

The Visual Design and Control of Trellis Display

Richard A. Becker
William S. Cleveland
Ming-Jen Shyu

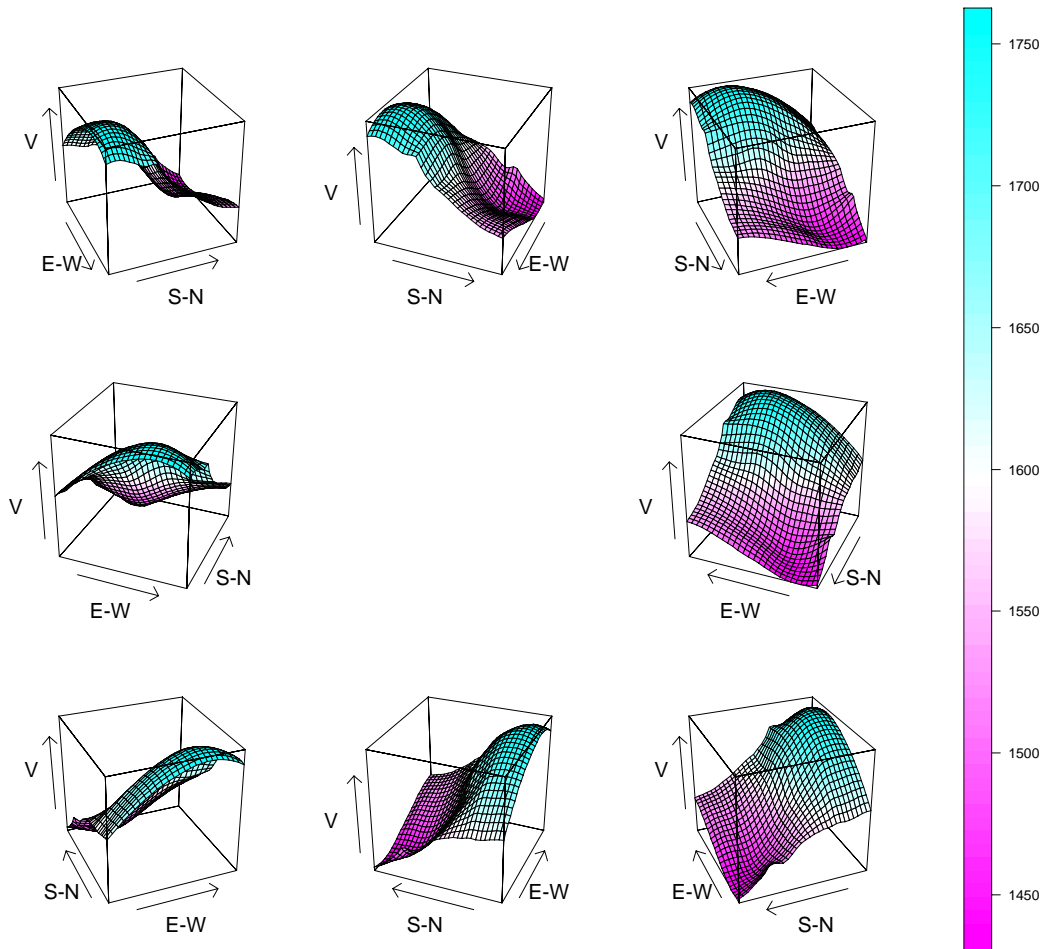


Figure 1: Trellis display is a framework for the visualization of data. One important application is uncovering the relationships of variables in multivariable data sets. In this figure, which will be described in detail later in the paper, a velocity surface is graphed as a function of longitude and latitude on eight panels, each for a different viewing direction.

Abstract

Trellis display is a framework for the visualization of data. Its most prominent aspect is an overall visual design, reminiscent of a garden trelliswork, in which panels are laid out into rows, columns, and pages. On each panel of the trellis, a subset of the data is graphed by a display method such as a scatterplot, curve plot, boxplot, 3-D wireframe, normal quantile plot, or dot plot. Each panel shows the relationship of certain variables conditional on the values of other variables.

A number of display methods employed in the visual design of Trellis display enable it to succeed in uncovering the structure of data even when the structure is quite complicated. For example, Trellis display provides a powerful mechanism for understanding interactions in studies of how a response depends on explanatory variables. Three examples demonstrate this; in each case, we make important discoveries not appreciated in the original analyses.

Several control methods are also essential to Trellis display. A control method is a technique for specifying information so that a display can be drawn. The control methods of Trellis display form a basic conceptual framework that can be used in designing software. We have demonstrated the viability of the control methods by implementing them in the S/S-PLUS system for graphics and data analysis, but they can be implemented in any software system with a basic capability for drawing graphs.

1 Introduction

1.1 Barley Data: The Detection of A Probable Error

In the 1930s an experiment was run in the state of Minnesota in the United States. At six sites, ten varieties of barley were grown in each of two years. The data collected for the experiment are the yields for all combinations of site, variety, and year, so there are $6 \times 10 \times 2 = 120$ observations. The experiment is of historical interest because it is one of the early field trials that incorporated R. A. Fisher's ideas on randomization and the analysis of variance. The agronomists published the data and an analysis of them in a 1934 paper [11]. Fisher published the data in his classic book, *The Design of Experiments* [10], but he did not present an analysis. Fisher's publication gave the data a large exposure, and many others tried their hands at analyzing them to illustrate new statistical methods [1, 2, 6]. We will do the same here, using the data to illustrate Trellis display. The visualization using Trellis reveals an important happening in the data — there appears to be a major error, one that survived undetected for six decades [4].

1.2 Trellis Display of the Barley Data

Figure 2 is a Trellis display of the barley data. Each panel displays the yields of the ten varieties for one year at one site.

Figure 2 uses an important display method: *main-effects ordering of category levels*. For these barley data, the explanatory variables are categorical. (Since there are only two years, the year variable is also treated as categorical.) The unique values of each categorical variable will be referred to as *levels*. For example, the levels for the site variable are Grand



Figure 2: A dotplot of the barley data showing yield against variety given year and site.

Rapids, Duluth, and so forth. The level medians are a measure of the main effects, and we have arranged that the levels for each variable are ordered based on level medians. On each panel the varieties are ordered from bottom to top by the variety medians; Svansota has the smallest median and Trebi has the largest. The site panels have been ordered from bottom to top by the site medians; Grand Rapids has the smallest median and Waseca has the largest. Finally, the year panels are ordered from left to right by the year medians; 1932 has the smaller median and 1931 has the larger. Later, we will discuss why main-effects ordering is important.

Visually scanning up each column of Figure 2, we can see an anomaly: for each year, the values for Morris appear out of place. Because of the main-effects ordering, the site medians increase from bottom to top. The ordering is preserved in each year separately except for Morris. But the visual impression is that if we were to interchange the years at Morris, the site would then fit into the patterns formed by the other sites.

This suggests another display. In Figure 3, the data are graphed again, but this time the 20 values for each site are graphed on a single panel with the year variable encoded by the plotting symbol. Now we can see clearly that at each site except Morris, the overall yield for 1931 is greater than 1932, but the reverse is the case for Morris. However, something else quite critical is also apparent. At Morris, the overall level of the absolute differences between the years has a value that is commensurate with the corresponding values at the other sites. (This is actually the same observation from Figure 2, that Morris would fit the pattern were we to interchange its years.) This suggests that there might be an error in the data at Morris, a reversal of the years. Either there is an error, or nature just happened to reverse effects at Morris in such a way that 1932 exceeds 1931 by an amount similar to the amounts that 1931 exceeds 1932 at the other sites. We will probe this issue later with other Trellis displays.

1.3 Trellis Basics

The salient visual aspect of Trellis display is a three-way rectangular array of panels with *columns*, *rows*, and *pages*. In Figure 2 there are 12 panels, 2 columns, 6 rows, and 1 page. In Figure 3 there are 6 panels, 1 column, 6 rows, and 1 page. Later, we will show a Trellis display with more than one page. We refer to the rectangular array as the *trellis* because it is reminiscent of a garden trelliswork.

Each panel of a trellis display shows a subset of the values of *panel variables*; these values, are formed by conditioning on the values of *conditioning variables*. In Figure 2 the panel variables are variety and yield, and the conditioning variables are site and year. On each panel, values of yield and variety are displayed for one combination of year and site. For example, the lower left panel displays values of variety and yield for Grand Rapids in 1932. In Figure 3 the panel variables are variety, year, and yield and there is one conditioning variable, site.

In Figure 2 the descriptions of the values of the year and site for a panel are given in *strip labels* at the top of the panel. The strip labels for each variable have a dark bar that indicates the value of the variable. This conveys in a graphical way how the values of the conditioning variables are changing over the trellis.




Figure 3: A dotplot of the barley data showing yield against site and year given variety.

