed by the factor `Country`. The lattice function `xyplot()` can accommodate any of the following possibilities:

- Different symbols and/or colors are used for different drinks, within the one panel. Different panels must then be used for different countries, as in Figure 1.² Or if different countries are shown in the same panel, then different panels must be used for the different drinks.

²The data (dataset `grog`, from *DAAGxtras*) are 1998 – 2006 Australian and New Zealand apparent per person annual consumption (in liters) of the pure alcohol content of Beer, Wine and Spirit. Data, based on Australian Bureau of Statistics and Statistics New Zealand figures, are obtained by dividing estimates of total available alcohol by number of persons aged 15 or more.

- Use a 3 drinks \times 2 countries, or 2 countries \times 3 drinks layout of panels.

Where plots are superposed in the one panel and, e.g. regression lines or smooth curves are fitted, this will be done separately for each different set of points. The separate sets may be distinguished by colour, and/or they may be distinguished by different symbols and/or line styles.

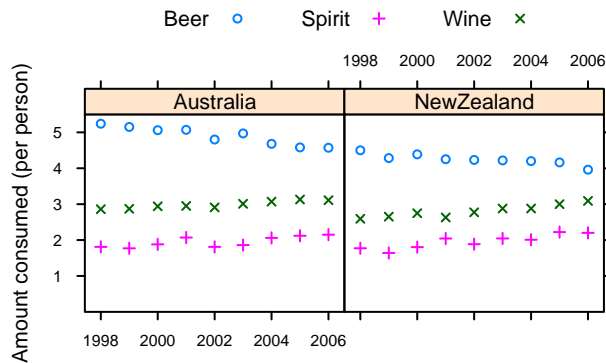


Figure 1: Australian and New Zealand apparent per person annual consumption (in liters) of the pure alcohol content of liquor products, for 1998 to 2006.

Simplified code for Figure 1 is:

```
xyplot(Beer+Spirit+Wine ~ Year | Country, data=grog, outer=FALSE,
       auto.key=list(columns=3))
```

To obtain various enhancements that give Figure 1, specify:

```
grogplot <-
  xyplot(Beer+Spirit+Wine ~ Year | Country, data=grog, outer=FALSE,
        auto.key=list(columns=3))
update(grogplot, ylim=c(0,5.5),
       xlab="", ylab="Amount consumed (per person)",
       par.settings=simpleTheme(pch=c(1,3,4)))
```

The footnote³ has alternative code that updates the object, then uses an explicit `print()`.

Observe that:

- Use of `Beer+Spirit+Wine` gives plots for each of Beer, Spirit and Wine. With `outer=FALSE`, these appear in the same panel.
- Conditioning by country (`| Country`) gives separate panels for separate countries.
- The function `simpleTheme()` is convenient for setting or changing point and line settings.

The plot superimposes the separate plots (two panels each):

```
xyplot(Beer ~ Year | Country, data=grog)      # Plot for Beer
xyplot(Spirit ~ Year | Country, data=grog)   # Plot for Spirit
xyplot(Wine ~ Year | Country, data=grog)     # Plot for Wine
```

For separate panels for the three liquor products (different levels of Country can now use the same panel), specify `outer=TRUE`:

```
xyplot(Beer+Spirit+Wine ~ Year, groups=Country, outer=TRUE,
       data=grog, auto.key=list(columns=2))
```

³## Update trellis object, then print
frillyplot <-
 update(grogplot, ylim=c(0,5.5),
 xlab="", ylab="Amount consumed (per person)",
 par.settings=simpleTheme(pch=c(1,3,4)))
print(frillyplot)

Here is a summary of the syntax:

	Overplot (a single panel)	Separate panels (conditioning)
Levels of a factor	$\text{Beer} \sim \text{Year}, \text{groups}=\text{Country}$	$\text{Beer} \sim \text{Year} \mid \text{Country}$
Columns in parallel	$\text{Beer+Wine+Spirit} \sim \text{Year},$ $\text{outer}=\text{FALSE}$	$\text{Beer+Wine+Spirit} \sim \text{Year},$ $\text{outer}=\text{TRUE}$

3 Lattice Parameters and Graphics Features

Lattice parameter settings

1. Changes to the defaults for points and lines are most easily made using the function `simpleTheme()` (in recent versions of *lattice*).
2. Axis, axis tick, tick label and axis label parameters are conveniently set using the parameter scales in the function call.
3. Lattice objects can be created, then updated to incorporate changes to parameter settings.
4. Note also the parameters `aspect` (aspect ratio) and `layout` (`# rows × # columns × # pages`).

