

Linked Micromap Plots: Named and Described

By Daniel B. Carr, Anthony R. Olsen,
Jean-Yves P. Courbois, Suzanne M. Pierson,
and D. Andrew Carr

1. Introduction

In this article we describe a new template for the display of spatially indexed statistical summaries and give it a name: linked micromap plots (LM plots). Readers can obtain a general notion about the template from a quick glance at the three examples in this article. The LM template emerged from our environmental graphics research and has quickly jumped to other applications such as the display of federal statistics summaries. We foresee a wide range of applications that extend even to the study of gene communications networks. Our definitive paper on the template emphasizes diverse environmental examples and is nearing completion. Our enthusiasm is such that we want to promote the template immediately. Thus we are adapting some of our material for this article.

LM plots resulted directly from Tony Olsen's challenge to link row-labeled plots (Carr 1994a, Carr 1994b, and Carr and Olsen 1996) to large maps. Micromaps were our linking solution, but we immediately observed that micromap sequences were informative without the large maps. Our first big presentation combining micromaps and large maps (Olsen et al, 1996) was a 4 x 8 foot poster. While we have shown this poster summarizing millions of observations at many conferences, we have been slow to publish. What journal would publish a large poster? The newsletter article by Carr and Pierson (1996) and a paper by Carr (1997) provide the first published example LM plots. Neither article provides a name for the template nor calls attention to the scope of design variations and applications. This sequel does both.

The structure of the article is as follows. Section 2 describes the LM template in the context of three examples. Section 3 revisits the examples to make comments about plot interpretation and continuing design issues. Section 4 comments on the newness of the template, credits many that have influenced our thinking and closes with indications of software available and challenges for the future.

2. The Template for Linked Micromap Plots

The template for LM plots has four key features. First, LM plots include at least three parallel sequences of panels that are linked by position. As illustrated in Figure 1, one sequence consists of micromaps, one sequence consists of linking legends, and one sequence consists of familiar statistical graphics, in this case times series panels. As in Figure 2, the template can include more sequences of panels as long as it has these three types of panels. Sorting the units of study (nations in Figure 1 and states in Figures 2 and 3) is the second feature. The third feature partitions the study units into panels to focus attention on a few units at a time. The fourth feature links representations for highlighted study units across corresponding panels of the sequences. We discuss these features below. With the three types of parallel sequences in mind, we discuss sorting and grouping of study units first, and end this section with discussions of across panel linking and LM plot labeling.

2.1 Study Units, Sorting, and Perceptual Grouping

The study units in LM plots have names, map locations, and summary statistics. In Figures 1, 2 and 3 the study units are nations, states, and states, respectively. One goal of LM plots is to group information into manageable units for human interpretation. Thus the number of study units in LM plots is typically modest, often 50 or fewer. The desire to represent many more study units suggests a hierarchical organization of multiple LM plots. For example one national LM plot might show summaries for all the U.S. states while multiple state LM plots show summaries for the counties within each state. Since overview summaries can show extrema as well as central tendency, say via boxplots with outliers, overviews do not necessarily hide interesting clues about where to look next.

LM plots sort the study units. Often, plotted summary estimates provide the basis for sorting. The unemployment ratio provides the basis for sorting states in Figure 3. The time series in Figures 1 and 2 are multivariate summaries. In both examples we sort by the time series average value. Many other multivariate sorting options are available. Cleveland (1993) suggests using the median, while Carr and Olsen (1996) discuss additional methods such as spanning tree traversals and sorting by the first principal component. Sorting can also be based on variability, spatial position and many other variables. Different sorting methods can reveal and accentuate different relationships. An interactive setting facilitates rapid re-expression using different sorting criteria. However, the easy ability to re-express LM

plots does not obviate the need to think about the sorting order and the choice of panel representations. For example, carefully sorting two parallel sequences of bar plots may fail to show structure as clearly as a simple scatterplot.

Kosslyn (1994) recommends creating small perceptual groups so that we can focus attention on a manageable number of elements. He cites literature indicating four items as a good number for direct comparison. We often chose to emphasize five items per panel since counting in fives is convenient. In the discussion below we refer to the emphasized items in a panel as highlighted items. In some cases, such as the middle panels of Figure 3, we include six items. This is still consistent with the cartographic literature that recommends using six or fewer classes for classed choropleth maps. Six items makes color discrimination and other tasks substantially more difficult than five. Still six is within reason when space constraints are present as in Figure 3.

Creating small perceptual groups is closely related to the chunking of information that appears in the psychological literature. Creating small perceptual groups is also consistent with the human computer interface mantra of focus and context. In LM plots, an individual panel focuses attention on a few study units while the full sequence of panels provides one facet of context.