The *panel display method* used in Figure 2 is a dot plot. The panel display method used in Figure 3 is a dot plot with the year encoded by plotting symbol. The information sent to each panel for rendering by the panel display method is a *packet*. The packet includes the values of the panel variables that are graphed on the panel, but only for those observations that should appear on the panel, as determined by the conditioning variables. In Figure 2 there are 12 packets and in Figure 3 there are 6 packets.

Trellis display can be used for small data sets — displaying the data with a simple single-panel display. It can be used for large data sets — displaying the data perhaps using hundreds of panels. But in the interest of journal space, we will need to restrict ourselves in this account to small data sets such as the barley data.

1.4 Display Methods and Control Methods

A variety of *display methods* employed in the visual design of Trellis display enable it to succeed in uncovering the structure of data. Several *control methods* are also essential. A control method is a technique for specifying information so that a Trellis display can be drawn. The control methods of Trellis display form a basic conceptual framework that can be used in designing software. We have demonstrated the viability of the control methods by implementing them in the S/S-PLUS system for graphics and data analysis [3], but they can be implemented in any software system with a basic capability for drawing graphs.

The display methods and the control methods have been developed in concert. Here is one example. Control methods specify the geometry of the trellis — that is, the layout into columns, rows, and pages — and then the assignment of the packets to the panels. These control methods are enabled by the use of the strip labels, a display method.

In the coming sections, we describe the display methods and the control methods of Trellis display. Because the precise boundary between display and control is sometimes fuzzy, we do not always delineate the two in the discussion.

2 Main-Effects Ordering

Main-effects ordering, described in Section 1, often allows us to perceive structure in the data that cannot be seen easily without it. Figure 2 provides one example. The ordering of the site levels from bottom to top by the site medians establishes a visual pattern from which Morris deviates, and this allows us to discover the anomalous behavior. By ordering the variety levels from bottom to top on each panel (and by maintaining the same horizontal axis for all panels) the 12 collections of plotting symbols on the panels form 12 gestalts that are alike and thus are more readily compared than if the levels were ordered in some other way. In Figure 4 the data are graphed again with the sites and varieties ordered alphabetically. We can no longer readily perceive the anomalous behavior at Morris, and the randomness of the patterns of plotting symbols within each of the 12 panels makes them harder to compare.

The main-effects ordering in Figure 3 lets us detect another property of the data, albeit

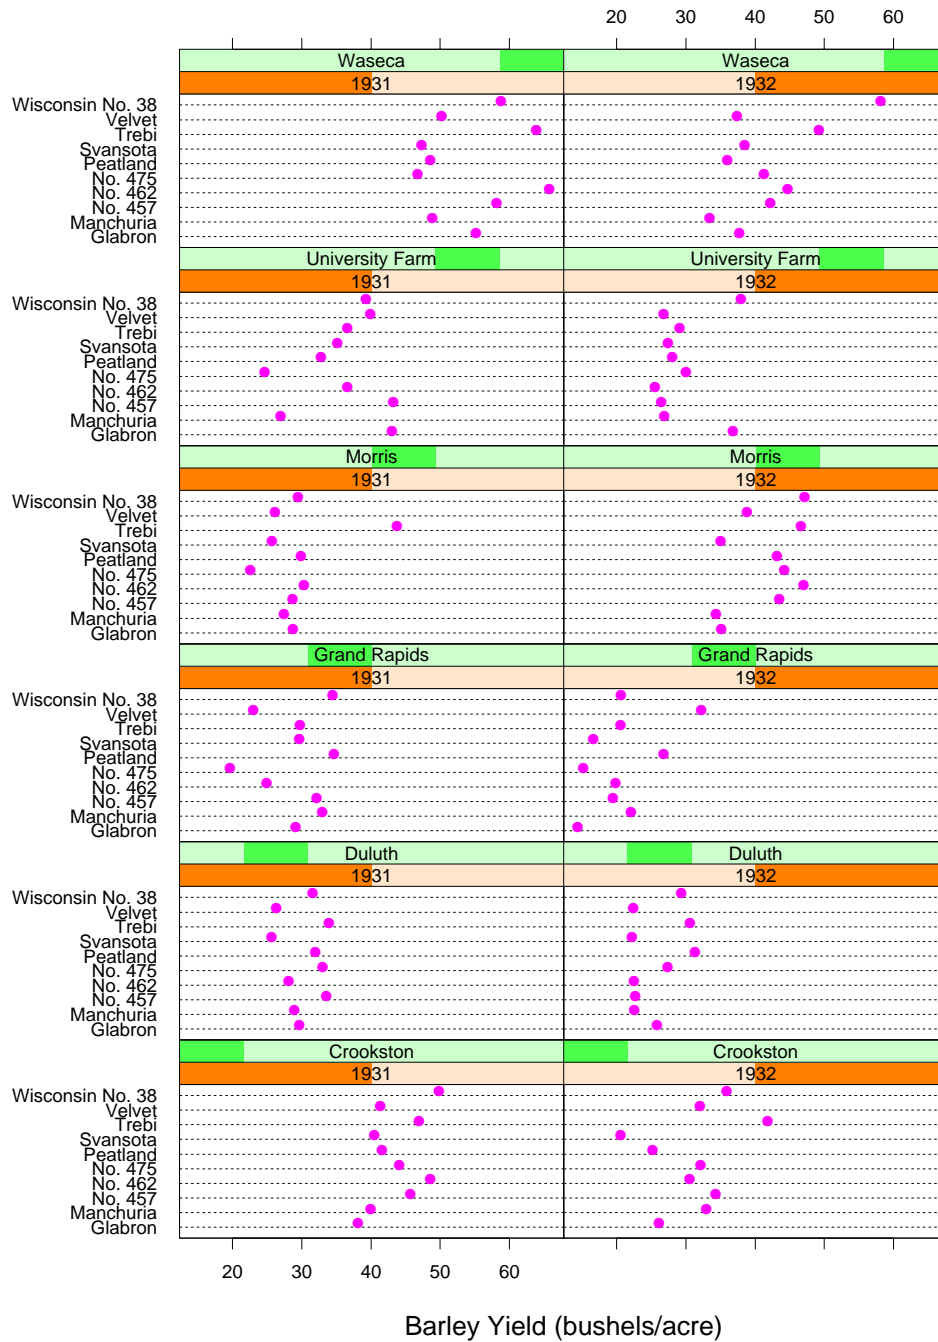


Figure 4: Yield against variety given year and site.

a somewhat subtle one. Consider the absolute value of the 1931 yield minus the 1932 yield for each site and variety. The average value of the 10 differences for each site appears to increase slightly from bottom to top. In other words, there appears to be a year by site interaction in which the yearly change is somewhat greater for sites with higher overall yields than for sites with lower yields. This is hardly a surprising interaction in data of this type. Furthermore, Morris appears to fit right into this interaction pattern. In fact, the site that deviates from the pattern is Grand Rapids, not Morris. This provides further evidence that the cause of the reversal at Morris is an error in the data.

Main-effects ordering is used extensively by Cleveland [4], who further demonstrates its usefulness. However, examples suggest that if a categorical variable is naturally ordered and there are more than two levels, then it is typically better to use the natural order of the variable instead of main-effects ordering.

3 Multiple Conditionings and Partial Residuals

Let us take a big step and suppose the barley data for Morris are incorrect, interchange the years, and continue the analysis. It makes sense to study Figures 2 and 3 again with the interchanged data, but we omit the new displays here in the interest of space, and continue with other displays of the interchanged data.

3.1 Multiple Conditionings

Figure 2 reveals not only the Morris anomaly, but also the relative performance of the different varieties. Trebi and Wisconsin No. 38, which have the largest variety medians, are typically among the highest yield varieties for each combination of site and year. An exception is Grand Rapids, which overall appears to behave quite differently from the other sites. At that site, Trebi is an average performer.

Figure 2 allows an incisive comparison of varieties for each combination of site and year because the 10 variety measurements are graphed together on the same panel. This graphing along a common scale allows highly accurate visual decoding of the relative values of the 10 yields. Similarly, in Figure 3 we were able to compare the two yearly values for each combination of site and variety because the two values are graphed along a common scale. But in neither of these displays can we compare the six values of yield for each combination of variety and year because the six values are on separate panels.

Thus we make another Trellis display in Figures 5 and 6, using the interchanged data. The panel variables are site and yield and the conditioning variables are variety and year. There are 20 panels and the trellis has two columns, five rows, and two pages. Now the six yields for each combination of year and variety are graphed along a common scale. Waseca, the site with the largest median, has close to the highest yield on all twenty panels, so it is consistently the leader. Grand Rapids has the lowest or nearly the lowest yield in all cases except Velvet in 1932 where it has an unusually high value compared to its performance elsewhere.




Figure 5: Yield against site given variety and year — page 1.




Figure 6: Yield against site given variety and year — page 2.

