CONDITIONING PLOTS AND DESIGNED EXPERIMENTS

Jeffrey S. Pontius

John W. Slocombe

John E. Boyer, Jr.

Recommended Citation

CONDITIONING PLOTS AND DESIGNED EXPERIMENTS

Jeffrey S. Pontius	Department of Statistics,
	Kansas State University, Manhattan KS 66506
John W. Slocombe	Department of Biological and Agricultural Engineering,
	Kansas State University, Manhattan KS 66506
John E. Boyer, Jr.	Department of Statistics,
	Kansas State University, Manhattan KS 66506

Abstract

Conditioning plots (coplots) are useful graphics for displaying values of response variables conditional on the values of given (conditioning) variables. We present a principles guide for construction of coplots when the data or statistics come from studies based on designed experiments, and illustrate the usefulness of these coplots in interpreting results. We have found coplots to be useful in our statistical consulting work, and illustrate our approach so that others may find them useful. Coplots in traditional and in trellis displays are provided.

1 Introduction

Evaluations of results from designed experiments can be enhanced by inspecting appropriately structured visual displays of the data or estimates. We suggest that visual displays be used along with formal analyses (e.g., anova F–tests) to obtain a deeper understanding of the results of the experiment. Crawley (2002), and Denby, Landwehr, and Mallows (2001) provide examples of using visual displays in association with hypothesis testing in designed experiments. Barton (1999) discusses visual displays in the planning stage of designing factorial and fractional–factorial experiments. Cleveland and Fuentes (1997) discuss conditioning trellis displays that correspond to modeling data from designed experiments.

We advocate the use of conditioning plots (coplots) to visually evaluate results from designed experiments as an integral task with formal statistical evaluations (e.g., hypothesis testing). This use of coplots, in particular, their trellis display versions, is not original with us. However, we have found these displays to be very useful in visually evaluating results from designed experiments, and we set out some guiding principles in constructing coplots for designed experiments. Coplots display values (or estimates) of response variables given the values of covariates (Becker, Cleveland, and Shyu 1996; Cleveland 1993). As applied to designed experiments, coplots are useful in visualizing patterns of response values organized with respect to the experimental design and the primary research hypotheses. In principle, construct the coplot with respect to

- the hierarchy of the experimental design. For example, organize the coplot by factor levels within blocks or by repeated measures factor levels nested within larger experimental units.

- the primary research questions. Emphasize results from factors that are considered to be the major focus of the experiment.

- the data or estimates within their factor levels. For example, arrange the levels of a factor horizontally by increasing level averages.

We illustrate these guidelines with examples from three basic experiments. Several displays may be needed to visualize the results of the experiment from several 'viewpoints', especially if the treatment and/or design structure is complex.

2   Example 1: Completely Randomized Design and Dry Matter Recovery

Our first example (Figure 1) is a good illustration of how important results from an experiment can be visually realized. The coplot is based on the trellis display (Becker, Cleveland, and Shyu 1996), and uses the dot plot graphic within each panel. The experiment was done using a completely randomized (CR) design structure with four factors in a factorial treatment structure. The main objectives were to compare percent dry matter recovery from two types of field mechanical wrapping and five types of storage on three types of hay. Each combination of wrapping and storage was done four times on each hay type. The values displayed are the means of percent dry matter recovered. As is conventional in dot plot graphics, the bottom horizontal axes contain the scales for the response (displayed) values. The bottom axis in each column applies to all panels (vertically) in that column. Because the experimental design is CR, the construction of the coplot reflects the importance of the factors under consideration. The coplot emphasizes the results of wrapping type across the four haying times. Each row in the display corresponds to a storage type. Each column corresponds to a hay type. The hay types are ordered left to right based on overall hay type means. Comparisons of haying times are secondary to comparisons of wrapping and storage types, so these are contained within each storage and hay type panel.

Except for a few cases, the two wrapping types do not seem give noticeably different means of percent hay recovery. Nylon wrapping may produce slightly higher means of percent dry matter under the storage types OSU, OSC, ORU, and ORC. The twine wrapping may have a slight edge in storage type I, but not significantly. Easily detectable is the smaller variability in the means of the prairie hay type as compared to the means of the fescue and alfalfa types. Storage types do not seem to have a major impact on hay recovery. There is some variability in means of percent dry matter across haying times for the fescue and alfalfa hay types. However, the means for the prairie hay type are very consistent across storage and wrapping types, and also produced the larger overall means of percent dry matter recovery.