Creating small perceptual groups can be done in different ways. Carr, Somogyi, and Michaels (1997), for example, grouped genes by gene function first and then created subgroups with four or fewer genes. Using logical criteria for sorting and grouping is often helpful. The cartographic practice of using gaps between values of the sorting variable to partition elements is often advantageous for bringing out spatial patterns. Different objectives motivate different groupings.

Breaking a long list into smaller perceptual units can simplify visual appearance and provide additional visual entry points that might be of interest. However, many small perceptual units can still constitute a long list. Thus small perceptual units can be grouped into larger perceptual units. The design strategy of Carr and Pierson (1996) reflected in Figure 3 uses a 5-1-5 grouping of panels. To our knowledge there is no general theory concerning iterated grouping for long lists.

2.2 Micromaps

The primary task of micromaps is to show the spatial location that corresponds to a statistical estimate. We use five distinct saturated hues as a rapid link between statistical estimates, study unit names, and spatial locations. The cyclic use of five distinct hues often causes

confusion when people first encounter LM plots. We discuss this further in Section 2.4. This coupling of name and location in micromaps serves as a reminder for some readers and as an educational device for others. When working with states, some people need an occasional reminder. When encountering U.S. ecoregions for the first time, the micromaps are instructive.

Carr and Pierson (1996) indicate that micromaps are often spatial caricatures designed to serve specific purposes. Figures 2 and 3 show simple, generalized state boundaries. The design purposes in the two figures include providing sufficient shape detail for state recognition, preserving neighbor relationships (DC is an exception), and enlarging small states to facilitate perception of color.

Micromaps can handle multiple tasks such as showing supplemental information. In environmental monitoring applications involving small areas, micromap details may include streams and various landmarks. Figure 1 illustrates display of three information layers, foreground, middle ground and background. The foreground layer consists of highlighted nations appearing in saturated color. The middle layer uses light yellow with black outlines to indicate OECD nations that are not highlighted in the particular panel. The background layer uses light gray with white outlines to indicate non-OECD nations.

Micromaps are particularly effective when the sequential highlighting of sorted study units reveals spatial patterns. We can often augment micromaps with contours to bring out spatial patterns. Figures 2 and 3 use light yellow and gray to distinguish states above the median from states below the median. In Figure 3 the light yellow contour serves as a spotlight that emphasizes high unemployment states in the upper half of the plot and low unemployment states in the lower half of the plot. In the lower half of the figure the Midwest, Northern Plains, and Southern Coastal states are readily evident as parts of the low rate contour.

The color encoding in Figure 3 is not immediately obvious to all readers. The double use of color for contouring and highlighting can be confusing at first glance. Still, the distinction between unsaturated and saturated color is easy to make and interpretation is easy once the encoding is understood. Using light yellow to emphasize above the median contour for states at the top of the page and the below median contour at the bottom of the page can also be a confusing encoding. An alternative is to use a light blue or other unsaturated color for the below median contour. We are reluctant to introduce yet another color. We conjecture that the notion of shifting

a light yellow spotlight is easily learned. Our purpose is to introduce a middle ground contour rather than two equal emphasis background contours.

Figure 3 calls attention to just two contours, one above and one below the median. This reinforces the statistical concept of median and targets a broad portion of the public that has at least some interest in statistical summaries. Examination of the parallel dot plot panels shows many dots close to the median. The display of other contours would be more consistent with gaps in the unemployment rate distribution.

We indicate three color options of many for showing more contours in Figure 3 to sophisticated audiences. The first uses a shifting light yellow spotlight to focus on the region that combines highlighted units in panels immediately above and below the current panel. The second uses a multicolor spotlight to cover units from more panels. The third approach uses a different color for each state. One such pattern starts with a spectral hue sequence for the panels. Warm colors are on top, bright yellow represents the median and cool colors are below. A lightness ramp within each panel attempts to distinguish the states. Distinguishing states is difficult with fifty-one color schemes, so at present we retain our five distinct hue approach.

Readers might be surprised at the mention of spectral ordering. After citing extensive literature arguing against spectral order, Brewer (1997) cites perceptual studies indicating that some instances of spectral order work quite well. Lightness remains the primary basis for ordering. The spectral order works when used as divergent spectral scheme with bright yellow in the middle. Brewer also discusses color scales for the color blind so the paper is of considerable interest.

In some cases the micromap design itself is less than ideal for showing spatial patterns. For example, Figure 1 illustrates a caricature developed to show OECD nations. The topological distortion and use of two insets is not conducive to properly observing spatial patterns. The Figure 1 micromap sequence might even be deleted since spatial location is not central to the story.