3.1 Point, line and fill color settings

Lattice point, line & related settings

First use `simpleTheme()` to create a “theme” with the new settings:

```
miscSettings <- simpleTheme(pch = c(1,3,4), cex=1.25)
```

Alternatives are then:

- (i) Supply the “theme” to `par.settings` in the function call.

```
xyplot(Beer+Spirit+Wine ~ Year | Country, outer=FALSE,
       auto.key=list(columns=3), data=grog,
       par.settings=miscSettings)
```

[This stores the settings with the object. These stored settings over-ride the global settings at the time of printing.]

- (ii) Supply the “theme” to `trellis.par.set()`, prior to plotting:

```
trellis.par.set(miscSettings)
xyplot(Beer+Spirit+Wine ~ Year | Country, outer=FALSE,
       auto.key=list(columns=3), data=grog)
```

[Makes the change globally, until a new trellis device is opened]

The function `simpleTheme()` creates a “theme”, i.e., a list of parameter settings, in a form that can be supplied: (i) in the argument `par.settings` in the graphics function call; or (ii) in the argument `theme` in a call to `trellis.par.set()`, prior to calling the graphics function; or (iii) in the argument `theme` to `trellis.device()`.

Note the use of the function `trellis.device()` to open a new graphics device. By default, it has `retain=FALSE`, so that parameters are reset to their defaults for the relevant device.

The following changes the plotting symbols to symbols 1, 3 and 4, as in Figure 1. It also increases the size of points by 25%:

```
miscSettings <- simpleTheme(pch = c(1,3,4), cex=1.25)
```

```
xyplot(Beer+Spirit+Wine ~ Year | Country, outer=FALSE,
       auto.key=list(columns=3), data=grog, ylim=c(0,5.5),
       par.settings=miscSettings,
       xlab="", ylab="Alcohol consumption (per person)")
```

Where there are a small number of points, it can be helpful to show them as large solid dots. The following affects all subsequent plots, until changed or until a new device is opened:

```
trellis.par.set(simpleTheme(pch = 16, cex=2))
```

When there are a large number of points, it may be helpful to set the background transparency `alpha` (c.f., also, `alpha.points` and `alpha.line`) to a value less than 1, so that regions where there are many overlapping points can be readily identified.

Where changes go beyond what `simpleTheme()` allows, it is necessary to know the names under which settings are stored. To inspect these, type:

```
> names(trellis.par.get())
 [1] "fontsize"          "background"        "clip"
 . . .
[28] "par.sub.text"
```

The settings that are of interest can then be inspected individually. Section 14.12 of DAAGUR has brief details. For a visual display that shows default settings for points, lines and fill colour, try the following:

```
trellis.device(color=FALSE)
show.settings()
trellis.device(color=TRUE)
show.settings()
```

The following sets the fontsize. Note that there are separate settings for text and symbols:

```
trellis.par.set(list(fontsize = list(text = 7, points = 4)))
```

3.2 Parameters that affect axes, tick marks, and axis labels

Axis, tick, tick label and axis label settings – the scales argument

- Tick positions and tick labels:

```
jobplot <- xyplot(Ontario+BC ~ Date, data=jobs)
## Half-length ticks, each quarter, Label years, Add key
tpos <- seq(from=95, by=0.25, to=97)
tlabs <- rep(c("Jan95", "", "Jan96", "", "Jan97"),
            c(1,3,1,3,1))
update(jobplot, auto.key=list(columns=2), xlab="",
       scales=list(tck=0.5, x=list(at=tpos, labels=tlabs)))
```

A logarithmic scale, and/or relation="sliced"

- Use a Logarithmic scale (here, natural logarithmic)