This change in what we can readily perceive in the Trellis displays for the barley data is a general phenomenon: when we seek to determine the dependence of a response on two or more explanatory variables, we typically see most effectively the dependence on the explanatory variables selected as panel variables, and we see how this dependence changes as the values of the conditioning variables change. Typically we do not perceive as effectively the dependence on the conditioning variables. Thus, in analyzing such multivariable data we want to make multiple Trellis displays so that each explanatory variable appears at least once as a panel variable.

3.2 Partial Residuals

Earlier we saw in Figure 3 a suggestion of an interaction between year and site — the higher the overall yield for a site, the greater the difference between 1932 and 1931. Let us focus the analysis on the year variable by adjusting for main effects of the other two explanatory variables. We could do this in the following way: (1) take means across the two years for each combination of site and variety; (2) subtract each mean from the two measurements; (3) graph the residuals by Trellis displays with the residuals taken as the response. Such Trellis displays are partial residual plots [7]. In this particular example, because there are just two levels of the year variable, each residual is one-half the difference of a 1931 value and 1932 value, and the 1931 residual for a particular site and variety is the negative of the 1932 residual for that site and variety. Thus, we can simplify the display by just graphing the 1931 values minus the 1932 values.

Figure 7 graphs the differences using the interchanged data. We have maintained the ordering of sites by overall site median. The vertical line on each panel is the 10% trimmed mean of the yields for the panel, that is, the mean with the smallest and largest differences removed. Now we see clearly the interaction between site and year that was hinted at in Figure 3. The overall level of the differences increases with the site median; Morris fits nicely into the pattern and Grand Rapids now emerges as the unusual site. This makes it seem even more likely that the original data are in error.

The effectiveness of displaying partial residuals is a common occurrence. By removing gross main effects, they can allow subtler effects to emerge. Partial residuals have a long history in statistics. For the most part, their usage, at least by the name of “partial residuals”, has been in studies where the explanatory variables are numerical [7]. But, in fact, partial residuals are a special case of a more general principle that emerges with greatest clarity in the exploratory methods of Tukey where effects are removed — fitted by models often known not to fully explain the data — to reveal other effects with more clarity [19, 20].

4 Trellising

The analysis of the barley data in the previous section illustrates a general phenomenon of the use of Trellis display; the structure of the data is revealed by the combined information of many Trellis displays, not just a few.

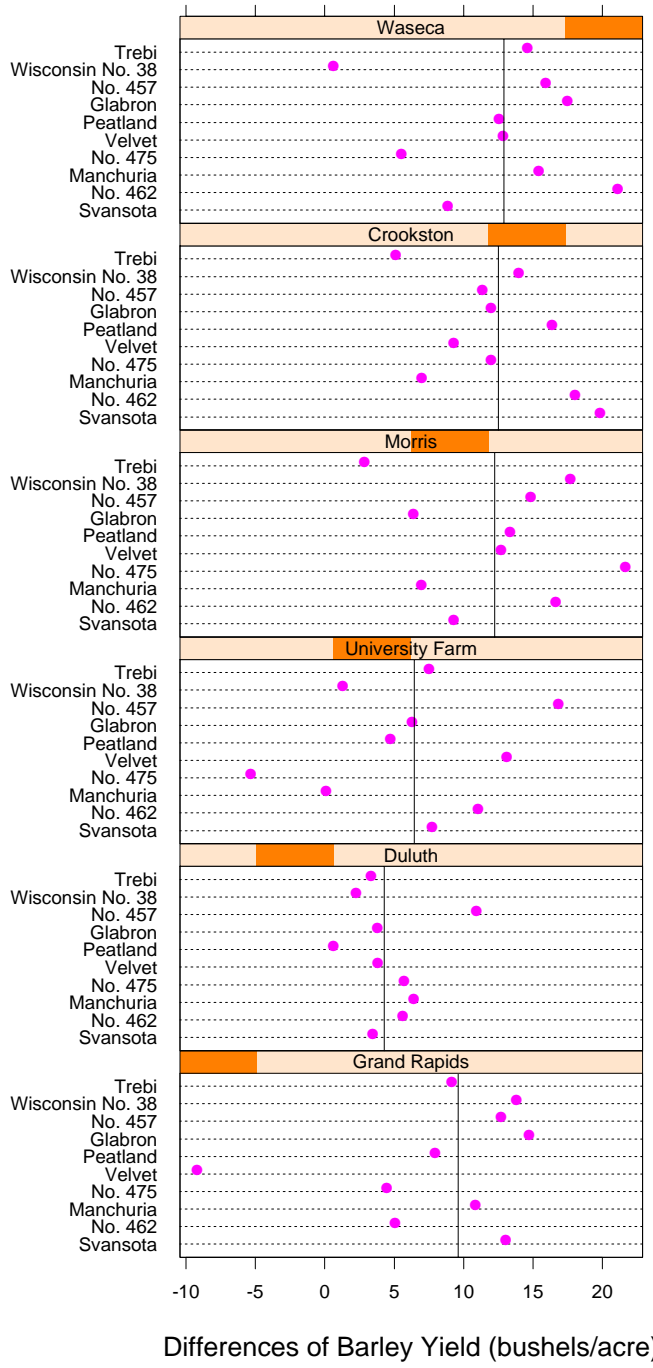


Figure 7: Differences of barley yield against variety given site.

To produce any Trellis display we must specify the trellis dimensions and the assignment of the packets to the panels of the trellis. We refer to this as *trellising*. For example, the trellising for Figure 2 is shown in Figure 8. For Trellis display to be useful in practice, we

Waseca	Waseca
1932	1931
Crookston	Crookston
1932	1931
Morris	Morris
1932	1931
University Farm	University Farm
1932	1931
Duluth	Duluth
1932	1931
Grand Rapids	Grand Rapids
1932	1931

Figure 8: Trellising.

need a conceptual framework for trellising that is both flexible and simple. No matter how powerful the visual transmission of information of Trellis display, it would not be of much use if trellising were complex and difficult to think through in each case. Fortunately, there is a flexible and simple control mechanism. The ingredients of this mechanism, which we explain in detail in this section are the following:

1. We specify the dimensions of the trellis — columns, rows, and pages.
2. We specify an order for the conditioning variables and an order for the levels of each conditioning variable.
3. Based on our specification of these orderings, the packets are ordered by a simple rule (packet order).
4. Another simple rule governs the order of the panels (panel order).
5. Packets are assigned to panels according to the packet order and the panel order.
6. While we typically make the selections of the dimensions and the orderings of the conditioning variables in a coordinated way, Trellis display has been designed not make a connection. This is a critical feature that makes trellising both simple and flexible.
7. A mechanism for specifying skipping allows us to assign packets with an irregular structure to the rectangular trellis in an understandable way.
8. A mechanism for adding space between any two adjacent columns or any two adjacent rows allows us to visually organize the panels more effectively in certain applications.

4.1 Dimensions, Packet Order, Panel Order, and Packet Assignments to Panels

In making a Trellis display we specify the trellis dimensions, an order for the conditioning variables, and an order for the levels of each conditioning variable. Consider Figure 2 whose

trellising is shown in Figure 8. The dimensions are (2, 6, 1) — 2 columns, 6 rows, and 1 page. There are two conditioning variables, year and site; the first with two levels and the second with six. The order of the conditioning variables is (1) year; and (2) site. The order of the year levels is (1) 1932; and (2) 1931. The order of the site levels is (1) Grand Rapids; (2) Duluth; (3) University Farm; (4) Morris; (5) Crookston; and (6) Waseca. Recall that the orders of the levels of site and year came from using main-effects ordering, but, of course, in any application we are free to order levels of the variables in any way that is sensible for the data at hand.

The following rule now orders the packets using the orderings of the conditioning variables and their levels: The levels of the first conditioning variable vary the fastest, the levels of the second conditioning variable vary the next fastest, and so forth. For Figure 2, the order of the packets is (1) 1932 Grand Rapids; (2) 1931 Grand Rapids; (3) 1932 Duluth; (4) 1931 Duluth; (5) 1932 University Farm; (6) 1931 University Farm; (7) 1932 Morris; (8) 1931 Morris; (9) 1932 Crookston; (10) 1931 Crookston; (11) 1932 Waseca; and (12) 1931 Waseca.

The panels of a Trellis display are ordered by the following rule: The bottom left panel of the first page is panel one; from there we move fastest through the columns, next fastest through the rows, and the slowest through pages. The following shows the panel order for Figure 2, which has two columns, six rows, and one page:

```

11 12
 9 10
 7  8
 5  6
 3  4
 1  2

```

The rule for assigning packets to panels is to match the packet order and the panel order. Packet one goes into panel one, packet two goes into panel two, and so forth.

The panel ordering rule is like a graph, not a table. The origin is at the lower left, and as we move from left to right or from bottom to top, the panel order increases. Holding to a graphical convention is important because we often have numerical values changing as either rows or columns change. For example, in Figure 2 the result of the main-effects ordering of the site levels and the panel ordering convention is an increase in the site medians in a graphical way — from the bottom row to the top row.

