

Using Layering and Perceptual Grouping In Statistical Graphics

By Dan Carr and Ru Sun

1. Introduction

This paper concerns the use of layering and perceptual grouping to modify the appearance of statistical graphics. Our motivating example is a Trellis Graphics dot plot with strip labels as shown in Figure 1 (see page 29). The object of our attention is the color-filled strip label boxes. Our subjective opinion is that the color-filled boxes are visually dominant, drawing more visual attention than the dot plot panels. Our impression is of looking through the trellis of strip labels to see the dot plots. We conjecture that the intervening foreground strip labels slow visual comparison of dots from different panels in the same column. While many chose the simple solution of turning off the strip label fill, our re-design goal is to move the dot plot panels into the foreground. We do this conceptually by putting dot plots on a thin marble block. This leaves the strip labels back in flatland. While examples similar to Figure 1 provided our motivation, the methods described in our redesign are applicable to a wide class of statistical graphics.

The notions of figure and ground, foreground and background, and layering are not new. Examples of figure and ground reversal have long intrigued lay people as well as researchers in the field of human perception and cognition. Similarly, researchers in cartography have studied the closely related notion of information layering on maps (for discussions of both topics see MacEachren 1995). The notion of map layers is a basic part of geographic information systems. Tufte (1990) devotes a full chapter to layering and separation. Our first goal is to promote the use of existing layering methods to enhance the appearance of statistical graphics.

We live in a competitive world in which appearances are important. Appearance influences our choices in all facets of our visual life from the food we eat, to the clothes we wear, to the partners we choose. Appearance is not necessarily our strongest criterion. When hungry most of us would choose a spotted banana over a flower despite our notion of what is pretty. Nonetheless at some point of other things being roughly equal, appearance can be the deciding criterion. Our opinion is that

statistical community could pay more attention to appearance of graphics. With appearances being roughly equal perhaps the public will choose the graphics with content.

Addressing the appearance of graphics is a bit risky in a profession that strongly focuses on content and perceptual accuracy. For example Cleveland (1993a) presents graphical methods in the context of data examples. This data centric approach is persuasive. He is interested in science and his methods continue to build upon the foundations of perceptual accuracy of extraction promoted by Cleveland and McGill (1984).

Our promotion of methods for the sake of appearance can be viewed with suspicion. We are promoting a three-dimensional appearance and this may seem to run counter to some sound guidance. For example Tufte (1983) disparages area and volume encodings that correspond to the square and cube of variable respectively, and discusses the lie factor. He also discourages display of extra dimensions even when they are constant. The phrase “dimensional puffery” now appears in the human computer interface literature (Card, MacKinlay, and Shneiderman 1999). There is no doubt that frivolous embellishment can detract.

We promote the notion that some embellishment can give the appearance of value added at little perceptual accuracy cost. In particular this paper suggests use of modest 3-D effects to layer information. Further, we consider “perceptual enhancements” that seem inconsistent with common guidance about graphics design. We recommend use of perceptual grouping for its own sake. This runs the risk of people thinking that the groups mean something more. We also connect points to guide the visual flow from point to point and to strengthen the perceptual grouping. Critics will claim that some readers (not themselves) interpret the presence of connection lines as implying linear interpolation, however nonsensical. Nonetheless, this paper suggests that the benefits of layering, perceptual grouping and guidance of visual flow may well outweigh the cost of confusion that can arise until new conventions are established.

The structure of the paper is as follows. Section 2 provides a little background on dot plots, Trellis graphics and strip labels. Section 3 concerns the use of multiple perspectives in plot construction. Section 4 touches on diffuse lighting, shadows, and the appearance of depth. Section 5 revisits labeling issues. Section 6 provides a brief discussion of perceptual grouping and connecting lines. Section 7 concludes by providing links to software and thoughts spawned by Figure 2 (see page 30).

2. Dot plots, Trellis Graphics, and Layering of Strip Labels

Consider dot plots. Every time I (Dan) return to the work of Cleveland and McGill (1984) or Cleveland (1985), I am impressed by the simple elegance of the dot plot concept and the numerous design variations that such simplicity facilitates. My work with federal agencies started with an attempt to promote dot plots. So far most of these efforts have failed. Other than the row-labeled dot plots that I promoted for use in the Atlas of United States Mortality (Pickle et al. 1997), I have yet to see a dot plot in a major federal publication. Perhaps the biggest reason is that dot plots are not a push button option in widely used spreadsheet graphics.

