



DSC 2003 Working Papers
(Draft Versions)

<http://www.ci.tuwien.ac.at/Conferences/DSC-2003/>

Some Notes on Lattice

Deepayan Sarkar
Department of Statistics
University of Wisconsin – Madison
deepayan@stat.wisc.edu

Abstract

Trellis Graphics is implemented in S-PLUS using the traditional S graphics engine. In contrast, Lattice, the implementation of Trellis Graphics in R, uses Grid as the underlying graphics mechanism. The decision to use Grid has obvious disadvantages, stemming from the fact that pre-existing graphics code cannot be used unchanged. We shall discuss some of the benefits of using Grid.

Other important differences between Trellis and Lattice are in the API. Although the Lattice API is largely based on S-PLUS, it also sports several extensions, as well as some new features. We discuss the motivation for some of these differences. An elementary knowledge of Trellis Graphics is assumed.

1 Introduction

Trellis Graphics, a prominent feature of S-PLUS, had been missing from R until recently. One of the reasons for this was that the standard R graphics system is not very well suited to producing Trellis-like graphs. The Grid graphics package (Murrell, 2002) provides several powerful features which make a clean implementation of Trellis much easier. Here, we discuss the role played by Grid, in particular “grob”s (graphical objects) and the Grid unit mechanism, in implementing Lattice.

The high level functions in Lattice are modeled on the ones in S-PLUS. These are described quite thoroughly in several papers available at Bell Labs’ web site (e.g., Becker et al. (1996)). Although the use of Grid makes some differences unavoidable, the descriptions of Trellis in S-PLUS are applicable for the most part to Lattice as well.

But even apart from the technical differences implied by the use of Grid, there are also some design aspects in which Lattice consciously deviates from Trellis, mostly in the form of extensions to the Trellis API. These could be useful to users coding exclusively for R, but will often lead to incompatibilities with Trellis. We give examples and brief justification for some of these extensions.

2 The role of Grid in Lattice

Trellis Graphics is different from traditional S graphics in many ways. It uses a formula interface to specify the form of a plot, it has a completely new system for customizing parameters, and of course, it has in-built the functionality to handle an arbitrary number of conditioning variables. But none of this is any more difficult to implement than it should be with traditional S graphics.

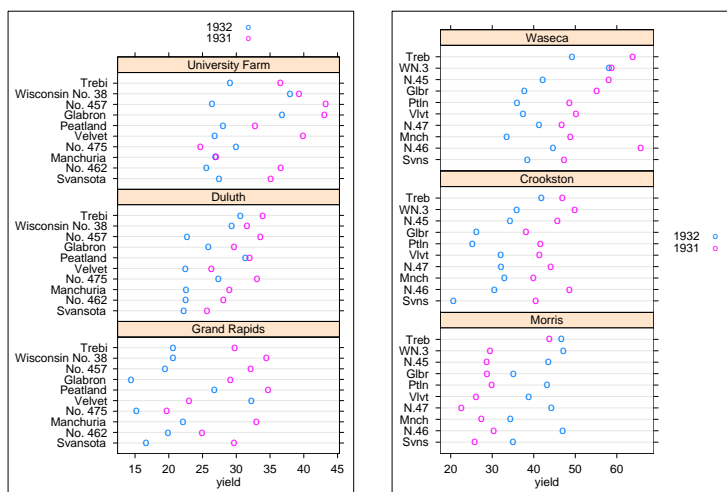


Figure 1: Two lattice plots, produced by essentially identical calls, except for different placement of the key, and abbreviated labels for the y-axis in the right hand plot.

However, Trellis Graphics is different from traditional S graphics in another aspect: it always tries to make as much use as possible of the available plotting space. Space will be allocated for a key (a.k.a. legend) or a title only if there is one. Exactly that much space will be used for tick marks and axis annotation as is required. Consider the two plots in Figure 1, representing two halves of the barley data set (Cleveland, 1993). Both are drawn inside Grid viewports of identical dimension, with almost identical calls. They differ only in the location of the key, and in that the Y-axis labels have been abbreviated in the second plot. Notice that in both cases, every component of the graph except the actual Dot Plots inside the panels uses the minimum space necessary, allowing the panels to expand to use the

remaining space. This is something that would be very difficult to get right without Grid.