3   Example 2: Repeated Measures and Alfalfa Yields

Our next example (Figure 2) illustrates the principle of coplot construction based on the nesting of the design structure. The coplot is based on the trellis display, and uses the dot plot graphic within each panel. The experimental design was a randomized complete block (RCB) with one repeated measures factor (month) nested within another repeated measures factor (year). Five alfalfa cultivars were planted in each of the three blocks (Kuehl, 2000). The main objective was to compare the yields (lb/plot) of the alfalfa cultivars under the two harvest time factors.

Because a block is the 'largest' unit in the design, the coplot is organized with three rows of five panels, with a row corresponding to a block. Blocks (rows of panels) are ordered by increasing mean yields per block from the bottom row to the top row. For this experiment, the block means are very similar (19.2, 20.2, 20.2), so there are no visual differences between blocks based on mean yields. Each column corresponds to one of the cultivars. One can think of each dot plot panel as a 'field plot'. Because dot plots are used in each panel, the horizontal axes contain the scales for the yields. The horizontal relationship of the cultivar panels (columns) is ordered by the mean yields of the cultivars. Each panel is partitioned into two sections (indicated by the dot lines) representing the two years (each year is nested within a 'field plot'). The month of cutting (April or May) is indicated by a circle or triangle within each year. The location of the month symbol relative to the bottom axis indicates the yield.

Overall, cultivar mean yields increase as one moves in graphical order (i.e., left to right). One can easily detect that the yields in April are less than the yields in May for all cultivars and for both years. Also, the yields in 1987 are usually less than the yields in 1986 relative to the months of cutting.

Using a split split unit analysis with compound symmetry covariance produced the ANOVA results in Table 1. One observation is that there is a significant 3–way interaction, and there is a significant 2–way interaction between year and month. The 2–way interaction is easier to visualize than the 3–way interaction and can be detected by inspecting the year and month yield relationships within each panel. Each panel can be viewed as a 'slice' of an overall year by month interaction plot. Notice that in many of the panels the April 1987 yields tend to be more distant from the April 1986 yields as compared to the May yields. The 3–way interaction is difficult to visualize from this coplot, and a consolidated plot may be more useful to visualize the 3–way interaction. Another observation is that the 2–way cultivar by year interaction is marginally significant. This 2–way interaction can be detected by visually constructing 'mean yields' within each year (per panel), and then visualizing these 'means' left to right across the alfalfa cultivars.

4   Example 3: Analysis of Covariance and Ascorbic Acid

Prior to the use of trellis displays for coplots, Cleveland (1993) introduced the construction of coplots as illustrated in Figure 3. The top panel displays the ordered values of the given (conditional) variable (here, means of percent ascorbic acid). The dependent panels contain the values of the responses associated with the values in the given panel. The dependent panels are connected to the ordered values in the given panel by graphical ordering. Graphical ordering begins with the left panel in the bottom row, moves to the right across rows, moves vertically to the adjacent row, and then moves across this adjacent row. Graphical ordering continues across subsequent upper rows. The strips between the horizontal dashed lines in the given panel correspond to the dependent panel rows. For example, the third panel from the left in the second row corresponds to the seventh ordered value in the given panel (which is approximately equal to 88 and is within the second given panel strip). Note that the given panel provides a visual indication of the separation of the given values (here, the means of ascorbic acid content of the cultivars). For simple treatment and design structures this 'traditional' coplot construction often facilitates easier visualization of the experimental results, and is typically easier to construct as

compared to trellis displays.

The objective of this experiment was to test whether there were differences in the means of percent ascorbic acid across eleven cultivars of lima beans 'adjusted' for percent lima bean dry matter (covariate). The experimental design was RCB with five blocks. Note that a 'circle' in a panel corresponds to a block. Steel, Torrie, and Dickey (1997) chose to test for parallel regression slopes across the cultivars by using the within cultivar values across blocks. A typical visual display of this test is a scatter plot of all of the values, sometimes labeling the treatments and sometimes displaying the treatment regression lines (see Steel, Torrie, and Dickey p. 437 (Figure 17.2) for an illustration). The scatter plot display can be difficult to interpret when there is a high degree of overlap in values and when the regression lines criss–cross or overlay. The coplot lays out the regression lines separately. Because all of the dependent panels have the same horizontal axis scale, the orientations (perceived slopes) of the lines can be easily compared across the panels. For example, we can easily detect that most of the lines appear to have similar orientations, and that the last ordered cultivar (the cultivar with the largest observed mean) has a line with a much different orientation from the other lines. Also, there appears to be a 'high' correlation between percentages of dry matter and ascorbic acid contents as visually indicated by the density orientation of the response circles.