The Trellis graphics in Splus represent a major step forward from the 1970's business graphics that are characteristic of spreadsheet software. The color versions can be quite attractive while conveying information with integrity. The Trellis design provides a general approach to multiway, multipanel and even multipage graphics. This generality necessitated the development of a general approach to labeling. Strip labels emerged as a elegant solution to multiway panel labeling challenge. However, our opinion is that this general labeling solution is not necessarily optimal for special cases.

As indicated in the introduction we think that strip labels between graphics panels impede comparison down the panels within the same column. For this reason, we often prefer some of Cleveland's (1985) earlier dot plot designs that keep the labeling in the margins. Our redesign goal here, however, is to retain the panel strip labels (at least in some form), while attempting to promote comparisons across vertical panels by bringing the panels into the foreground. This is a task in relayering information.

There are different ways to go about relayering information. We conjecture that the strip labels appear in the foreground because the color fill calls attention to areas. It seems that area symbols dominate over line and point symbols. One relayering option simply avoids the color fill. Another option makes the dots even bigger so they effectively become area symbols that compete on an equal footing with strip labels. We think the strip labels provide a visual barrier to vertical visual flow even though they appear in the same visual layer as the dots. Thus we strive to reduce the impact of this visual barrier by using 3-D effects that bring the graphics panels into the foreground.

3. The Marble Block and Perspective Views

As indicated above and shown in Figure 2 (page 30), our simple idea for bringing the graphics panels into the foreground is to place them on thin marble blocks. Simple perspective views that include the block surface and sides provide a 3-D appearance. The details about perspective views can be found in numerous sources such as Foley, van Dam, Feiner, and Hughes (1990), Carr (1994), and Blinn (1996). We want to show the marble blocks in a way that minimizes the area that is consumed by the 3-D construction. Toward this end we digress for a moment.

It is fun to play with perspective views. The impossible tribar of Penrose (1958) and the artwork of M. C. Escher (for example see Ernst 1976) have delighted people around the world. Humans have an amazing ability to maintain a local perspective that is not globally consistent. For many situations local is good enough. For example the construction of the "infinity enlarged" stereo pairs projection (see Carr 1994) effectively moves the center of projection (in this case the mid-point between the eyes) in front of every single point in the plot before projecting the point onto the surface. This provides the key to a space saving view of our marble block.

To save space we position the viewpoint so we see only one block side. As shown in Figure 2, our viewpoint choice is below the center of each block and beyond the block boundary so we see only the lower block side. Maintaining the same relative viewpoint for each block provides the basis for constructing Figure 2. This raises the question of what to do when panels touch left and right. Currently, we move the viewpoint again so the shared edge is right in front of us.

Figure 2 shows our hopping about perspective construction of juxtaposed panels. Some previous examples (see Cleveland 1993b and discussant comments) show the two years of data superimposed on the same panels. For comparing barley yield across years we also prefer superposition to juxtaposition and note that the direct display of differences is helpful. Figure 2 has not been optimized for comparing years.

4. Diffuse lighting, shadows, and the appearance of depth

In Figure 2 we are dealing with a solid block and not a wireframe. Thus we need to pay attention to surface rendering. An actual marble texture would be distracting so we chose to use a plain gray surface. Given the surface, the next step involves the use of lighting models. Lighting can create or enhance the appearance of depth and provide a sense value added relative to flat graphics.




Figure 3 Exploring ambient light and diffuse reflection.

Computer graphic classics (Foley, van Dam, Feiner, and Hughes 1990; and Blinn 1996) describe illumination models and surface rendering. Those interested in rigorous treatment should refer to such sources. In brief, there are four basic facets to illumination models, ambient light, diffuse reflection, atmospheric attenuation, and specular reflection. Ambient light is the minimal light emitted by surfaces where the intensity does not depend on a specific light source. Phong shading (normal vector interpolated shading) and specular reflection are useful in producing 3-D appearance. These have been used effectively in density representations (see Wegman, Carr and Luo 1993). Here we focus on the use of diffuse reflection and shadows.