2.1 Why is this difficult to implement ?

When using the traditional S graphics engine, a graphics device (conceptually) consists of a plotting region, along with figure margins used for the title, axis labels, annotations, etc. The dimensions of these margins can be controlled via the `par()` function. In the case of Trellis Graphics, this would need to be adjusted dynamically by the plotting function. A more sophisticated alternative, more suited to the other demands of Trellis Graphics, is the `layout()` function. Whatever the mechanism, to dynamically adjust the dimensions, we need to determine those dimensions in a form that would be understood by `par` or `layout`. Unfortunately, the figure margins can be only specified either as the number of lines of text or inches in `par`, and in centimetres in `layout`.

Consider the problem of determining the amount of space required for the key. The key itself can be quite sophisticated. It could consist of an arbitrary combination of columns of text, lines, rectangles or points, possibly divided into blocks. The width of each of these columns can be controlled separately, and the widths and heights of the text components of course would depend on the text itself. Further complications arise when these texts are expressions. Technically, of course, it is possible to obtain the physical dimensions of each of these components. The `strwidth()` and `strheight()` functions, for instance, can return widths and heights in inches. However, dealing with these would clearly be quite messy.

2.2 How does Grid help ?

Grid allows a tremendously flexible mechanism for specifying lengths. To see how this becomes useful, let's delve a little into how Lattice plots are produced.

A Lattice plot is always drawn inside a grid viewport, which can be an arbitrary rectangular region on the plotting device. Once the layout (arrangement of the panels) is decided, the first thing that `print.trellis` (the Lattice function responsible for the actual plotting) does is to create a layout using `grid.layout`. This layout is at the heart of the actual display, and consists of rows and columns that correspond to different components of the graph.

To take a simple case, consider Figure 2. A simplified description of the rows in the layout that created this plot is given in Table 1.

Note that several rows are for components that are actually missing from the plot, the heights for these set to 0 “lines”, and so do not occupy any space. There are rows for the key both at the top and the bottom, but only the top one is used, for which the height is specified simply as 1 “grobheight” – i.e., just enough space to hold the ‘grob’ (Grid graphics object) specified as ‘data’. In this case, this ‘grob’ is produced by an independent function called `draw.key()` which produces the key. Thus, the Lattice code never has to explicitly deal with the physical dimensions of the key.

```

> data(iris)
> print(xyplot(Sepal.Length ~ Petal.Length, data = iris,
+   groups = Species, auto.key = list(columns = 3),
+   main = list("Iris Data", cex = 1.2), xlab = "Petal\nLength",
+   ylab = "Sepal\nLength"))

```

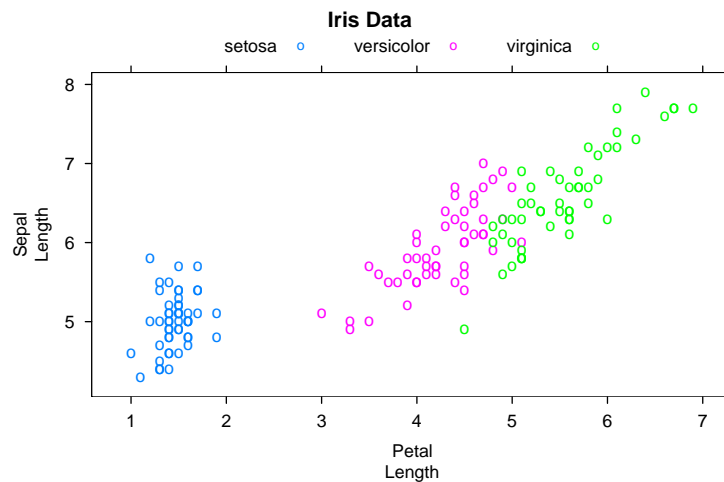


Figure 2: A simple plot using the Iris Data.

Component	height	units	data
Main Label	1.2	"strheight"	"Iris Data"
Key	1	"grobheight"	'Grob' from draw.key()
Tick labels	0	"lines"	
Ticks marks	.5	"lines"	
Panel	1	"null"	
Tick marks	.5	"lines"	
Tick labels			<i>See text</i>
X Label	1	"strheight"	"Petal\nLength"
Key	0	"lines"	
Sub Label	0	"lines"	

Table 1: A slightly simplified description of the rows that go into the Grid layout for Figure 2, top to bottom. The height of the tick labels is more complicated, being the maximum of several "strheight" units.

The `draw.key()` function itself uses another powerful feature of Grid – the ability of a “grob” to pack within itself several smaller “grob”s. The “grob” that is finally returned by `draw.key()` is initially created as a “frame” (`grid.frame()`) object with a specified layout. Smaller “grob”s are then created – lines, rectangles, text, etc – and placed inside the frame (`grid.place()`).