4.2 Different Trellisings

Let us consider different trellisings of the barley data to further illustrate how different dimensions and orderings alter the trellis.

Figure 2 was made tall and narrow to fit into a graphics region with a portrait shape. Just to review, here is the specification that produced its trellising: Dimensions = (2, 6, 1); Variable Order = year, site; Year Level Order = 1932, 1931; Site Level Order = Grand Rapids, Duluth, University Farm, Morris, Crookston, Waseca.

Suppose, however, that the graphics region has a landscape shape. Then to use space effectively we might want to interchange rows and columns, producing the trellising shown in Figure 9. We get this by the following specification: Dimensions = (6, 2, 1); Variable

1931	1931	1931	1931	1931	1931
Grand Rapids	Duluth	University Farm	Morris	Crookston	Waseca
1932	1932	1932	1932	1932	1932
Grand Rapids	Duluth	University Farm	Morris	Crookston	Waseca

Figure 9: Trellising.

Order = site, year; Site Level Order = same as previous trellising; Year Level Order = same as previous trellising. Thus we have altered the dimensions and reversed the order of the conditioning variables.

The Trellis display shown in Figures 5 and 6 has the two-page trellising shown in Figures 10 and 11. We get this trellis by the following specification: Dimensions = (2, 5, 2); Variable Order = variety, year; Variety Level Order = Svansota, No. 462, Manchuria, No. 475, Velvet, Peatland, Glabron, No. 457, Wisconsin No. 38, Trebi; Year Level Order = 1932, 1931.

1932	1932
Wisconsin No. 38	Trebi
1932	1932
Glabron	No. 457
1932	1932
Velvet	Peatland
1932	1932
Manchuria	No. 475
1932	1932
Svansota	No. 462

Figure 10: Trellising — page 1.

1931	1931
Wisconsin No. 38	Trebi
1931	1931
Glabron	No. 457
1931	1931
Velvet	Peatland
1931	1931
Manchuria	No. 475
1931	1931
Svansota	No. 462

Figure 11: Trellising — page 2.

Suppose, however, that the graphics region has a landscape shape. We might try to squeeze the panels on a single page by the trellising shown in Figure 12. We get this by the following specification: Dimensions = (5, 4, 1); Variable Order = same as previous trellising; Variety Level Order = same as previous trellising; Year Level Order = same as previous trellising. Thus we have changed only the dimensions.

1931	1931	1931	1931	1931
Peatland	Glabron	No. 457	Wisconsin No. 38	Trebi
1931	1931	1931	1931	1931
Svansota	No. 462	Manchuria	No. 475	Velvet
1932	1932	1932	1932	1932
Peatland	Glabron	No. 457	Wisconsin No. 38	Trebi
1932	1932	1932	1932	1932
Svansota	No. 462	Manchuria	No. 475	Velvet

Figure 12: Trellising.

4.3 Flexible Trellising

Earlier we emphasized the importance of the independence of the specification of packet order and the specification of the trellis dimensions. In particular, the numbers of levels of the conditioning variables and the trellis dimensions are independent. This allows flexible trellising. Of course, we often want to coordinate the two to enhance our ability to perceive the assignment of the levels of the conditioning variables to the trellis. When possible, we would like to have a column, row, or page dimension changing with the values of one conditioning variable and to have one variable per dimension. For example, in Figure 2, the trellising (shown by itself in Figure 8) was carried out so that the year levels change along the columns and the site levels change along the rows. But because the area of display devices is finite, and because we can find ourselves conditioning on more than three variables, we cannot always have such coordination. In Figures 5 and 6, whose trellising is shown by itself in Figures 10 and 11, the site variable changes both with the rows and with the columns. The independence of the packet order and the panel order is in part responsible for the flexibility that allows such a trellising.

The flexible trellising is also a result of a visual design feature of Trellis display — the strip labels. The labeling of the levels of conditioning variables that produce a panel is part of the panel. Thus we can arrange panels in any order and not worry about where labels are placed. The flexibility would not have been possible had we attempted labeling in the margins, the method typically used for multipanel displays.

4.4 Breaking

Two adjacent columns of a Trellis display meet at a *column join*, and two adjacent rows meet at a *row join*. Breaking the join — that is, adding white space between two columns or rows — can enhance our perception of the levels of conditioning variables in cases where more than one variable changes along the column dimension or along the row dimension. Consider the trellising in Figure 12 where 1932 packets are in the bottom two rows and 1931 packets are in the top two rows. Figure 13 breaks between rows 2 and 3 to enhance our perception of this division of panels by year.

1931	1931	1931	1931	1931
Peatland	Glabron	No. 457	Wisconsin No. 38	Trebi
1931	1931	1931	1931	1931
Svansota	No. 462	Manchuria	No. 475	Velvet
1932	1932	1932	1932	1932
Peatland	Glabron	No. 457	Wisconsin No. 38	Trebi
1932	1932	1932	1932	1932
Svansota	No. 462	Manchuria	No. 475	Velvet

Figure 13: Trellising with added white space.

4.5 Skipping

Suppose we have two conditioning variables, the first with 11 levels and the second with 19 levels. We might want a trellising with 19 pages, one for each level of the second variable. This means that we need to fill each page with 11 panels. Since 11 has no divisors except 1 and 11 our only choices for the trellis dimensions would be (1, 11, 19) or (11, 1, 19), unless we allow panel skipping. By skipping we can do the following: make the trellis dimension (3, 4, 19) and then on each page fill 11 panels with packets and then skip one. On each page the upper right panel would be empty.

The Trellis display skipping rule is the following: (1) a sequence whose values are “fill” or “skip” is specified; (2) if the sequence is smaller than the number of panels (the product of the dimensions) then the sequence is repeated. Thus the skipping sequence for the example in the previous paragraph is 11 values of “fill” followed by 1 “skip”.

5 Conditioning on A Numeric Variable with Repeat Measurements at Discrete Values

5.1 Fuse Data: A Missed Effect

Sheesley reports an experiment to study the properties of a thermal fuse in an electronic device [16]. The response is the operating temperature, F , of the fuse. There are three explanatory variables. The first is the ambient temperature, A , of the environment in which the fuse is located during a run. There are two levels for this variable, 75° and 110° . The temperature of the fuse is expected to increase as the temperature of the environment increases. The second is an ordered categorical variable, S , that describes the start condition of the fuse in a run. The value “cold” means that the fuse begins essentially at the temperature of the environment, and the value “hot” means that it was previously in operation at the start of the run. Starting “hot” is expected to produce a higher fuse temperature. The third variable is voltage, V . There are three levels — 110, 120, and 126 volts. The fuse temperature should increase with increasing voltage.

In the experiment, the $12 = 2 \times 2 \times 3$ combinations of the explanatory variables were tested with 10 replications (different fuses) for each combination. Due to cost considerations, once the experiment was set up to run a particular combination, all replications were carried out in parallel. While this saves money, it makes the experiment vulnerable to drift in other

variables that might affect F ; a special condition for one of these variables during the runs of one combination would be confounded with the effect of the combination. The use of Trellis display will show an effect missed in the reported analysis; either there is a genuine three-factor interaction or some special condition occurred during the testing of one or more combinations.

5.2 Categorical Variables and Numeric Variables Treated as Categorical

For the barley data, the variety and site variables are categorical. The year variable was treated as categorical, and not numerical or ordered categorical because there are only two values and because we cannot conclude that barley yield would likely increase or would likely decrease due to the mere march of time. Thus, for the barley data, all three explanatory variables are categorical and in making the Trellis displays, we conditioned on the levels of the variables, that is, the levels of the categories. For example, when site was a conditioning variable we made Trellis displays by conditioning on Grand Rapids, Duluth, etc.

For the fuse data, S is an ordered categorical variable with two levels, A is numerical with two levels, and V is numerical with three levels. We expect an increase in F with an increase in any one of these explanatory variables. For either A or V , because we have discrete values with repeat measurements at each value, we can condition on the distinct values in the same way that we condition on the levels of a categorical variable. Of course, as with an ordered categorical variable, we would almost always want to order the numerical levels by their values and not by main-effects ordering.

If the measurements of V , say, had not been under strict experimental control and varied over many values with few or no repeat values, we would need a special method for conditioning on V . Such a method will be discussed in Section 6.

5.3 Trellis Display of the Fuse Data

Figure 14 is a Trellis display of F against S given A and V . The levels of each of the explanatory variables are in their natural order, and as stated earlier, we expect F to increase as the level of any of the variables increases. The panel display method is a dot plot that uses open circles to graph the temperatures of the 10 fuses for each combination, and uses a “+” to graph the mean of the 10 values. The vertical line on each panel is at the mean of the 120 observations. The display shows the expected; $S = \text{“hot”}$ results in higher mean F than $S = \text{“cold”}$. But there is a strong interaction between S and A . For the lower level of A , the effect of S is substantially larger.