A test of the null hypothesis of all slopes equal to a zero slope under RCB (Littel *et al*, 1996) produced an $F–value = 12.11$ with a corresponding $p–value < .0001$, indicating that not all slopes are zero, which is evident from the coplot. A test of the null hypothesis of all slopes equal produced an $F–value = 2.12$ with a corresponding $p–value = .0564$, marginally indicating that not all slopes are the same. There is some indication of this from the coplot (i.e., panels associated with the smallest and largest cultivar means), although an overall visual assessment indicates that most of the slopes are very similar in orientation. At this point, one recommendation would be to test for differences in means of ascorbic acid at several values of percent dry matter.

5   Summary

We have presented several illustrations of coplots for visually evaluating results from some basic designed experiments. We have found them to be very useful in revealing aspects of experimental results that may not be as evident when using only 'formal' statistical analyses. S–plus (2001) was used to produce these displays. Basic coplots can be constructed relatively easily using the functions provided in the S–plus distributed software. The software package R contains functions that will produce traditional coplots (e.g., Figure 3) and trellis displays (using the lattice package for R).

6   References

Barton, R. 1999. *Graphical Methods for the Design of Experiments.* Lecture Notes in Statistics #143. Springer.

Becker, R. A., Cleveland, W. S., and Shyu, M.-J. 1996. The visual design and control of trellis display. *Journal of Computational and Statistical Graphics* 5:123–155.

Cleveland, W. S. 1993. *Visualizing Data*. Hobart Press.

Cleveland, W. S., and Fuentes, M. 1997. Trellis display: Modeling data from designed experiments. *Technical Report, Bell Labs*.

Crawley, M. J. 2002. *Statistical Computing: An Introduction to Data Analysis using S–Plus*. Duxbury Press.

Denby, L., Landwehr, J., and Mallows, C. 2001. An exercise in the real world of design and analysis. *The American Statistician* 55: 263–271.

Littel, R. C., Milliken, G. A., Stroup, W. W., and Wolfinger, R. D. 1996. *SAS System for Mixed Models*. SAS Institute.

Steel, R. G. D., Torrie, J. H., and Dickey, D. A. 1997. *Principles and Procedures of Statistics: A Biometrical Approach*. The McGraw–Hill Companies.

——. 2001. *S-Plus 6 for Windows User's Guide*. Insightful Corp.

7   Acknowledgement

Table 1: An ANOVA table for the alfalfa repeated measures experiment (Example 2).