`xlab`, the label for the x variable, is an example of how the “strheight” unit is used. In this case, this is a string containing a newline escape sequence. In general it could be a more complicated mathematical expression. `print.trellis()` doesn’t need to care what it is and how many inches it occupies, it simply delegates that problem to Grid.

2.3 Example: A rotated histogram

The decision to use Grid brings with it the ability to do some useful things that would not have been otherwise possible. Although it’s not very common, one might want to obtain a side-up histogram. Rather than writing a new function or modifying an existing one, simply rotating the viewport in which the plot is made produces the required plot (Figure 3)

```
> push.viewport(viewport(w = 2/3, angle = -90))
> print(histogram(rnorm(500), ylab = NULL, xlab = NULL,
+   scales = list(x = list(rot = 90), y = list(rot = 90),
+   alternating = 2))), newpage = FALSE)
> invisible(pop.viewport())
> grid.text(lab = "i.i.d. Normal Variates", rot = 90,
+   vp = viewport(x = 1/6))
```

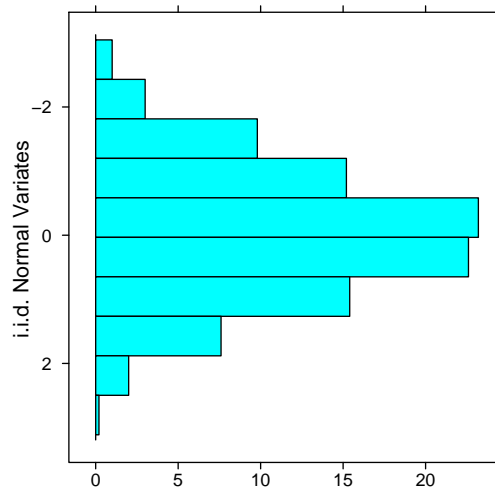


Figure 3: Rotated histogram of a Normal random sample

2.4 Example: Fixed scale plots

I recently had an e-mail exchange with Federico Calboli, who was trying to produce a LOD scores plot for QTL mapping data, similar to those in [Macdonald and Goldstein \(1999\)](#). The overall figure consists of 3 individual plots, each with one panel. The catch was that the 3 plots had different ranges on the X-axis, but were to have the same scale — i.e., one inch should correspond to the same increment on each plot. This was easy to achieve approximately, by placing the plots in regions that had width proportional to the x ranges. However, it was impossible to exactly account for the labels and tick marks.

The solution was a simple extension to `print.trellis()`. In the layout for a Lattice plot, the rows and columns that correspond to the panels are typically assigned “null” units, which allows them to expand to fill the plotting regions. This is easily generalised to allow the user to specify these dimensions explicitly, using the grid unit system. Figure 4 was obtained by specifying the widths of each panel in millimeters, proportional to the X-axis limits.

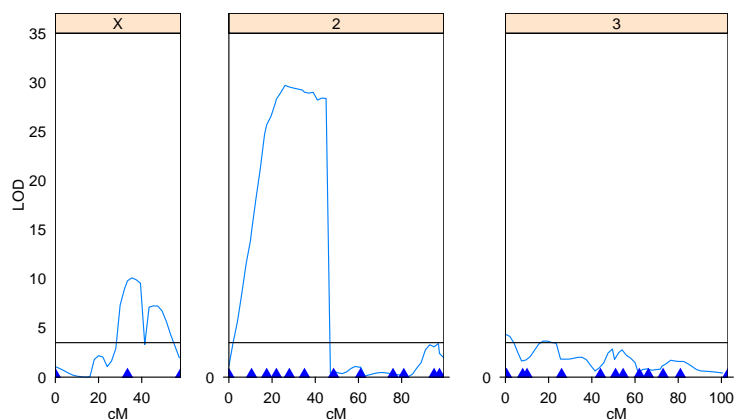


Figure 4: LOD scores plot: the 3 plots correspond to 3 chromosomes of different lengths. The panel widths are specified explicitly, proportional to the X-axis limits, so that all 3 plots have the same scale. (This plot uses dummy data since the original work is still unpublished.)