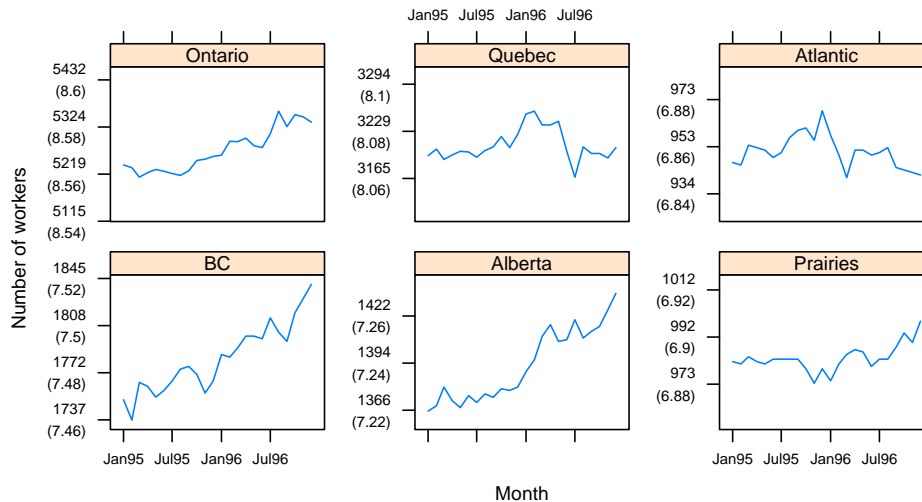
```
logplot <- xyplot(Ontario+BC ~ Date, data=jobs, outer=TRUE,
                 xlab="", scales=list(y=list(log="e")))
```

- Slice the scale

```
update(logplot, scales=list(y=list(relation="sliced")))
```

Scales may have `relation="fixed"`, or `relation="sliced"`, or `relation="free"`

Figure 2: Labor force numbers (1000s) for various regions of Canada. Labels on the vertical axis show both numbers and \log_e of numbers. Distances between ticks are 0.02 on the \log_e scale, i.e., a change of almost exactly 2%.



The data frame `jobs` (`DAAG`) has numbers in the Canadian labour force, for each of six different regions, by month over January 1995 to December 1996. The regions appear as columns of the data frame `jobs`. The following plots these in parallel, on the one panel:

```
xyplot(Ontario+Quebec+BC+Alberta+Prairies+Atlantic ~ Date, data=jobs,
       ylab="Number of jobs", type="b", outer=FALSE,
       auto.key=list(space="right", lines=TRUE))
```

To get a good visual indication of relative changes over time, however, a “sliced” logarithmic scale is needed. The following saves for later enhancement a simplified version of the plot shown in Figure 2, giving it the name `jobs.xyplot`:

```
## Save the graphics object, for later updating
jobs.xyplot <-
  xyplot(Ontario+Quebec+BC+Alberta+Prairies+Atlantic ~ Date,
         data=jobs, type="b", layout=c(3,2), ylab="Number of jobs",
         scales=list(y=list(relation="sliced", log=TRUE)), outer=TRUE)
```

The sliced scale gives each panel the slice of the scale that is needed for points in that panel. A logarithmic scale makes equal relative changes equidistant.

The labeling can be greatly improved. In Figure 2 the y -axis label shows number of jobs, with the logarithms of the numbers given in parentheses. Additionally, dates of the form `Jan95` label the x -axis. In the following code, observe the use of `update()` to specify tick positions and tick labels for the graphics object `jobs.xyplot`.

```
ylabpos <- exp(pretty(log(unlist(jobs[, -7])), 100))
ylabpos <- paste(round(ylabpos), "\n(", log(ylabpos), ")", sep="")
## Create a date object 'startofmonth'; use this instead of 'Date'
atdates <- seq(from=95, by=0.5, length=5)
datelabs <- format(seq(from=as.Date("1Jan1995", format="%d%b%Y"),
                      by="6 month", length=5), "%b%y")
update(jobs.xyplot, xlab="", between=list(x=0.5, y=1),
```