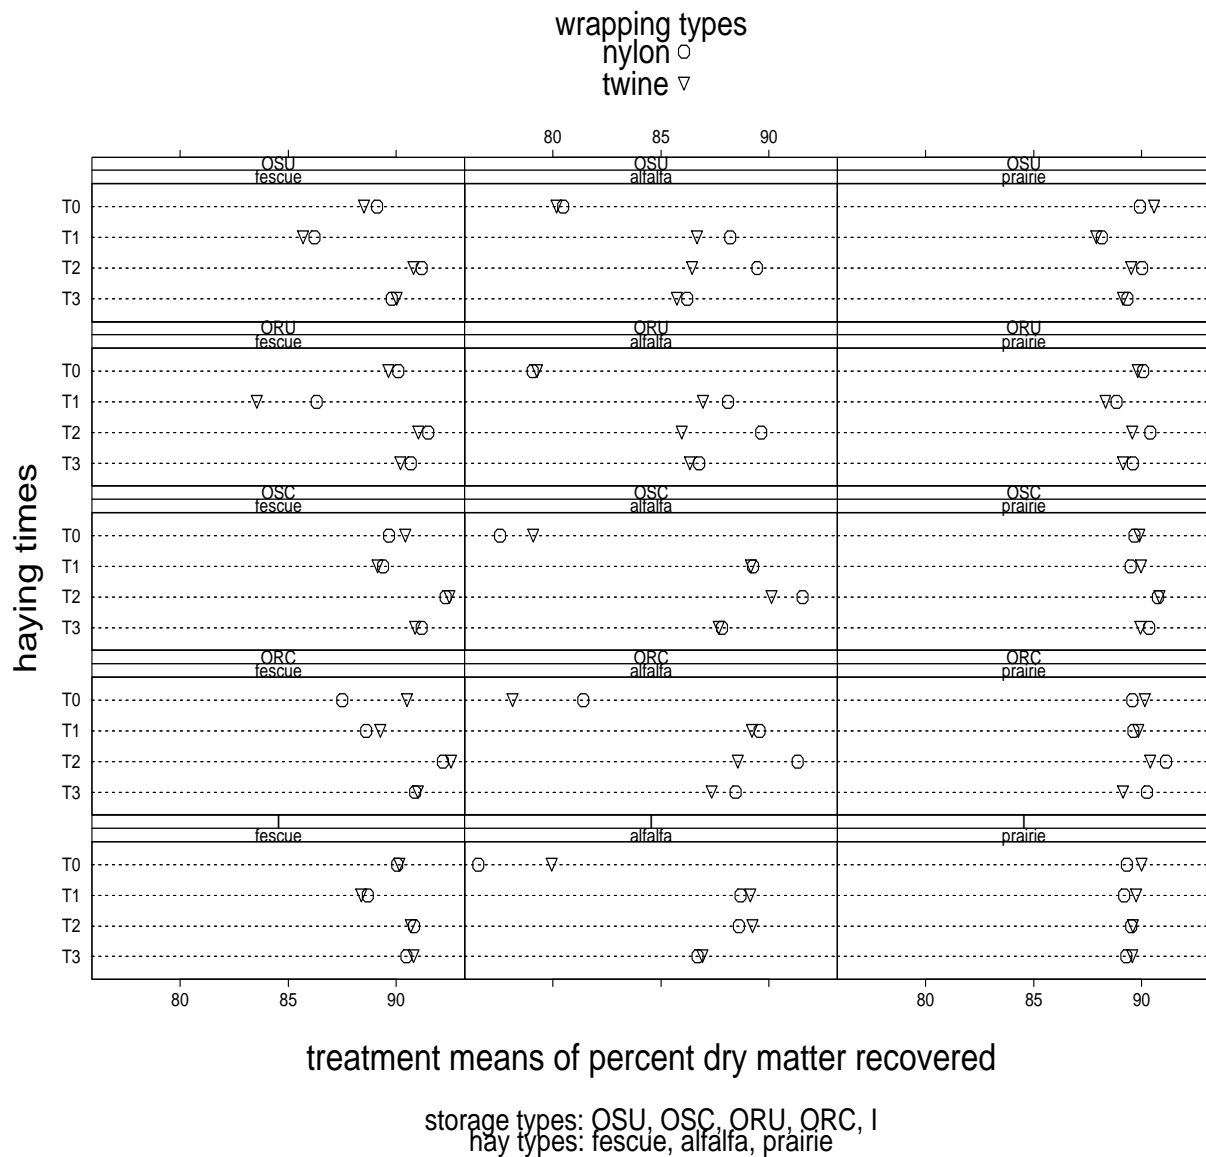| source	numerator df	denominator df	F–value	p–value
cultivar	4	8	2.8	.100
year	1	10	216.2	<.001
cultivar X year	4	10	3.7	.042
month	1	20	238.3	<.001
cultivar X month	4	20	1.4	.258
year X month	1	20	28.0	<.001
cultivar X year X month	4	20	3.5	.025