3 Some features of the Lattice API

3.1 Grouped Displays

Trellis Graphics in S-PLUS has all the basic support necessary for grouped displays. However, it has only one panel function — `panel.superpose()` — that allows grouped displays, and even that has to be explicitly specified when required. Lattice takes the approach that whenever the `groups` argument is specified in a call, the

display should be grouped by default, provided grouping makes sense in the context of that plot. Such functions include `xyplot`, `dotplot`, `stripplot`, `barchart`, `densityplot`, `cloud`, `wireframe` and `splom`. Figure 5 shows an example of this in a `barchart` call, and Figure 6 in a `wireframe` plot.

```
> print(barchart(yield ~ variety | year, data = barley,
+   groups = site, stack = TRUE, reference = TRUE,
+   ylab = "Barley Yield", auto.key = list(rect = TRUE,
+   points = FALSE, columns = 3), scales = list(x = list(rot = 45))))
```

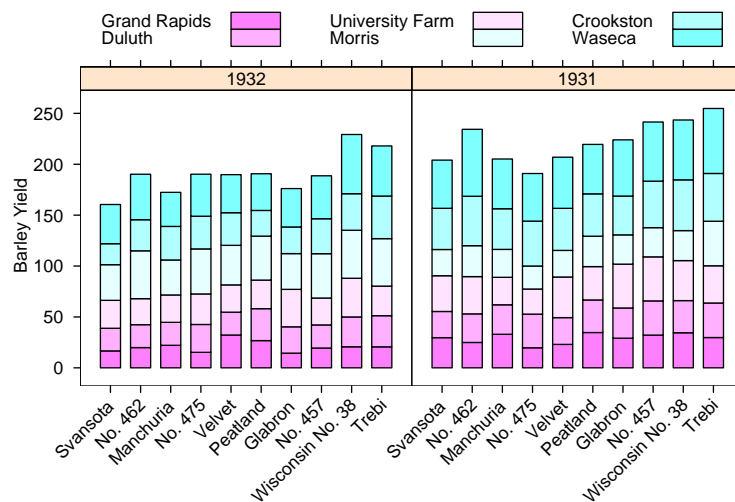


Figure 5: Stacked barchart for the barley data

One problem with the Trellis API for grouped displays is that it requires the data frame to be in the “long” format (to borrow terminology from the `reshape` function). Suppose we want Kernel Density Estimates for the Iris data, with separate panels for each Species. It is easy to obtain this for, say, the Sepal Length, with a call like

```
densityplot( ~ Sepal.Length | Species, iris)
```

But suppose we want density estimates for all four variables in each panel, differentiated by color or line type. `densityplot` can produce grouped displays, so this is not too difficult, but it requires the four data vectors to be concatenated into one long vector, and the creation of a new ‘grouping variable’. This can all be done with `reshape()`, but that is at best inconvenient, and at worst obscure option. Lattice extends the Trellis `formula` interface to make this easier to code — by allowing more than one variable to be specified on either side of the `~` in the formula. Figure 7 illustrates the usage for the Iris example above.

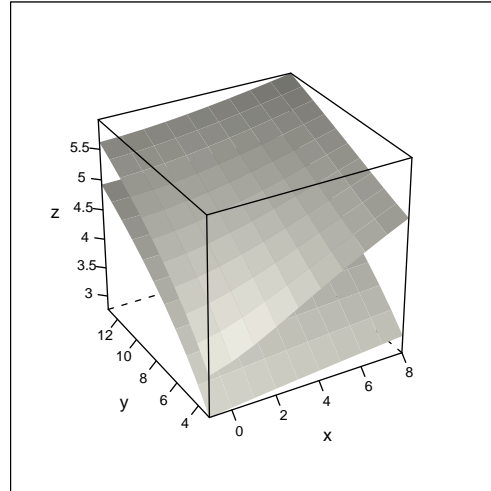


Figure 6: A grouped Wireframe plot (this works only with non-intersecting surfaces, and even then not perfectly, but it can be useful for simple applications).

```
> print(densityplot(~Sepal.Length + Sepal.Width +
+   Petal.Length + Petal.Width | Species, iris,
+   allow.multiple = TRUE, plot.points = FALSE,
+   xlab = "Measurements on 3 species of Iris",
+   scales = list(alternating = FALSE), auto.key = list(line = TRUE,
+   point = FALSE, columns = 4)))
```