Given the reflective properties of a surface, the relative intensity of diffuse reflection from a point light source shining on a planar surface depends upon the angle between a line from the light source and the normal to surface. The light intensity (lightness, value or brightness depending on the literature) decreases with the cosine of this angle. In contrast to specular reflection, the diffuse light intensity does not vary with viewer's position (except when a surface becomes hidden). Figure 3 provides a simple illustration of ambient light and diffuse reflection. The lighting can make the figure appear as a truncated 3-D hexagon-based pyramid.

Where should the light source be positioned? The convention established by walking under the sun is that light comes from above. The computer graphics convention for simple single-light-source cases is that the light comes from the upper left. Figure 3 should appear to poke out of the surface toward the reader. Putting the light source at the lower right should make the pyramid poke into the surface. However, our visual system can do interesting things. The perception of directionality, above or below the surface, is fragile, and some readers experience the opposite of what is expected.

The construction for Figure 3 provides a nice review of vector operations and is suitable as a student exercise. Suppose we have a right handed coordinate sys-

tem with the positive z -axis pointing toward the viewer. Suppose the center point of the polyhedron is $(0,0,7)$ and the vertices are located at $(\cos(ai), \sin(ai))$, where $ai = 0, 60, 120, 180, 240,$ and 300 degrees. The consecutive vertices of a triangular face can define the vectors corresponding to two edges. The cross product of these vectors defined the normal to the face. The dot product of the unit length normal and a unit direction vector (say $[-.544,.314,.778]$), from a light source at infinity gives the relative diffuse intensity (di) for each triangular surface. Letting the gray level on a 0 to 1 scale be something like $.2 + .7 * di$ provides plausible shading. This can be viewed as a mixture of ambient light and diffuse reflection. The ambient light provides a minimum value for the surfaces. A rough interpretation is that the ambient intensity is $.1 + .8 * \min(di)$. The other gray levels are lighter due to diffuse reflection from the light source. We have avoided using pure white by choosing $.7$ rather than $.8$. This leaves a little room for specular highlights.

Figure 2 (page 30) conveys depth by 3-D perspective. The diffuse lighting comes into play when we color the quadrilateral below each panel. We do not bother with a precise diffuse reflection calculation but simply use a darker shade of gray.

Similarly we attempt to provide the appearance of grid lines etched into the surface by pairing the light and dark lines. The side of a groove away from the light (right side for vertical lines and bottom side for horizontal lines) should appear lighter since those surface catch more light. Ideally, a 3-D design should construct thin and shallow "v" shaped trenches and use proper diffuse lighting. The dots are more important than the reference lines, so the construction with etched lines is proper in terms of depth. Care must be taken in terms of line width and contrast so the lines are less salient than dots.

We choose to use red dots because the focal length for red makes the dots appear closer. Shadows for the dots provide another way to convey depth and make the dots appear closer. As Blinn (1996) indicates, simple shadow construction is often sufficient to convey depth. Given that the light source is in the upper left shadows of an object appear to the lower right of the object. Drawing a dark gray image to the lower right before drawing the object provides the appearance of depth. The dots in Figure 3 illustrate this.

Lifting the dots off surface in Figure 3 causes some ambiguity. Should the reader judge the center of the dot or the center of its shadow when determining its value? The description here is explicit: the position of

the shadow is altered. However, without an established convention the viewer may reason that shadow on the surface should be judged against grid lines. If shadows are used, the shift should be slight. Our statistical Puritanism may incline us not to use shadows, but we confess that a little bit of shadow is fun.

Note that web page design makes heavy use of lighting models to produce modest 3-D effects. Boxes routinely have light edges on the left and top and dark edges at the right and bottom. Such effects are sometimes applied with little thought. In particular standard tabling methods put 3-D ridges around numbers. This impedes visual flow and deters comparison of numbers. The choice of ridges is consistent with drill-down mentality. We are supposed to be happy just to have found a number. The notion that graphics (include tabular graphics) is about comparison has not yet been reinvented some communities.

5. Labels Revisited

Labeling is perhaps the hardest of design challenges. The labels for barley variety in Figure 2 (page 30) posed a difficult challenge until we figured out that they had to go on the marble block to keep aligned with dots. This leaves the strip labels back in flatland by themselves. This is not so bad since strip labels refer to the whole panel below rather than to specific dots. Bump-mapped text could put the strip labels on the same visual plane as the dots. From our point of view the strip labels deserve their fate for getting between our precious graphics panels.