Our notion of micromaps is meant to be general. The examples here involve area representations. Some environmental applications involve monitoring sites that are represented as points. A map may be something other than areas or points on the surface of the earth. For example locations might be position within a building, nodes or links in a formal graph, or even a position in a transition matrix.

Students at George Mason University have developed

some of their own micromap variations. A student in a Statistical Graphics and Data Exploration class won an external poster competition by showing sequences of Virginia maps overlaid with pie glyphs at county centroids. The pie glyphs represented crime rates for different classes of crimes summarized at the county level. (The class did not promote pie glyphs for making comparisons). A student in a Scientific and Statistical Visualization class provided a much better example. He redesigned a statistical summary from World War I concerning the effects of mustard gas. The micromaps were caricatures of the human body and clearly showed the susceptibility of exposed and moist locations. Micromaps can take many forms.

2.3 Statistical Summary Panels

The statistical summary panels can take many forms such as dotplots, barplots, boxplots, times series plots, scatterplots, cdf plots, perspective views, stereo pairs plots and so on. While most of these plots are familiar, that does not mean there is a lack of graphical design issues to address.

Statistical summary panels are typically small, so overplotting remains a problem even though the number of highlighted elements in a panel is small. In Figure 1, the time series overplot substantially. Since there are no missing values the reader can infer values for hidden points. When there is missing data, as in Figure 2, sometimes the overplotting is not too bad. When something must be done, less than elegant solutions include plotting dots of different size with large dots plotted first, plotting symbols that remain identifiable when overplotted, staggering plotting locations and so on.

Space constraints continually come into play. It is advantageous to keep a LM plot to one page. We often forego Cleveland's (1993) guidance about banking to 45 degrees and are tempted to skimp on labeling. Uncomfortable compromise, of course, is not unique to LM plots.

Scaling and resolution issues are recurrent in statistical summary panels. Figure 1 deals with the scaling and resolution issue in two ways. First, the selected unit of measure is tons per person. This reduces some of the disparity between small and large population countries. Second, the top panel has a different scale than the other panels. This compromise is problematic. It is difficult to compare time series between panels on different scales. Since we ordered the nations by the time series mean it is not necessarily the case that values for a specific year in the top panel are above those in the second panel. A helpful option is to show the times series from all panels

in the middle ground, clipping the series appropriately for the scale of each panel. The LM plot suggests the presence of a different scale by the separation between the two panels where the change occurs but this is too subtle. Some additional labeling or special scale warning convention would be helpful.

Focus and context issues apply to a variety of statistical summary panels as well as to a variety of micromaps. As another summary panel example, it can be advantageous to plot all the points in all scatterplots as a middle layer and then overplot the highlighted points for the particular panel. In some cases a translucent middle layer helps by keeping background grids visible.

While there are many design issues, statistical summary panels can nonetheless accommodate much information that helps in interpretation. Carr and Pierson (1996) discuss the design of confidence bounds shown in Figure 3. Note also the dashed line that represents U.S. average as a reference value. While statistics are beyond some segments of the public, the LM plot design makes the notion of median almost self-explanatory and the difference between the median state value and the national average is readily evident.

2.4 Linking Legends and Visual Guides

The linking legend typically uses close juxtaposition to link the name of each unit of study to a symbol such as a dot. In the three figures dot color links the names to micromap regions and to elements in the statistical summary plots. When the study units have point locations, symbol shape may also serve as a link. In Figure 3, the statistical panels are dot plots. For this case the state names link directly to statistical summary elements by vertical position as well as by color. People often have a difficulty when they first encounter LM plots because saturated color only links across panels highlighting a few units of study. Their prior experience suggests that color will retain the same meaning for the full sequence of micromaps and statistical graphics panels. This adds to the learning curve for LM plots but the learning curve is usually short.

Figure 2 illustrates a different LM plot design that includes visual guides (lines) from names to symbols. Something needed to be done to address over-plotted values for annual estimates. Rather than stagger the plotting locations, the design choice uses connecting lines from the state names and to annual estimates. Extending the lines to the right of the plot and placement of the tic labels between panels enables quick reading of annual values.

2.5 LM Plot Labeling

Labeling is one of the most difficult challenges in graphical design. An ill-chosen word can confuse the reader. Lack of explanation can leave the reader confused or with a totally wrong interpretation. At the same time space is at a premium and words can interfere with the power of the eye brain system to perceive graphical patterns. We do not claim to have a general solution, but attempt to address problems as they arise. In Figure 3 a legend concerning confidence bounds and reference line now appears at the bottom of the panel sequence in which they appear. There is now a label for the median state. Labels also suggest that the light yellow calls attention to states above the median for panels about the median and to states below the median for panels below the median.