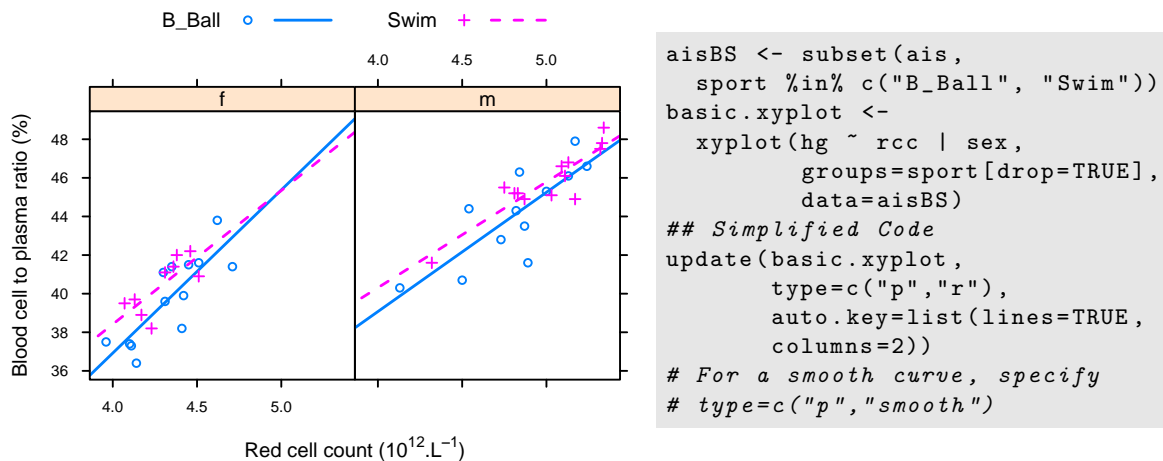


Figure 3: Blood cell to plasma ratio (hc) versus red cell count (rcc), by sex (different panels) and sport (distinguished within each panel). The argument `type=c("p", "r")` displays both points ("p") and regression lines ("r").

```

scales=list(x=list(at=atdates, labels=datelabs),
  y=list(at=ylabelpos, labels=ylabels), tck=0.6)

```

Notice the use of `between=list(x=0.5, y=1)` to add horizontal and vertical space between the panels. The addition of extra vertical space ensures that the tick labels do not overlap. Specifying `tck=0.6` reduces the length of axis ticks to 60% of the default. This can be a vector of length 2.

3.3 A further example

The dataset `ais` (*DAAG*) has a number of physical and biological measurements on 202 athletes at the Australian Institute of Sport. See `help(ais)` for details of the measurements, which include a variety of blood cell counts. Data are for ten different sports. Here is a breakdown of numbers:

```

> with(ais, table(sex, sport))
  sport
sex B_Ball Field Gym Netball Row Swim T_400m T_Sprnt Tennis W_Polo
f   13     7   4     23   22   9     11     4     7     0
m   12    12   0     0   15  13     18    11     4    17

```

The data were collected with the aim of examining possible differences in blood characteristics, between athletes in endurance-related events and those in power-related events.

Figure 3 plots blood cell to plasma ratio (%) against red cell count, for two sports only. The two sports appear within the one panel, distinguished by different symbols and/or colours. Females and males appear in separate panels. Figure 3 shows a suitable plot, adding also regression lines. Again, note the creation of an initial basic plot, which is then updated. Code is:

```

basic1 <- xyplot(hc ~ rcc | sex, groups=sport[drop=TRUE],
  data=subset(ais, sport %in% c("B_Ball", "Swim")),
  xlab="", ylab="Blood cell to plasma ratio (%)")
basic2 <- update(basic1, par.settings=simpleTheme(pch = c(1,3),
  scales=list(tck=0.5), lty=1:2, lwd=1.5))
update(basic2, type=c("p", "r"), auto.key=list(columns=2, lines=TRUE),
  xlab=expression("Red cell count (10^{12}*"."*L^{-1}*")"))

```

Again, there are details that require explanation:

- `type=c("p","r")` gives both points and fitted regression lines.
- In `groups=sport[drop=TRUE]`, the `drop=TRUE` is optional. If not included there will be one legend item for each of the 10 sports. Subsetting a factor leaves the levels attribute unchanged, even if some of the levels are no longer present in the data.
- As in base graphics, graphical annotation (tick labels, axis labels, labels on points, etc.) can be given using the function `expression()`. In this context, “expression” is broadly defined; thus expressions can have terms that are character strings. See `help(plotmath)`.

3.4 Keys – `auto.key`, `key` & legend

The argument `auto.key=TRUE` gives a basic key that identifies colors, plotting symbols and names for the groups. For greater flexibility, `auto.key` can be a list. Settings that are often useful are:

- `points`, `lines`: in each case set to `TRUE` or `FALSE`.
- `columns`: number of columns of keys.
- `x` and `y`, which are coordinates with respect to the whole display area. Use these with `corner`, which is one of `c(0,0)` (bottom left corner of legend), `c(1,0)`, `c(1,1)` and `c(0,1)`.
- `space`: one of `"top"`, `"bottom"`, `"left"`, `"right"`.

Use of `auto.key` sets up the call `key=simpleKey()`. If not otherwise specified, colors, plotting symbols, and line type use the current trellis settings for the device. Unless text is supplied as a parameter, `levels(groups)` provides the legends.