Figure 1: (EXAMPLE 1) Conditioning plot of the means of percent dry hay recovered given two wrapping methods, three hay types, five storage types, and four sequential haying times (T0 is the earliest time). The focus is on the visual comparison of the wrapping types within the other treatment combinations. The graphical layout is a trellis display that uses dot plots for each panel.
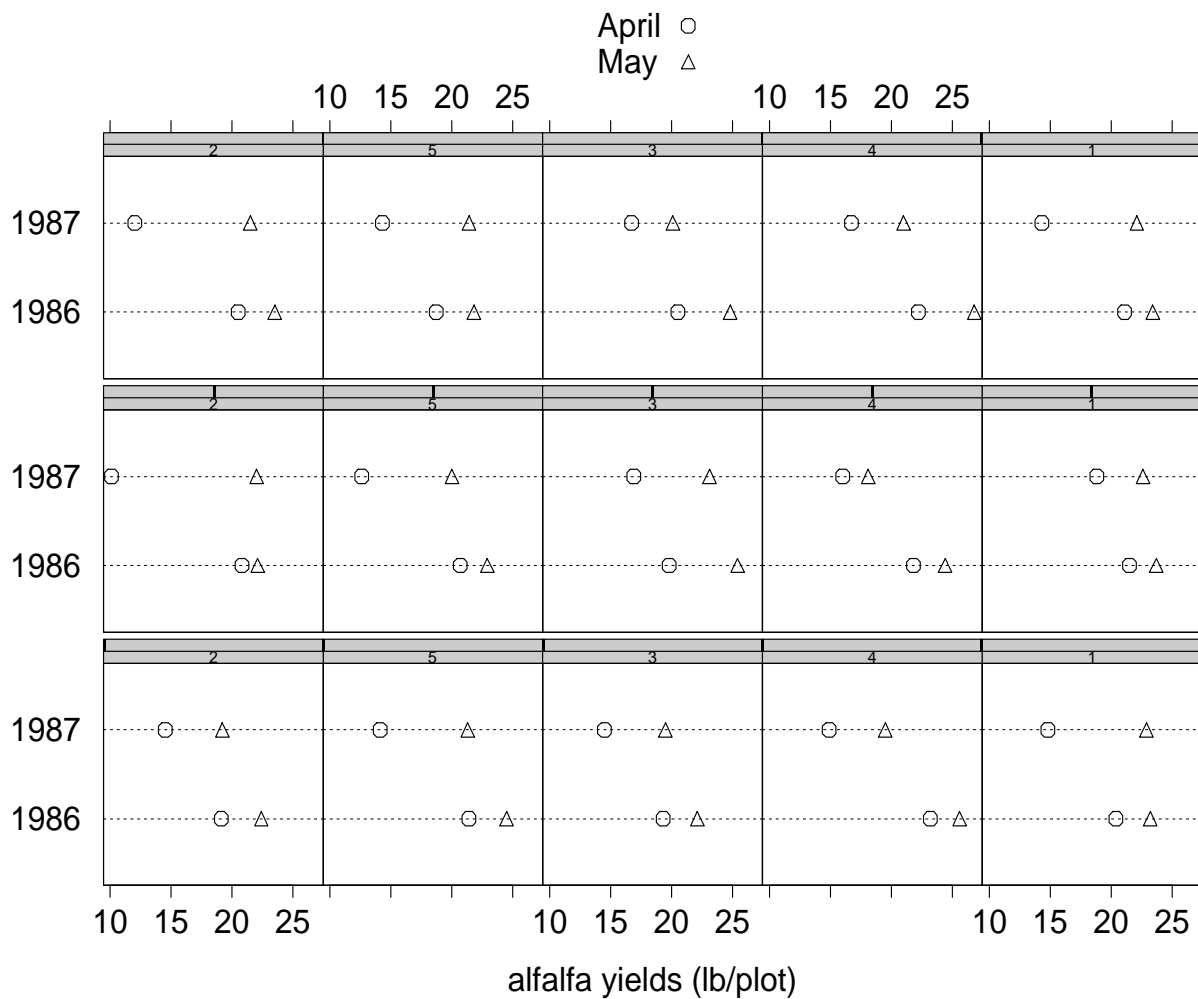
Figure 2: (EXAMPLE 2) Conditioning plot of the yields (lb/plot) of alfalfa. Panels represent the cultivars of alfalfa (numbered) and each row represents a block. Notice how the coplot is organized with respect to the nesting of units in the experimental design. Month of haying is nested within year of haying, which is nested within cultivar and block. The graphical layout is a trellis display that uses dot plots for each panel. Data are from Kuehl (2000).