Figure 15 is a Trellis display of F against A given S and V . The same panel display method is used as in Figure 14. Again, we see the interaction between S and A but a different aspect is clearly portrayed. For $S = \text{“cold”}$, F increases as A increases, but for $S = \text{“hot”}$, A appears to have no effect.

Figure 16 is a Trellis display of F against V given S and A . The panel graphical method is a scatterplot of F against V with the means for each value of V connected by line segments. The shapes of the dependence on V in the top two panels are very nearly the same, both




Figure 14: Fuse temperature against start given ambient temperature and voltage. Means are shown by “+”




Figure 15: Fuse temperature against ambient temperature given start and voltage.

concave. The shape in the lower right panel is also concave, but with a smaller difference between the two slopes than in the top panels. The shape in the lower left panel is markedly convex. This is, of course, a three-factor interaction. We will use partial residuals to study it further. For each of the four combinations of S and A , there are 30 observations. In each case we take the mean of the 30 values and subtract it from the values. The new values are the partial residuals for a model with main effects and an interaction for S and A . Figure 17 is a Trellis display of the partial residuals against V given S and A using the same panel display method as in Figure 16. Careful study of this display together with our conclusions from the other displays reveals that something quite unusual is happening, probably for one or more of the runs with $S = \text{“cold”}$ and $A = 75^\circ \text{ F}$, the lower left panel in Figure 17.

Overall, the Trellis displays of the fuse data show a nonincreasing marginal effect: as the level of an explanatory variable increases, F increases, but the marginal effect of now increasing one of the other explanatory variables is the same or lessened. The unusual behavior in Figure 17 is a deviation from this pattern. Let us look at this in detail.

Figure 14 shows that for A at the greater temperature level there is a lessened effect of increasing S . Figure 15 shows that for S at “hot”, there is a lessened effect of A , in fact no effect. In Figure 17, consider the effect of going from $V = 120$ volts to $V = 126$ volts. The biggest change is for the lower levels of the other two variables, $S = \text{“cold”}$ and $A = 75^\circ$, the lower left panel of Figure 17. If we move to any other panel of the figure — which means that we increase the level of either S or A or both, the effect of V at the two selected levels decreases. If we move from the upper left panel to the upper right, the effect of V at the two levels stays the same. If we move from lower right panel to the upper right panel the effect of the two levels of V decreases. Now consider the effect of going from $V = 110$ volts to $V = 120$ volts. If we move from the lower right panel to the upper right panel, the effect decreases slightly. If we move from the upper left panel to the upper right, the effect is nearly the same. But now consider the effect in the lower left panel. It is the smallest of the four on the display. But given the overall law of nonincreasing marginal effect that obtains everywhere else, we would expect the effect in the lower left to be the greatest. It is possible that some very peculiar law of electricity and thermal conduction takes over for these two combinations. But it is more likely that some special condition in the operating environment occurred, confounding the condition with the other explanatory variables.

The original analysis of these data missed the unusual combination, which is a three-factor interaction. A plausible model for the data suggests unusual behavior is not random variation. The model is the following: main effects and an interaction for S and A ; a quadratic in V ; and a dummy variable that is 1 for the 10 observations of the special condition ($V = 120$, $A = 75^\circ$, $S = \text{“cold”}$) and -1 elsewhere. The dummy variable is significant at the 1.2% level.

6 Shingles

6.1 Haystack Data: A Solution to a Problem of Ezekiel and Fox

Before the introduction of hay-baling machines in the 1940s, hay was often sold loose in round stacks. Two aspects of a round stack that a farmer could easily measure with a rope

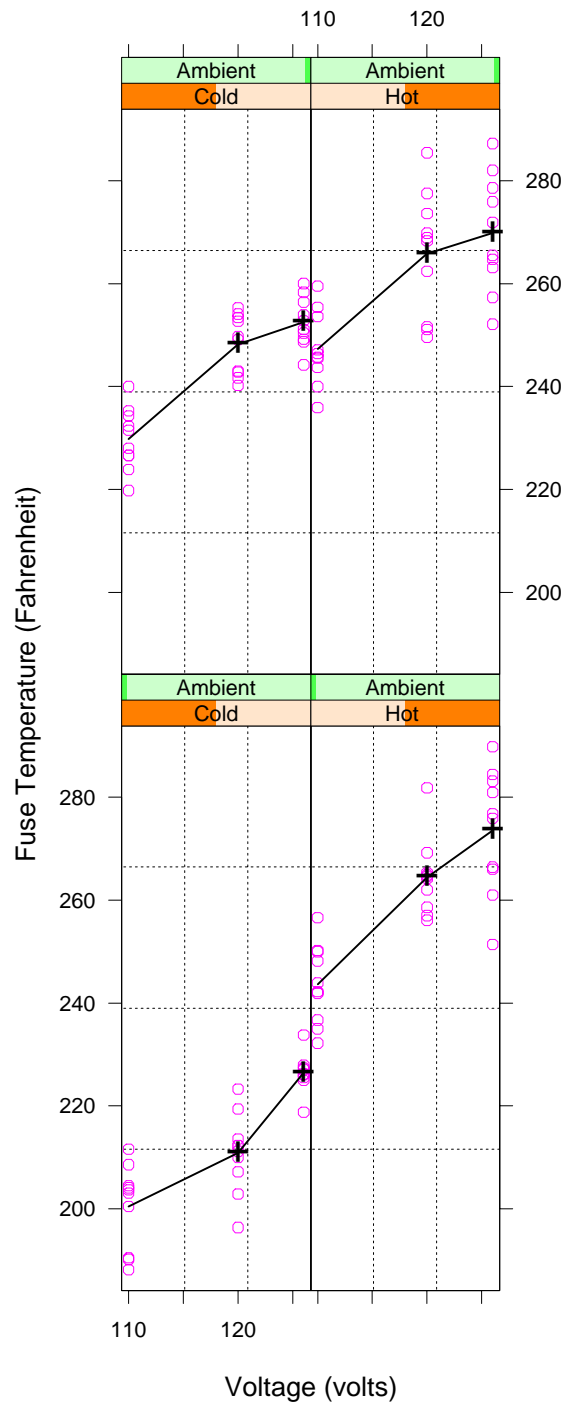


Figure 16: Fuse temperature against voltage given start and ambient temperature.

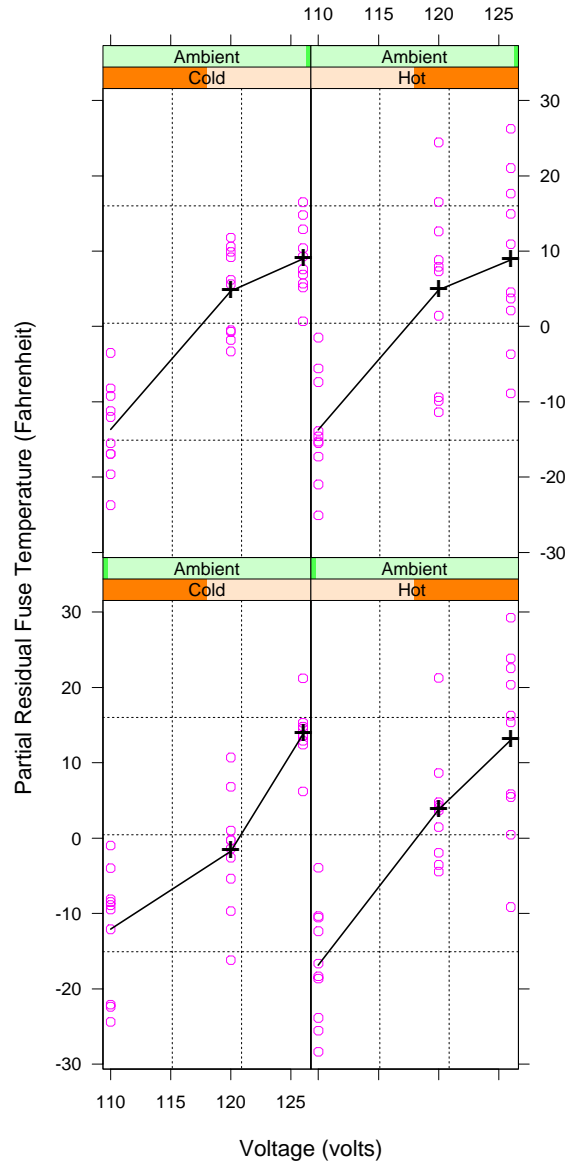


Figure 17: Partial residual fuse temperature against voltage given start and ambient temperature.

were the *circumference* of the base and the *over*, the distance from the base on one side, over the top at the center, to the point on the base on the opposite side. The circumference and the over were then used to estimate the volume.