When updating, use `legend=NULL` to remove an existing key, prior to adding a different key.

4 Panel Functions and Interaction with Plots

Further flexibility is added, in the creation of plots, by the use of a user’s own panel functions. A number of panel functions are provided that can be incorporated. A further possibility is interaction with the panels, or other graphics features, of a graph that has just been printed.

4.1 Panel functions

Each lattice command that creates a graph has its own panel function. Thus `xyplot()` has the panel function `panel.xyplot()`. The following are equivalent:

```
xyplot(ACT ~ year, data=austpop)
xyplot(ACT ~ year, data=austpop, panel=panel.xyplot)
```

The user’s own function can be substituted for `panel.xyplot()`. Panel functions that may be used, either in combination with functions such as `panel.xyplot()` or separately, include:

- `panel.points`, `panel.lines()` and a number of other such functions that are documented on the same help page as `panel.points`);
- `panel.abline()`, `panel.curve()`, `panel.rug()`, `panel.average()` and a number of other functions that are documented on the same help page as `panel.abline()`.

The following gives a version of Figure 3 in which the lines for the two sports are parallel:

```
xyplot(hg ~ rcc | sex, groups=sport[drop=TRUE], data=aisBS,
       auto.key=list(lines=TRUE, columns=2), aspect=1,
       strip=strip.custom(factor.levels=c("Female","Male")),
       # In place of level names c("f", "m"), use c("Female", "Male")
       panel=function(x, y, groups, subscripts, ...){
```



```

panel.superpose(x,y, groups=groups,
                subscripts=subscripts, ...)
b <- coef(lm(y ~ groups[subscripts] + x))
lcol <- trellis.par.get()$superpose.line$lcol
lty <- trellis.par.get()$superpose.line$lty
panel.abline(b[1], b[3], col=lcol[1], lty=lty[1])
panel.abline(b[1]+b[2], b[3], col=lcol[2],
            lty=lty[2])
})

```

When there are groups within panels, `panel.xyplot()` calls `panel.superpose()`. The customized panel function has the structure:

```

panel=function(x, y, groups, subscripts, ...){
  panel.superpose(x,y, groups=groups,
                 subscripts=subscripts, ...)
  . . . .
  panel.abline(b[1], b[3], col=lcol[1], lty=lty[1])
  panel.abline(b[1]+b[2], b[3], col=lcol[2],
              lty=lty[2])
})

```

As the function `panel.superpose()` lacks an option to fit and display multiple parallel lines, the user must supply the needed code. The following calculates the regression slope estimates:

```

b <- coef(lm(y ~ groups[subscripts] + x))
# NB: groups (but not x and y) must be subscripted

```

The lines are then drawn one at a time, taking care that the line settings agree with those that will appear in the key.

A further enhancement (omitted from the above code) adds axis labels, using an expression for the x -axis label.

```

xlab=expression("Red cell count (1012*L-1)")
ylab="Blood cell to plasma ratio (%)"

```

4.2 Interaction with Lattice Plots

Focusing and unfocusing:

- Following the plot, call `trellis.focus()`.
- In a multi-panel display, click on a panel to select it.
- Use functions such as `panel.points()`, `panel.text()`, `panel.abline()`, `panel.identify()`.
- Call `trellis.focus()`, as needed, to switch panels.
- When finished, call `trellis.unfocus()`.

For non-interactive use of `trellis.focus()`, turn off highlighting, i.e., the call becomes `trellis.focus(highlight=FALSE)`.

Use the call `trellis.panelArgs()` to identify the arguments that are available to panel functions following a call to `trellis.focus()`.

Viewports:

A lattice plot is made up of a number of “viewports”. In the call to `trellis.focus()`, the default is `name="panel"`.

Other choices of `name` include `"panel"`, `"strip"`, `name="legend"` and `"toplevel"`. For `name="legend"`; `side` should be indicated.

Here is an example of the interactive labeling of points:

```
xyplot(log(Time) ~ log(Distance), groups=roadORtrack,
       data=worldRecords)
trellis.focus()
## Now click (maybe twice) on a panel
panel.identify(labels=worldRecords$Place)
## Click near to points that should be labeled
## Right click to terminate
trellis.unfocus()
```

Lattice is built on top of the `grid` package. This implements viewports, which are arbitrary rectangular regions within which plotting takes place. In the course of plotting, the focus moves from one viewport to the next, as needed to build up the plot.