In Figure 2 we do not show the factor levels, and that appear as dark shaded rectangles in Figure 1 (page 29). These shaded rectangles complicate the reading of text. Our visual system sees lines at the edges of regions with contrasting lightness (see Friedhoff and Benson 1991). When text overplots such edges it is like plotting a character on a line. Adjusting to a background of changing contrast when reading is also extra work. A possible solution, when there is space, is to plot little dots near the edge of the strip frame. For example plotting one dot could indicate level 1, two dots level two, and so on. Perceptual grouping in units of five would aid in fast perception for numbers over five. Printing the levels as text might be considered but may get confused with the labels. We do not consider the numeric level information important in Figure 1 so do not show it.

Cleveland (1993a) discusses on multivariate sorting by medians to establish the plotting order for barley varieties, sites and years. Sorting is important to simplify appearance and facilitate comparison. Figure 1 is sim-

ilar to a figure that appears in a copyrighted user manual (Mathsoft 1995). In the effort to avoid the appearance of producing an exact copy of such there are some changes. The obvious one is putting 1931 to the left of 1932. In this particular case we lean toward following the left to right reading convention rather to maintain full consistency in the sorting of factors.

6. Perceptual Grouping

The paper also presents a second facet of graph redesign, perceptual grouping. Having seen the merits of perceptual grouping illustrated by Kosslyn (1994), it is hard not to promote wider usage of various forms of perceptual grouping in statistical graphics. A few examples now appear in newsletter articles (Carr and Pierson 1996, Carr and Olsen 1996, and Carr et al. 1998).

One direct application of perceptual grouping is in table-lookup process that Cleveland (1993b) calls matching. Consider the visual task of matching labels and points. Figure 1 uses lines across the panels to connect labels and dots. This makes the plot look busy. Further horizontal lines inhibit vertical visual flow. The line linking works to the extent that one has the patience to track a line all the way across the page. Figure 2 groups the labels and dots in a 3-4-3 pattern. It is trivial to pick out the middle label in the top group of three and the middle dot in the top group of three on the right panel. We suggest that for remote points this label to point matching is easier than individual line tracking. It is also trivial to pick out the middle group or the bottom group on each panel and match the respective labels and points within each group.

There is a possible drawback to the perception grouping of barley varieties in Figure 2. Readers may look for some other meaning to the grouping than simple perceptual grouping. This is in part a matter of convention. If perceptual grouping were routine, the default interpretation for unlabeled groupings would suffice.

The design of Figure 2 strengthens the perceptual grouping by connecting points in a group with lines. This makes it relatively easy to compare the middle four points in the top panel with the middle four points in the fourth panel below. One can easily focus attention of the third points in these two sets of four points to make comparison. We think it is harder to find, remember and compare the two 6th points using Figure 1 with all symbols. The connecting lines help to guide the visual flow down the page. Each line leads the reader to the next point. Searching to find the next point is necessary only for disconnected points. Thus connecting points serves a worthy purpose.