Ezekiel and Fox report data from an experiment run in the 1920s to determine a method for estimating the volume of a round stack from its circumference and over measurements [8]. The volumes of 120 stacks were precisely determined, and measurements were made of the circumferences and the overs. Ezekiel and Fox analyzed the data by fitting log volume, V , to log circumference, C , and log over, O .

Figure 18 is a *scatterplot matrix*: scatterplots of all combinations of the three \log_2 haystack variables. Notice that O by itself is a very good predictor of V . Also, the measured values of O and C are correlated; their squared correlation coefficient is 0.40. Consider the basal diameter and the height of a haystack. C , of course, is a simple function of the diameter. But O also depends on the diameter, hence the correlation. O also depends on the height.

As we will see shortly, there is an interaction between C and O in explaining V . The interaction is discussed by Ezekiel and Fox:

The basal area is a function of the square of the basal circumference; the “over” is a function of both the basal diameter and the height — but attempts to separate the two have been unsuccessful. . . . The problem is evidently one where the relation may best be expressed by a joint function such as

$$\text{volume} = f(\text{circumference, over})$$

In other words, Ezekiel and Fox saw no way to eliminate the interaction. One might expect that if there were a way to define the height of a stack from the over and circumference measurements, then there might be no interaction in a model using the diameter and the height. But as they state, Ezekiel and Fox saw no way to do this. They went on to describe the function using a nonparametric, or smoothing, approach to fit V as a function of C and O . We will use Trellis display to reveal the structure of the interaction; this leads to method for determining a definition of haystack height which, in turn, leads to a model for the data with no interaction.

6.2 Conditioning on Intervals and the Equal Count Algorithm

For the haystack data, we would like to study conditional dependence — the dependence of V on C given O , and the dependence of V on O given C . But Figure 18 shows that C and O take on many different values, with insufficient repetition at each value to allow us to condition on single numerical values as we can for a discrete numerical variable with many observations at each distinct value. For variables such as C and O we must condition on each of a collection of intervals of values. For example, we must study the dependence of V on O given C in an interval, and vice versa.

For a categorical variable or for a discrete numerical variable we have referred to the values on which we condition — that is, the distinct values — as the *levels*. For a numerical

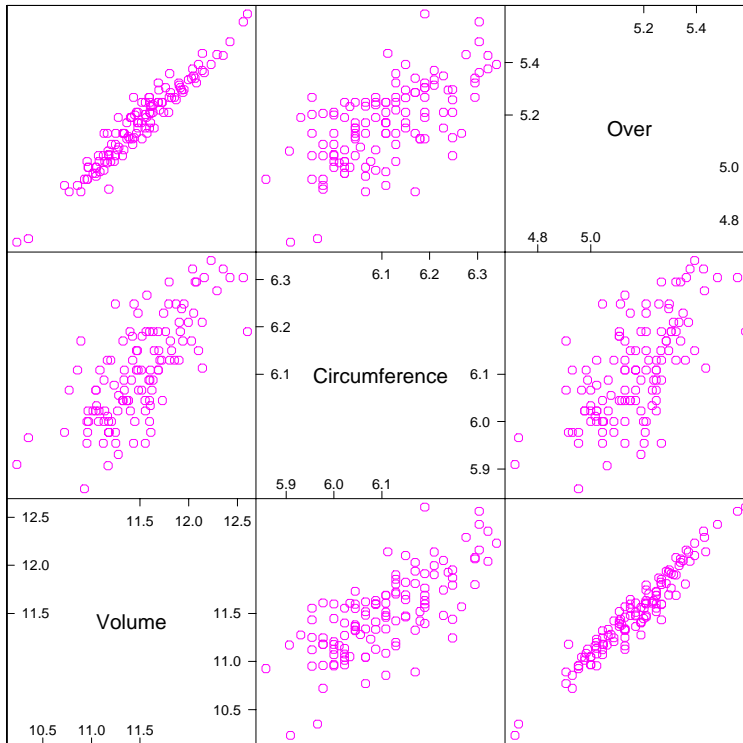


Figure 18: Scatterplot matrix of log base 2 of the haystack data.

variable in which we condition on intervals, the intervals become the levels, which, as in the other cases, are ordered.

We could form intervals by simple binning: breaking up the range of a conditioning variable into disjoint intervals. Such binning of numerical data to study conditional dependence reaches deeply into the history of statistics. In fact, Ezekiel, in earlier editions of the Ezekiel and Fox book cited here [8], used such simple binning in looking at the haystack data.

But we will take a different approach, binning in overlapping intervals. Figure 19 is a *given plot*; it shows intervals of O that we will use for conditioning. The intervals were formed from the *equal count algorithm* which we will describe shortly [4]. The overlap of the intervals increases the resolution with which we can study conditional dependence. Suppose something special happens in the conditional dependence of V on C when O is equal to $5 \log_2 \text{ft}$ — for example, the maximum curvature occurs. We are likely to best perceive this dependence if the collection of intervals contains at least one with $5 \log_2 \text{ft}$ at or near the center. We want successive intervals to slide along rather than jumping by large amounts. A good analogy is moving averages for a time series; there, to increase resolution, the averaging intervals overlap, sliding along to provide the best picture of the local behavior at each time point.

The intervals for a numerical variable together with the measured values of the variable form a data structure that we call a *shingle* because the intervals often overlap. In an implementation of Trellis display, it is convenient to store intervals and values together in a shingle data structure. We will discuss that point later in Section 7.

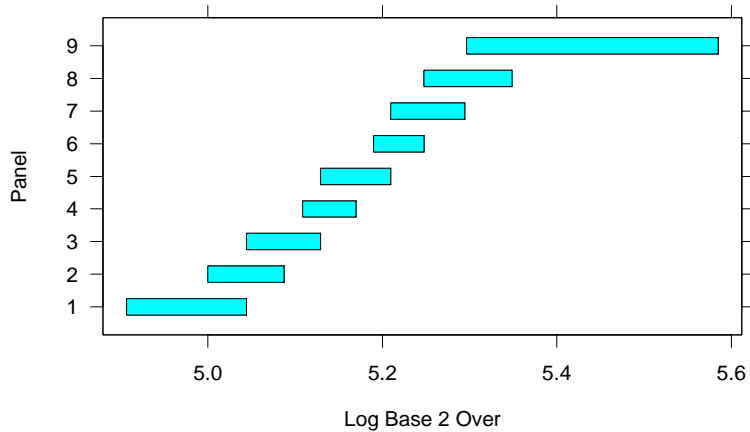


Figure 19: Conditioning intervals.

Two parameters are given to the equal count algorithm — the number of intervals and the approximate fraction of points shared by two successive intervals. In Figure 19, the parameters are 9 and 0.5. The algorithm then picks interval endpoints that are values of the data; the left endpoint of the lowest interval is the minimum of the data, and the right endpoint of the highest interval is the maximum of the data. The endpoints are chosen to make the counts of points in the intervals as nearly equal as possible, and the fractions of points shared by successive intervals as close to the selected fraction as possible. For Figure 19, the interval counts range from 11 to 18, and the numbers of points shared by successive intervals range from 5 to 13. The large variation in each of the two sets of numbers in this example arises from the large number of ties in the values of O .

6.3 Trellis Display of the Haystack Data

Figure 20 is a Trellis display of V against C given O using the equal-count intervals of Figure 19. The panel variables are V and C and the conditioning variable is O . The levels of O are ordered in the natural way, from the lowest interval to the highest. The dimensions of Figure 20 are (3, 3, 1). Unlike our previous examples, the packets are not disjoint because the intervals are not; the same observation may be in two adjacent packets. Because of the conventions for assigning packets to panels, the intervals go from lowest to highest as we move from left to right through the bottom row, then left to right through the middle row, and so forth.

The strip labels in Figure 20 contain a graphical display of the intervals with a scale going from the left edge of the strip to the right, but with no tick marks. The scale value at the left edge is equal to the value of the left endpoint of the lowest interval, and the scale value at the right edge is equal to the value of the right endpoint of the highest interval. The dark regions inside the strips portray the intervals. Compare their portrayal in Figure 19 with their portrayal on the strips.

Each panel of Figure 20 has a scatterplot of V against C for values of O in one of the intervals. A smooth curve is fitted to the data of each panel using loess with local linear




Figure 20: Log volume against log circumference given log over.

fitting, nearest-neighbor bandwidth selection, and a smoothing parameter of 1 [5]. The curves help us perceive the underlying patterns.

Figure 21 is a Trellis display of V against O given C . The panel variables are V and O and the conditioning variable is C . As with Figure 20, 9 equal-count intervals are used with a target fraction of overlap of 0.5. The loess curves use local quadratic fitting with a smoothing parameter of 1.

The underlying patterns in Figure 20 appear to be nearly linear but with slopes that tend to increase as the conditioning level of O increases, an interaction between O and C . This is the interaction referred to by Ezekiel and Fox that appeared to them to be unremovable. Figure 21 also shows the interaction. For low values of C there is a hint of concavity in the dependence of V on O for the lower values of C , but as C increases the dependence becomes nearly linear.