The function `trellis.focus()` can, once the printing is complete, be used to restore focus to a viewpoint within one of the current panels, or to the whole panel. Common choices for the parameter `name` are `"panel"` and `"strip"`, with `column` and `row` (by default, counting from the bottom up) identifying the column and row in the layout. For `name="legend"`; `side` must be indicated. The argument `name="toplevel"` gives access to the rectangular region within which the panels are placed.

For interactive use, the function `trellis.focus()` can be called without parameters. In a single panel display, this highlights the panel. In a multi-panel display, clicking on a panel will select it. The function `trellis.unfocus()` removes the highlighting and makes `"toplevel"` the current viewport.

Once the focus is on a panel, the user has access to the functions that were noted in the previous subsection, and to many others besides.

The functions `trellis.focus()` and `trellis.unfocus()` can be used in a non-interactive mode. The following prints the stripplot and a boxplot objects created in Subsection 5.1 one under the other on the same graphics page. Following the printing of each plot, the focus is placed on the top level viewport, and the function `grid.text()` (from the `grid` package) used to add a label:

```
library(DAAG)
library(grid)
## Positioning will be (xmin=0, ymin=0.46, xmax=1, ymax=1)
print(update(cuckoostrip, xlab=""), position=c(0, .46, 1, 1))
trellis.focus("toplevel", highlight=FALSE)
grid.text("A", x=0.05, y=0.935, gp=gpar(cex=1.15))
trellis.unfocus()
print(cuckoobox, position=c(0, 0, 1, 0.54), newpage=FALSE)
trellis.focus("toplevel", highlight=FALSE)
grid.text("B", x=0.05, y=0.935, gp=gpar(cex=1.15))
trellis.unfocus()
```

Observe that the rectangular regions on which the objects are printed have been chosen so that they overlap somewhat, reducing the space between.

5 Displays of Distributions

5.1 Stripplots, dotplots and boxplots

Because the syntax for `stripplot()` and `boxplot()` are very similar, we demonstrate suitable code side by side. (The function `dotplot()` is very similar to `stripplot()`, with differences that are mainly cosmetic.) The following code creates a stripplot object and a boxplot object, for the cuckoos data (from *DAAG*):⁴

```
cuckoostrip <- stripplot(species ~ length, aspect=0.5,
                        xlab="Cuckoo egg length (mm)", data=cuckoos)
cuckoobox <- bwplot(species ~ length, aspect=0.5,
                   data=cuckoos, xlab="Cuckoo egg length (mm)")
```

The `aspect` argument determines the ratio of distance in the y-direction to distance in the x-direction. The following demonstrates the use of `dotplot()`:

```
dotplot(variety ~ yield | site, data = barley, groups = year,
        xlab = "Barley Yield (bushels/acre) ", ylab = NULL,
        layout = c(1, 6), aspect = 0.5,
        auto.key=list(labels=levels(barley$year), space = "right"))
```

Try stretching the plot vertically so that the labels do not overlap.

The argument `type="h"` gives a line from the origin to the point. Both a line and a point may be given. This can be used to quite striking effect, as in the following:⁵

```
deathrate <- c(40.7, 36,27,30.5,27.6,83.5)
hosp <- c("Cliniques of Vienna (1834-63)\n(> 2000 cases pa)",
        "Enfans Trouves at Petersburg\n(1845-59, 1000-2000 cases pa)",
        "Pesth (500-1000 cases pa)",
        "Edinburgh (200-500 cases pa)",
        "Frankfort (100-200 cases pa)", "Lund (< 100 cases pa)")
hosp <- factor(hosp, levels=hosp[order(deathrate)])
dotplot(hosp ~ deathrate, xlim=c(0,110), cex=1.5,
        scale=list(cex=1.25), type=c("h","p"),
        xlab=list("Death rate per 1000 ", cex=1.5),
        sub="From Nightingale (1871) - data from Dr Le Fort")
```

5.2 Lattice Style Density Plots

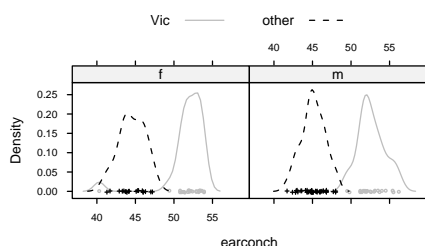


Figure 4: Lattice style density plot comparing possum earconch measurements, separately for males and females, between Victorian and other populations. Observe that the scatter of data values is shown along the horizontal axis.