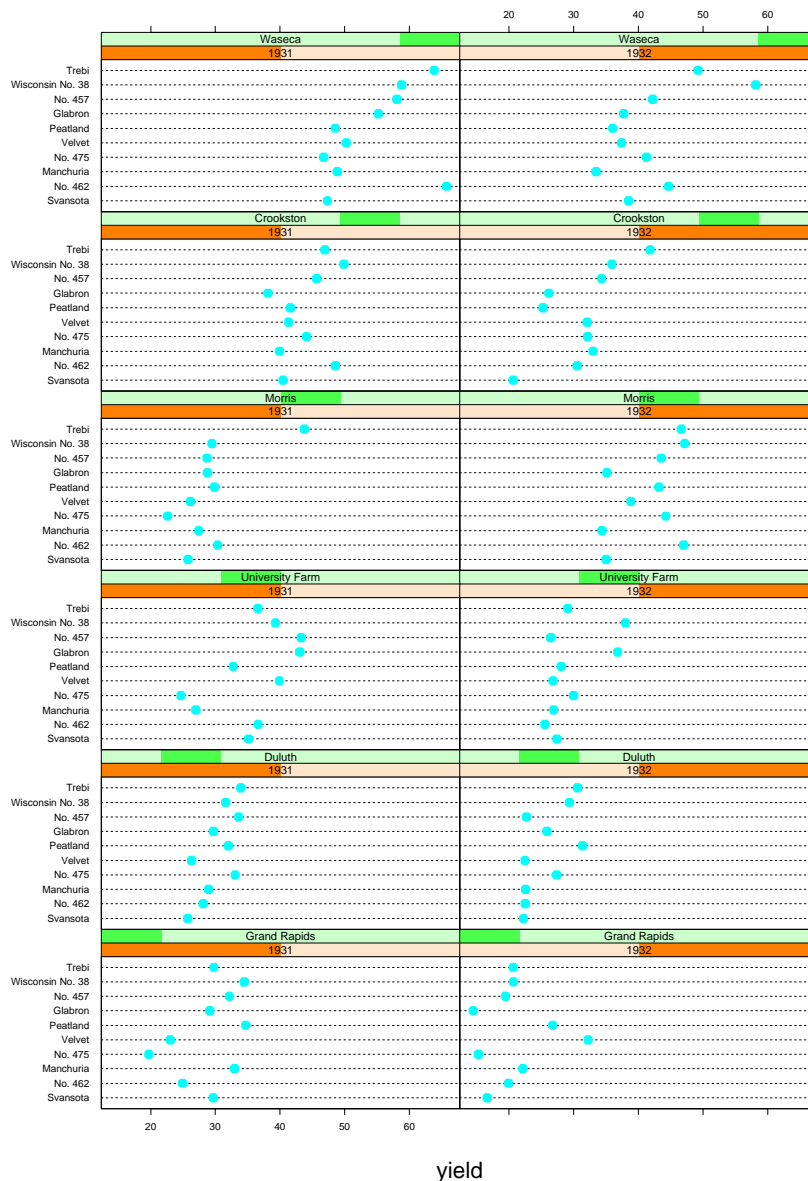


Figure 1: A Trellis Graphics dot plot with strip labels.

There are at least two drawbacks to the connecting lines. First, treatment of points is not uniform. Some consecutive points are connected and others are not. Connecting all the points in a panel solve this problem. However, experience with such examples (see Wilkinson 1993 and Carr 1994) suggests that ten items provide too complicated a shape to store in short term memory. It is not a small perceptual group. The unequal treatment becomes a price that we are willing to pay for the advantage of rapid comparison.

The second drawback is that there seems to be a graphic convention to avoid connecting points unless linear interpolation is implied. This convention is routinely vi-

olated in time series line plots. The designers of such time series graphic do not typically plan to defend the idea of precise linear interpolated between values even if such values are plausible. In Figure 2 we connect consecutive points based on the multivariate-sorting-established rank order. One could interpret the points along the connecting lines as being interpolated values for mixing barley varieties, but this is not what we intend. We want to guide the visual flow down the page. We are willing to pay the price of confusing those who might spatially interpolate population values between New York and Paris because the rank ordered population dots are connected on a graph.

Barley Yield




Figure 2: An enhanced dot plot, making use of lighting and perceptual grouping.

7. Access and Thoughts Spawned by Figure 2

We produced the graphics in this paper using Splus. Script files for the examples are available via anonymous ftp to galaxy.gmu.edu. Change directory to pub/dcarr/newsletter/lightlayer.

Lighting models, doubling buffering for animation, map layer composition and a host of other graphics tools have long been used in computer graphics or in GIS environments. The computer human interface community is experimenting with some very interesting interaction methods. This gives pause for thought. Yes, the business graphics in spreadsheet software are very old designs. However, the graphics that many of us statistical packages user cope with is also very old. One can go only so far by putting new wrappers on old graphics

engines. Our community could use more modern environments for combining statistical transformation with visual representations.

This paper skips over several variations and extensions of methods used in the design of Figure 2. Some very old density and texture applications that use 3-D lighting models have yet to see the light of publication. Looking forward, the power of using lines and “ribbons” and translucent surfaces in statistical graphics is far from been fully appreciated. Much work remains.

Some of what remains to be done may seem surprisingly simple. A perceptual question raised by Figure 2 is, “how close is the graphic.” Vocal presentations can be too soft, meet the listener half way, or be too loud. Graphics can go all the way from being too remote to being “in your face.” An example of the latter is a full

screen 3-D three-valued pie chart. Too loud or too close is often associated with lack of content and purposeful inhibition of thought. Now immersive graphics are an option. Where this fits in the spectrum remains to be seen.