Figure 21: Log volume against log over given log circumference.

6.4 Modeling the Haystack Data

The problem with the haystack data is the interdependence of C and O . The over measurement depends on the haystack height and on the diameter, and the circumference measurement is equivalent to the diameter. This results in a correlation between O and C and an interaction in the dependence of V on O and C . Next, attacking the problem posed by Ezekiel and Fox, we will present a method for deriving a height measurement; then we will model volume by height and basal radius, seeking to eliminate both the correlation in the explanatory variables and the interaction.

Haystacks are mounds. Suppose the shape of each haystack is a *frustum* of a right circular cone, that is, the region of the cone between the base and a plane parallel to the base and intersecting the cone. Let b_i be the radius of the circular base of the i th haystack frustum, and let t_i be the radius of the circular top, so $b_i > t_i$. Let h_i be the height. Suppose further that the ratio of the top radius to the bottom is a constant, $t_i/b_i = \alpha$. Let o_i be the over measurement and let c_i be the circumference measurement of the base. Then

$$\begin{aligned} b_i &= c_i/2\pi \\ t_i(\alpha) &= \alpha b_i \\ h_i(\alpha) &= \sqrt{(o_i/2 - t_i(\alpha))^2 - (b_i - t_i(\alpha))^2}, \end{aligned}$$

and the volume is

$$f_i(\alpha) = \frac{1}{3}\pi h_i(\alpha)(b_i^2 + b_i t_i(\alpha) + t_i^2(\alpha)).$$

We have indicated the dependence of t_i , h_i , and f_i on α which in the coming model for the data will be an unknown parameter.

We will model the actual \log_2 volumes, V_i , of the haystacks by the logs of the frustum volumes,

$$F_i(\alpha) = \log f_i(\alpha).$$

The model is

$$V_i = \mu + F_i(\alpha) + \epsilon_i$$

where the ϵ_i are independent with mean zero and variance σ^2 . The unknown parameters are α , μ , and σ . If the haystack shapes are well approximated by frustums with a constant ratio of top radius to base radius, then the estimated value of μ should be close to 0.

The model was fitted by nonlinear least squares. The estimates of the parameters and their standard errors are $\hat{\alpha} = 0.90 \pm 0.030$, $\hat{\mu} = 0.092 \pm 0.011 \log_2 \text{ft}^3$, $\hat{\sigma} = 0.091 \pm 0.0058 \log_2 \text{ft}^3$. The sample distribution of the residuals is well approximated by a normal distribution. The standard errors of the estimates were estimated by a simulation with the ϵ_i generated from a normal distribution with standard deviation 0.091.

The squared correlation coefficient between $\log_2 b_i$ and $\log_2 h_i(\hat{\alpha})$ is 0.065. Furthermore, the fitted function can be written as two additive functions of $\log_2 b_i$ and $\log_2 h_i(\hat{\alpha})$. Thus we have created two explanatory variables that are nearly uncorrelated and that act with no interaction in explaining the response.

The approximation of haystack shape by a frustum is quite reasonable. The estimate of μ is significantly different from zero but its magnitude is small. The residual plots suggest

that the fitted model slightly underpredicts the volumes of the largest haystacks, but the distortion is small.

The fitted log volume model can be written as a function of O and C . We will graph this function to compare its properties with the behavior in the Trellis displays of the data in Figures 20 and 21. And we will use Trellis display to show the function as well.

The Trellis display in Figure 22 graphs two fitted functions. One function is the fit just discussed, which is based on the fustrum model; this fit is portrayed by the heavy solid curves on the panels. The second function is a fit that is linear in C and O ; this fit is portrayed by the lighter lines on the panels. The purpose of the linear fit is simply to provide a visual reference so we can more readily judge the curvature and interaction of the fustrum fit.

The fustrum fit is graphed over a *cropped* rectangular grid. Figure 18 shows that C and O , which are correlated, vary over a roughly elliptical region. We begin with the full rectangular grid: 50 equally-spaced values of C from the minimum measured value of C to the maximum, and 12 equally-spaced values of O from the minimum to the maximum. This full grid contains the elliptical region of measured values of C and O . We crop the grid back to those points that lie inside the elliptical region. The result then is values of V for the cropped grid values of C and O . Figure 22 graphs V against C given O . Since C is discrete in this case, we condition on the distinct values of O . Figure 23 uses a similar method to graph V against O given C .

Figures 22 and 23 show that the qualitative properties of the fitted function are quite similar to the patterns observed in Figures 20 and 21. The interaction between O and C is shown clearly in Figure 22. As the conditioning values of O increase, the overall slopes of the curves increase. However one property that is not seen in Figure 20 is the concavity of the conditional dependence; in Figure 20 it appears linear. But the amount of noise about the underlying patterns in Figure 20 is large enough that it could easily mask the mild curvature shown in Figure 22. Figure 23 shows the same mild concavity shown in Figure 21.

7 Discussion

7.1 The Effectiveness of Trellis Display

This account of Trellis display has been organized around applications. We studied three data sets and in each case made important discoveries not appreciated in the original analyses. There are two reasons for this organization. First, it is far easier to describe the details of a visualization method through sets of data. Second, while we can appeal to the theory and experiment of visual perception to attack very specific issues of visualization, the only way to establish the validity of a general approach to visualization that contains a whole collection of graphical methods is to provide examples where the method allows us to detect important properties of data sets that would be difficult or impossible to see with other statistical methods.




Figure 22: Log volume against log circumference given log over.




Figure 23: Log volume against log over given log circumference.

7.2 Standard and Nonstandard Trellis Displays

In practice, most Trellis displays can be produced by straightforward specifications of dimension, variable order, and standard panel display methods such as scatterplots, dot plots, histograms, and so forth. But when needed, special displays can be produced by creative use of dimension, variable order, panel display methods, panel skipping and by adding extra break space between panels. Figure 1 at the beginning of the paper provides an example that uses some of these techniques.

The raw data on which Figure 1 is based are measurements of the velocity of a galaxy in the Southern Hemisphere [5]. When looked at from the earth, the galaxy fills a very small area on the celestial sphere. Measurements of the radial velocity (in a direction toward or away from Earth) were made at 323 locations over this area. The velocities are not constant because the galaxy appears to be spinning and has other complicated motions. The data are noisy, so a loess surface was fitted to the velocities as a function of longitude and latitude on the celestial sphere [4]. The fit was then evaluated over a 26×31 grid of longitudes and latitudes.

On each panel of Figure 1, the gridded loess fit is displayed by a 3-D wireframe plot from some viewing direction. The base plane of the box is longitude and latitude, and the axis perpendicular to the base plane is velocity. There are 8 viewing directions, all at 45° to the base plane. The directions make a complete circle around the surface in equal angular steps. As we move from the lower left panel counterclockwise, the equally spaced viewing directions go from 22.5° to 337.5° in steps of 45° .

The packet for each panel is the identical set of (x, y, z) coordinates. There are eight such packets of triples, one for each viewing direction. The trellising is achieved in the following manner. A new variable, viewing direction, is defined. For each triple (x, y, z) , the value of the new variable is the circular angle of the viewing direction described above, so it takes one of the eight values from 22.5° to 337.5° . This new variable becomes the conditioning variable; it is numeric and has 8 unique values for conditioning. But the eight levels are not ordered by their numerical values. Instead they are ordered $22.5^\circ, 67.5^\circ, 112.5^\circ, 337.5^\circ, 157.5^\circ, 292.5^\circ, 247.5^\circ, 202.5^\circ$. Furthermore, there is a skipping vector defined by (4 fills, 1 skip, 4 fills). The result of the ordering and the skipping is wireframe displays whose viewing directions progress around a circle as we proceed circularly around the panels of the Trellis display. The panel display method performs the perspective transformation based on the data and the viewing direction.

7.3 Graphical Perception: Orientation Estimation and Visual Assembly

The theory and experiment of graphical perception have played an important role in the visual design and the control design of Trellis display. Here we discuss two issues, aspect ratio, and color.

In many graphical displays we judge the orientations of line segments to decode information about the rate of change of one variable with respect to another. For example, in Figures 22 and 23 we judge the orientations of the short line segments that make up the curves to decode information about the relative steepness of the curves and the amount of

curvature.

This decoding is greatly affected by the aspect ratio of the graph. The *data rectangle* of a graph is a rectangle that just encloses all of the data. The aspect ratio is the physical height of the data rectangle (measured in cm, for example) divided by the width. The aspect ratio is a vital parameter for judging rate of change from orientations because the orientations change as the aspect ratio changes. The aspect ratio in Figure 24 is much smaller than that of Figure 22; although both graph the same data, we cannot as accurately decode information about the relative overall slopes and the curvature from Figure 24. Both the theory and experiment of graphical perception have led to the 45° principle: orientations of line segments are most accurately judged when the absolute slopes are centered on 45° [4]. This *banking to 45°* was used in Figures 22 and 23. Because of its importance, control of aspect ratio has been built into Trellis display at a fundamental level.