Here is a density plot (Figure 4), for data from the possum data set (*DAAG*), that compares sexes and Vic/other populations.

```
## For slightly improved labeling, precede with:
```

```
4levels(cuckoos$species) <- sub(".", " ", levels(cuckoos$species), fixed=T)
```

```
5Data are from Nightingale, F. (1871): Notes on Lying-in Institutions.
```

```
densityplot(~ earconch | sex, groups=Pop, data=possum,
            par.settings=simpleTheme(col=c("gray","black"),
            auto.key=list(columns=2))
```

Where `densityplot()` (and `histogram()`) have a formula as argument, a name is not allowed on the left of the `~` symbol. For `histogram()`, the `groups` argument is not available.

6 Lattice graphics functions – Further Points

6.1 Help on lattice functions

For an overview, type `help(Lattice)`. For help on the graphical parameters used by lattice functions, see `help(trellis.par.set)` and `help(simpleTheme)`. For other settings, see `help(lattice.options)`.

Several of the help pages for lattice functions are shared between more than one function. For example, `xyplot()`, `dotplot()`, `barchart()`, `stripplot()` and `bwplot()` share the same help page. As a consequence, typing `example(bwplot)` has the same effect as typing `example(xyplot)`.

6.2 Selected Lattice Functions

Functions that have already been demonstrated are `xyplot()`, `stripplot()`, `dotplot()`, `densityplot()` and `bwplot()`. Other "high level" functions include:

```
barchart(character ~ numeric,..)
histogram( ~ numeric,..)      # NB: does not accept groups parameter
densityplot( ~ numeric,..)    # Density plot; does allow groups
bwplot(factor ~ numeric,..)   # Box and whisker plot
qqmath(factor ~ numeric,..)   # normal probability plots
## Bivariate
qq(factor ~ numeric, ...)     # comparing two empirical distributions
                                # (Two factor levels identify the 2 distns)

## Multivariate
cloud(numeric ~ numeric * numeric, ...) # 3D surface
wireframe(numeric ~ numeric * numeric, ...) # 3D scatterplot
levelplot(numeric ~ numeric * numeric, ...) # cf image() in base graphics
contourplot(numeric ~ numeric * numeric, ...) # contour plot
## Hypervariate
splom( ~ dataframe,..)        # Scatterplot matrix
parallel( ~ dataframe,..)     # Parallel coordinate plots
## Miscellaneous
rfs()                          # Residual and fitted value plot (also see 'oneway')
tmd()                           # Tukey Mean-Difference plot
```

In each instance, users can add conditioning variables.

Reference

Sarkar, D. 2008. *Lattice. Multivariate Data Visualization with R*. Springer.
This is the definitive account of the *lattice* implementation of trellis graphics. It merits careful study.

Index of Functions

as.Date, 6

barchart, 12
boxplot, 11
bwplot, 11, 12

c, 3–8, 10–12
cloud, 12
coef, 9
contourplot, 12

demo, 1
densityplot, 12
dotplot, 11, 12

example, 12
exp, 6
expression, 7–9

factor, 11
format, 6
function, 8, 9

gpar, 10
grid.text, 10

help, 7, 8, 12
histogram, 12

image, 12
invisible, 2

levelplot, 12
levels, 8, 11
library, 1, 2, 10
list, 3–8, 11, 12
lm, 9
log, 6, 10

names, 5

order, 11

panel.abline, 8, 9
panel.average, 8
panel.curve, 8
panel.identify, 9, 10
panel.lines, 8
panel.points, 9
panel.rug, 8
panel.superpose, 9
panel.text, 9
panel.xyplot, 8, 9

parallel, 12
paste, 6
plot, 2
pretty, 6
print, 1–3, 10

qq, 12
qqmath, 12

rep, 5
rfs, 12
round, 6

seq, 5, 6
show.settings, 5
simpleKey, 8
simpleTheme, 3–5, 7, 12
splom, 12
strip.custom, 8
stripplot, 11, 12
sub, 11
subset, 7

table, 7
tmd, 12
trellis.device, 4, 5
trellis.focus, 9, 10
trellis.panelArgs, 9
trellis.par.get, 5, 9
trellis.par.set, 4, 5
trellis.unfocus, 9, 10

unlist, 6
update, 3, 5–7, 10

wireframe, 12
with, 7

xyplot, 2–8, 10, 12