The desire here has been to lift the information off the page but not very far. Our goal is to actively present the information (move it forward from flatland), yet present it gently. The visual flow and perceptual grouping correspond to efforts in clear articulation. That is, we don't want the reader to work harder than necessary to translate the message and want to provide the reader with the freedom to think about the message. There are simple facets of the graphics presentation process that seem to be little studied. For starters, how big should dots be for us to easily see the color and how small should they be for us to take the dot center for granted. More generally, what is the right balance in terms of perceptual proximity to serve the intended communication purpose?

To a designer's eyes graphics are rarely finished but rather opportunities for further experimentation. Perhaps readers will look at both Figures 1 and 2 and come up with something better.

Acknowledgements

S-Plus is a register trademark of Mathsoft, Inc. This work was supported by the EPA under cooperative agreement No. CR8280820-01-0. The article has not been subject to review by EPA, does not necessarily reflect the view of the agencies, and no official endorsement should be inferred.

References

Blinn, J. (1996) *Jim Blinn's Corner: A Trip Down the Graphics Pipeline*, Morgan Kaufmann, SF, CA.

Card, S. K., J.D. Mackinlay and B. Shneiderman, Editors (1999) *Reads in Information Visualization: Using Vision to Think*, Morgan Kaufmann, SF, CA.

Carr, D. B. (1993) "Production of Stereoscopic Displays for Data Analysis," *SCGN*, Vol.4 No. 1, pp. 2-7.

Carr, D. B. (1994) "Converting Tables to Plots," Tech Rpt No. 101, Center for Computational Statistics, GMU.

Carr, D. B. and A. R. Olsen (1996) "Simplifying Visual Appearance By Sorting: An Example Using 159 AVHRR Classes," *SCGN*, Vol. 7 No. 1 pp. 10-16.

Carr, D. B., A. R. Olsen, J. P. Courbois, S. M. Pierson, and D. A. Carr (1998) "Linked Micromap Plots: Named and Described," *SCGN*, Vol. 9 No. 1 pp. 24-32.

Carr, D. B. and S. Pierson (1996) "Emphasizing Statistical Summaries and Showing Spatial Context with Micromaps," *SCGN*, Vol. 7 No. 3 pp. 16- 23.

Cleveland, W. S. (1985) *The Elements of Graphing Data*, Hobart Press, NJ.

Cleveland, W. S. (1993a) *Visualizing Data*, Hobart Press, NJ.

Cleveland, W. S. (1993b) "A Model of Studying Display Methods of Statistical Graphics," *JGCS*, Vol. 2, No. 4.

Cleveland, W. S. and R. McGill (1984) "Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods," *JASA*, 79:531-554.

Ernst, B. (1976) *The Magic Mirror of M.C. Escher*, Ballantine Books, NY.

Foley, J. D., A. van Dam, S. K. Feiner, and J. F. Hughes (1990) *Computer Graphics: Principles and Practice, 2nd Edition*, Addison-Wesley, NY.

Friedhoff, R. M and W. Benson (1991) *The Second Computer Revolution - Visualization*, Freeman and Company, NY.

Kosslyn, S. M. (1994) *Elements of Graph Design*, Freeman and Company, NY.

MacEachren, A.M. (1995) *How Maps Work: Representation, Visualization, and Design*, The Guilford Press, NY.

Mathsoft (1995) *S-PLUS Trellis Graphics User's Manual, Version 3.3*, MathSoft, Inc. Seattle, WA.

Penrose, R. (1958) *British Journal of Psychology*, 49(1).

Pickle, L.W., Mingle, M., Jones, G. K. and White, A. A. (1997) *Atlas of United States Mortality*, Hyattsville, MD: National Center for Health Statistics.

Tufte, E. R. (1983) *The Visual Display of Quantitative Information*, Graphics Press, CT.

Tufte, E. R. (1990) *Envisioning Information*, Graphics Press.

Wegman, E. J., D. B. Carr and Qiang Luo (1993) "Visualizing Multivariate Data," *Multivariate Analysis: Future Directions*, ed. C.R. Rao, North Holland, New York. pp. 423-466.

Wilkinson, L. (1993) "Comment on A Model for Studying display Methods of Statistical Graphics," *JGCS*, 2(4).

Dan Carr
Center for Computational Statistics
George Mason University
dcarr@galaxy.gmu.edu

Ru Sun
Institute for Computational Sciences
and Informatics
George Mason University
rsun@osf1.gmu.edu