Figure 24: Log volume against log circumference given log over.

Visual assembly, the simultaneous perceiving of a collection of graphical objects as a gestalt or visual whole, is another important visual operation of graphical perception. For example, on Trellis displays we want to perceive the strip labels for a particular conditioning

variable as a visual whole, filtering out other aspects of the display. Color encoding is a powerful method for producing effective assembly. In Trellis display, default color encodings have been developed that are reasonably robust to the output device being used. The guiding principle is to choose colors that are simple combinations of cyan, magenta, and yellow, three of the four standard printing colors. The reason is that these simple combinations work well on color printers but also work well on screens; encodings that work well on screens do not necessarily move well to printers. Simple combination is defined as follows: (1) We can add white to any of these three primaries to alter their lightness or saturation; for example, 80% cyan is 80% cyan and 20% white. (2) We mix no more than two of these altered primaries. For example, in Trellis display the first four colors for color encoding plotting symbols or curves are the following: (1) 100% cyan; (2) 100% magenta; (3) 50% yellow + 50% cyan (green); (4) 75% magenta + 25% yellow (orange).

7.4 A High-Level Design for Software for Trellis Display

By “high-level” design for software we mean the design of the mechanism that a user employs to specify what is produced. We have developed a high-level design for Trellis display and have implemented the design in the S/S-PLUS system for graphics and data analysis [3].

We have already provided information that has been used in the high-level design:

1. The trellising mechanism of Section 4 serves not only as the conceptual framework for trellising; it is also the control mechanism for users.
2. Conditioning variables in Trellis display use a data structure that consist of the values of the variable and the ordered levels to be used for conditioning. For categorical variables, the data structure is a *category* and the levels are a set of unique values. For numerical variables the data structure is a *shingle* and the levels are intervals. In both cases the data structure simply carries along the levels as an attached attribute; routines that want the values but that do not know about the levels can still access the data structure as if the levels were not there.

There are many other aspects of the high-level design, too numerous to discuss here, but there is one more that warrants discussion because it plays a central role. The highest-level routines for Trellis display implement specific graphical methods such as dot plots, quantile plots, boxplots, and wireframes, together with a few rather general routines that, say, graph y against x . The job of each of these routines is to control the trellising, scaling, and aspect ratio; to create the packets; and to draw the trellis framework which is everything on the display except the data regions of the panels. The drawing in the data regions is carried out by a *panel function* that is called by the high-level routine. The panel function is a single prescription that takes the packet for a particular panel as input and then draws something in that panel. There is just one prescription that applies to all panels. Each high-level routine comes with a default panel function. Users can program a panel function, which is relatively simple in most cases because it is a prescription for a single-panel display. In contrast, users are not expected to program the routines which deal with the very complex task of establishing the trellis framework. This two-part division works quite well, providing flexibility without undue complexity.

7.5 Trellis Display and Other Work in Multipanel Data Display

Arranging panels in a nested way to see the relationships of certain variables conditional on others has arisen in many forms in many different scientific disciplines. Among the displays that Tufte refers to as *small multiples* are examples of such conditioning displays [17, 18]. The casement window displays of Tukey and Tukey [21] show four numerical variables using a rectangular array of scatterplots of two of the variables, conditional on disjoint intervals of the other two variables, and with the scales of two conditioning variables in the margins to show the conditioning intervals. Mihalisin, Timlin, and Schwegler present a systematic approach to both the design and control of nested displays with certain types of display methods on the individual panels [13, 14]. The system, which has evolved into TempleMVV, has also been designed so that it is able to display very large databases [15]. In work based on these ideas, LeBlanc, Ward, and Wittels use nested multipanel displays in their *dimensional stacking* for displaying functions of several variables [12]. In the “worlds within worlds” of Feiner and Beshers [9], nesting is generalized to three dimensions; 3-D displays become the panels and they are arranged in a 3-D array with perspective. Finally, in their recent paper about TempleMVV, Mihalisin *et al.* give an excellent discussion of general criteria and goals for multipanel display [15].

Trellis display brings substantial generality to multipanel display. Indeed, it serves as an overall framework within which all data display can be carried out. Any graphical method — whether it is a univariate, bivariate, or higher dimensional display technique — can be displayed in a trellis, and there can be many panels or just a single panel. For example, each panel of a Trellis display can be a scatterplot matrix. In this account we used dot plots, scatterplots, curve plots, and wireframes. The major features that allow the generality are the following:

1. The independence of the specification of packets and the specification of the trellis geometry.
2. The simple rules for forming the packet order, for the panel order, and for matching the two orders.
3. The use of strip labels to make panels self-contained.
4. A division of responsibility into (1) high-level routines that implement broad graphical methods, controlling features specific to the method, scaling, aspect ratio, labeling, and the trellising; and (2) programmable panel functions that carry out the rendering in the data regions of the panels.

Acknowledgements and More Information

Montse Fuentes played a vital role in applying Trellis display to the data sets discussed here. We are grateful for her incisive insights and suggestions for improving Trellis Display.

More information about trellis display is available at

<http://netlib.att.com/netlib/att/stat/info/trellis.html>

References

- [1] F. J. Anscombe. *Computing in Statistical Science through APL*. Springer, New York, 1981.
- [2] F. J. Anscombe. Looking at Two-Way Tables. Technical report, Department of Statistics, Yale University, New Haven, Connecticut, U.S.A., 1983.
- [3] R. A. Becker, J. M. Chambers, and A. R. Wilks. *The New S Language: A Programming Environment for Data Analysis and Graphics*. Chapman and Hall, New York, 1988.
- [4] W. S. Cleveland. *Visualizing Data*. Hobart Press, Summit, New Jersey, U.S.A., 1993.
- [5] W. S. Cleveland and S. J. Devlin. Locally-Weighted Regression: An Approach to Regression Analysis by Local Fitting. *Journal of the American Statistical Association*, 83: 596–610, 1988.
- [6] C. Daniel. *Applications of Statistics to Industrial Experimentation*. Wiley, New York, 1976.
- [7] C. Daniel and F. Wood. *Fitting Equations to Data*. Wiley, New York, 1971.
- [8] M. Ekekiel and K. A. Fox. *Methods of Correlation and Regression Analysis*. Wiley, New York, 1959.
- [9] S. Feiner and C. Beshers. Worlds within Worlds: Metaphors for Exploring n-Dimensional Worlds. *Proceedings of UIST '90 (ACM Symp. on User Interface Software)*, pages 76–83, 1990.
- [10] R. A. Fisher. *The Design of Experiments*. Hafner, New York, 9th edition, 1971.
- [11] F. R. Immer, H. K. Hayes, and Le Roy Powers. Statistical Determination of Barley Varietal Adaption. *Journal of the American Society of Agronomy*, 26: 403–419, 1934.
- [12] J. LeBlanc, M. O. Ward, and N. Wittels. Exploring N-Dimensional Databases. In *Proceedings of Visualization '90*, pages 230–237, Los Alamitos, California, 1990. IEEE Computer Society Press.
- [13] T. Mihalisin, J. Timlin, and J. Schwegler. Visualization and Analysis of Multi-variate Data: A Technique for All Fields. In G. M. Nielson and L. Rosenblum, editors, *Proceedings of Visualization '91*, pages 171–178, Los Alamitos, California, 1991. IEEE Computer Society Press.
- [14] T. Mihalisin, J. Timlin, and J. Schwegler. Visualizing Multivariate Functions, Data, and Distributions. *Computer Graphics and Its Applications*, 11: 28–35, 1991.
- [15] T. Mihalisin, J. Timlin, J. Schwegler, E. Gawlinski, and J. Mihalisin. Visual Analysis of Very Large Multivariate Databases. In *Proceedings of the Statistical Graphics Section*, Fairfax, VA, to appear. American Statistical Association.
- [16] J. H. Sheesley. Use of Factorial Designs in the Development of Lighting Products. In R. D. Snee, L. B. Hare, and J. R. Trout, editors, *Experiments in Industry*, pages 47–57. American Society for Quality Control, Milwaukee, 1985.
- [17] E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, Connecticut, U.S.A., 1983.
- [18] E. R. Tufte. *Envisioning Information*. Graphics Press, Cheshire, Connecticut, U.S.A., 1990.
- [19] J. W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, Reading, Massachusetts, U.S.A., 1977.
- [20] J. W. Tukey. *The Collected Works of John W. Tukey*. Chapman and Hall, New York, 1984.
- [21] P. A. Tukey and J. W. Tukey. Graphical Display of Data Sets in 3 or More Dimensions. In V. Barnett, editor, *Interpreting Multivariate Data*, pages 189–275. Wiley, Chichester, U. K., 1981.