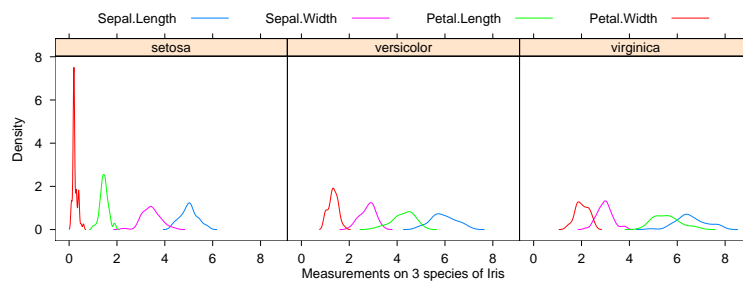


Figure 7: A grouped densityplot. The grouping is specified not in the usual way, but rather as multiple variables in the right hand side of the formula. Note that this needs the `allow.multiple` flag to be set, to ensure S compatibility

3.2 factors in the Trellis formula

Trellis is picky about what sort of variables it allows in the formula. `bwplot`, `dotplot`, `barchart` and `stripplot`, for example, allow only the `y` variable to be a factor. Lattice allows `x` to be the factor, which in fact enables the creation of the more conventional form of Box Plots and Bar Charts. A typical problem with such plots is that the tick labels (which for factors are its levels) are long and often overlap — Lattice deals with this by allowing the labels to be rotated and/or abbreviated through the `scales` argument (see Figure 5 for an example).

Most other Trellis functions, including `xyplot`, `levelplot` and `cloud` require all variables to be numeric. While this is consistent with the nature of these plots, it is often appropriate to use categorical variables there, albeit after coercing them to be numeric. However, we would still want the levels of the factor to provide tick mark positions and labels. Figure 8 shows what is hopefully a convincing illustration.

```
> print(levelplot(yield ~ year * variety | site,
+   barley, aspect = "xy", shrink = c(0.5, 0.9),
+   main = "Barley Yield (bushels / acre)", col.regions = terrain.colors(50),
+   scales = list(alternating = FALSE)))
```

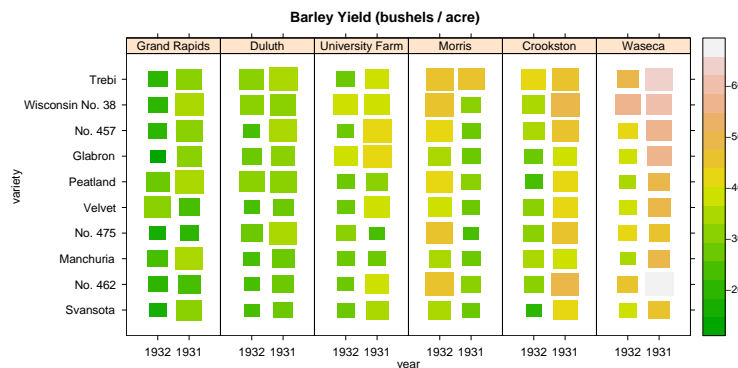


Figure 8: A levelplot (with the size of the rectangles scaled to reflect the value of the `z` variable) of the barley data set, with categorical variables on the X and Y axes. The Morris anomaly (an increase in the yield from 1931 to 1932, the opposite of the trend in all the other sites) is easily identifiable. Figure 1 is a more familiar representation of this data

It is of course possible to do this in Trellis by explicitly coercing the factors to be numeric, but then a lot of extra effort needs to go into manually specifying the tick mark positions and labels to something meaningful.

Lattice takes the approach that any variable can be anything when specified in the formula of a high-level function call, and the nature of this variable will

determine the tick positions and labels. The decision on how to plot this data is delegated to the `panel` function. Unfortunately, this introduces an incompatibility with Trellis (which assumes that `prepanel` and `panel` functions are always passed numeric arguments) — hopefully, the benefits outweigh the problems in this case.

4 Concluding remarks

When it comes to specific code options, this article is light on details. This is intentional, as Lattice is constantly evolving, and details may change over time. The documentation accompanying the package should always be current.

My thanks to Wolfram Fischer and Federico Calboli for several new ideas recently implemented in Lattice.

References

Richard A. Becker, William S. Cleveland, and Ming-Jen Shyu. The visual design and control of trellis display. *Journal of Computational and Graphical Statistics*, 5:123–155, 1996. URL <http://cm.bell-labs.com/cm/ms/departments/sia/project/trellis/display.writing.html>.

William S. Cleveland. *Visualizing data*. Hobart Press, 1993.

Stuart J. Macdonald and David B. Goldstein. A quantitative genetic analysis of male sexual traits distinguishing the sibling species *Drosophila simulans* and *D. sechellia*. *Genetics*, 153:1683–1699, 1999.

Paul Murrell. The grid graphics package. *R News*, 2(2):14–19, 2002.