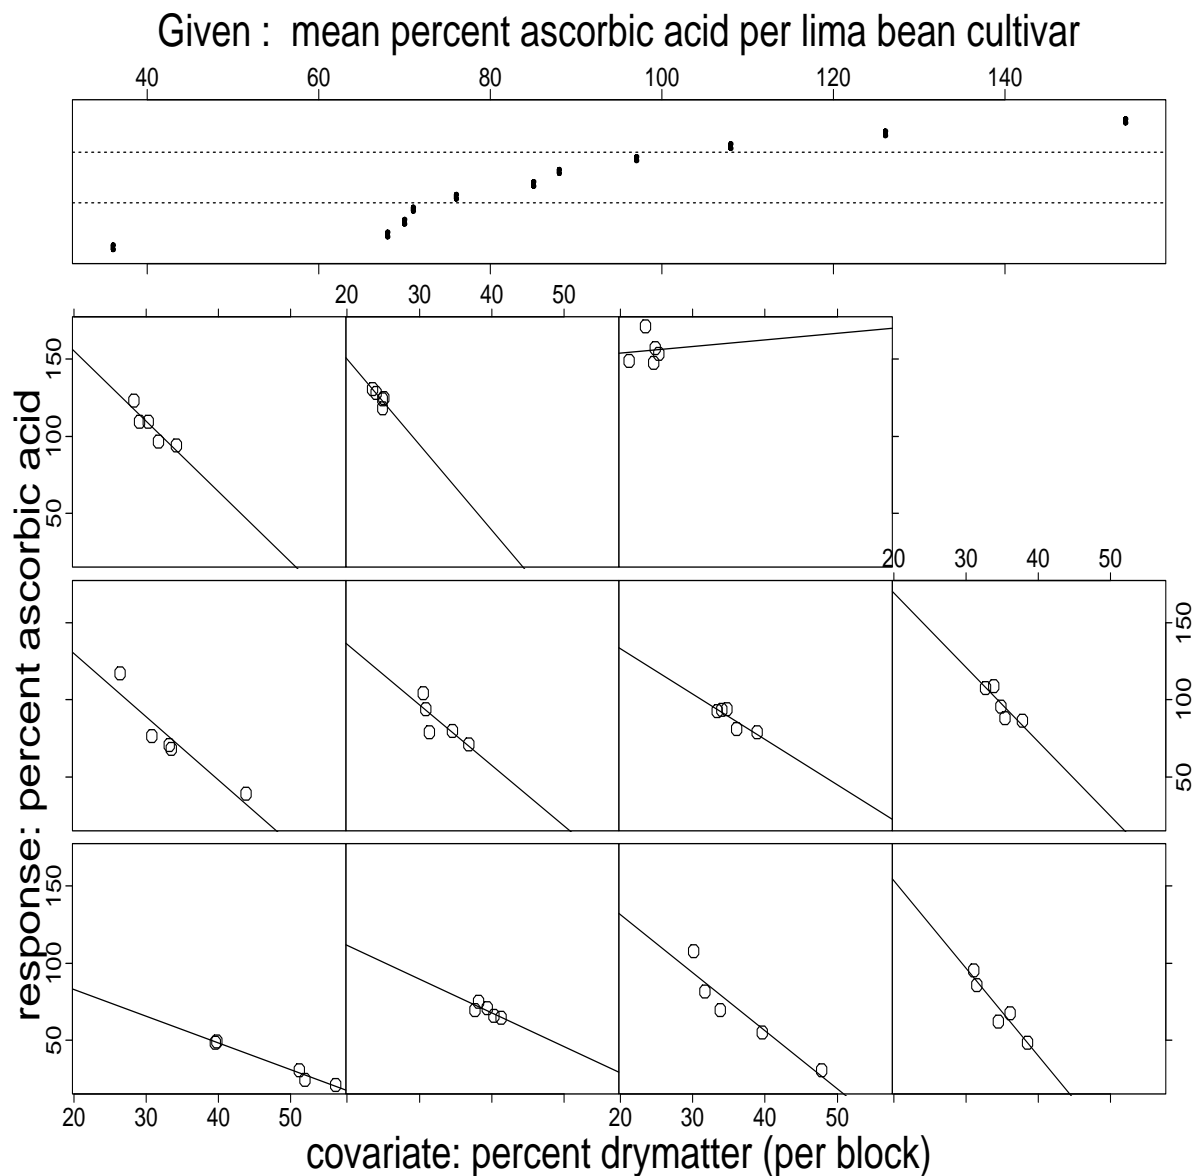Given : mean percent ascorbic acid per lima bean cultivar

Figure 3: (EXAMPLE 3) Conditioning plot of the percentages of ascorbic acid in response to percentages of dry bean matter (covariate). Each panel represents one cultivar where the percentages of ascorbic acid are regressed on percentages of dry bean matter. The cultivars (panels) are ordered according to the means of percent ascorbic acid. The graphical layout is a traditional coplot display that uses scatter plots for each panel. Data are from Steel, Torrie, and Dickey (1997).