Labeling placement can also be used to create perceptual groups. Three columns appear simpler than four columns. In Figure 3 we attempt to combine the micromaps and linking labels into one perceptual unit by centering the label over the two columns of panels. For most people this may make little difference, but if we can invite even a few more people into the world to statistical graphics, that is good.

3. Comments on Plot Interpretation

This article emphasizes the LM template. However, some comments about plot interpretation and remaining design issues for the examples seem appropriate.

3.1 Figure 1

Figure 1 shows an OECD (Organization for Economic Co-operation and Development) time series of per capita CO₂ emissions for energy use. The data on CO₂ emissions for energy use and population size comes from the OECD Environmental Data Compendium 1995, pages 39 and 283 respectively. The population data is incomplete, so we interpolated population values as necessary to produce yearly per capita estimates.

In terms of interpretation, population interpolation is a minor concern. While tabling values in OECD reports provides some pressure for international consistency, this pressure is minor compared to limitations in available methodology, limitations in assessment resources, and pressures of maintaining national images. Thus, comparisons across nations without a deep understanding of the nation specific estimation process can be misleading. With times series data, there is some hope for consistency for individual nations over time.

When first looking at Figure 1, the magnitude of 20 tons of CO₂ emissions per person in the U.S. was a shock.

That the U.S. was not top on the list was also a surprise. However, those familiar with Luxembourg describe special circumstances consistent with high values. The unusual value for Iceland is a rounding artifact with the tabled numerator being represented by a single digit. An encouraging hint of declining values appears for Germany, France, and Sweden. (German unification leads to interesting accounting issues in regard to future improvement.) The figure suggests increasing per capita rates for nations such as Ireland, Japan, New Zealand, Greece and Portugal.

In terms of total emissions, even flat per capita patterns are cause for concern when coupled with increasing population. For example the Mexican population is increasing rapidly. Of course the U.S. population continues to increase due to immigration and U.S. per capita values are much higher than Mexican values. Our perhaps injudicious interpretation is that nations tend to have CO₂ emission styles that are relatively stable for the reporting period, that total emissions are linked to population, and that as far as we know population growth is not under control.

As indicated by Wood (1992) plots reflect some agenda. The agenda behind showing amount per capita was both to reduce the tremendous range of values and to present the information on a personal level. In a chance airplane conversation, a consultant for U.S. utilities looked at the plot and suggested reporting CO₂ emissions per gross domestic product. His agenda was to make the U.S. appear as a waste minimizing (efficient) energy producer. The verbal battle over greenhouse gas emissions will continue.

Developing a micromap for OECD nations was a design challenge. Figure 1, takes a variety of liberties, not only slicing a way most of the Atlantic and enlarging small countries, such as Luxembourg, at the expense of others, but also by using two insets, one for Australia and New Zealand and one for Japan. Is the distortion too much for those familiar with maps of Europe? Will the inset placing Japan on Russia arouse political sensitivities? We have not addressed such issues, but note the micromap will need to be revised to incorporate the three new OECD nations that appear in the 1997 compendium. Revising an already published OECD view (The State of the Environment, OECD 1991, page 134) may provide a solution, but showing small nations is still a challenge.

3.2 Figure 2

Figure 2 re-expresses a portion of a table published on page A-29 in Agriculture Prices (release date January 31, 1994) by The National Agriculture Statistics Ser-

vice. The micromaps call attention to the higher wheat prices in the West Coast and Northern States. Are transportation costs involved? Is the total amount available at different times during the year a major factor determining price factor? The time series indicate missing data. Why is it missing? A big question that jumps out in the graphical representation concerns the mismatch between the marketing year average for each state (footnoted as being preliminary) and the monthly time series values. In some cases the marketing year average for a state is near extreme values for the state, and that implies heavy weighting of specific months. An explanation about the weighting seems appropriate. Likely the weighting explanation is available in related Department of Agriculture documents.











3.3 Figure 3

The primary source for Figure 3 is the Geographic Profile of Employment and Unemployment, 1995, U. S. Department of Labor Bureau of Labor Statistics, Bulletin 2486. Carr and Pierson (1996) propose a LM plot as a replacement for a choropleth map in that document. They discuss the relative merits of choropleth maps and LM plots and we encourage readers to read the article. Figure 3 takes a step further from the previous LM plot, showing contours, the U.S. average as a dashed line, and more labeling. The plot tells a reasonably complete story indicating estimates, estimate precision, estimate importance (the number unemployed), and estimate location. A deeper interpretation item that is not shown concerns the determination of who is excluded from the numerators and denominators in the determination of rates.

4. LM plot history, Connections to other Research and Challenges

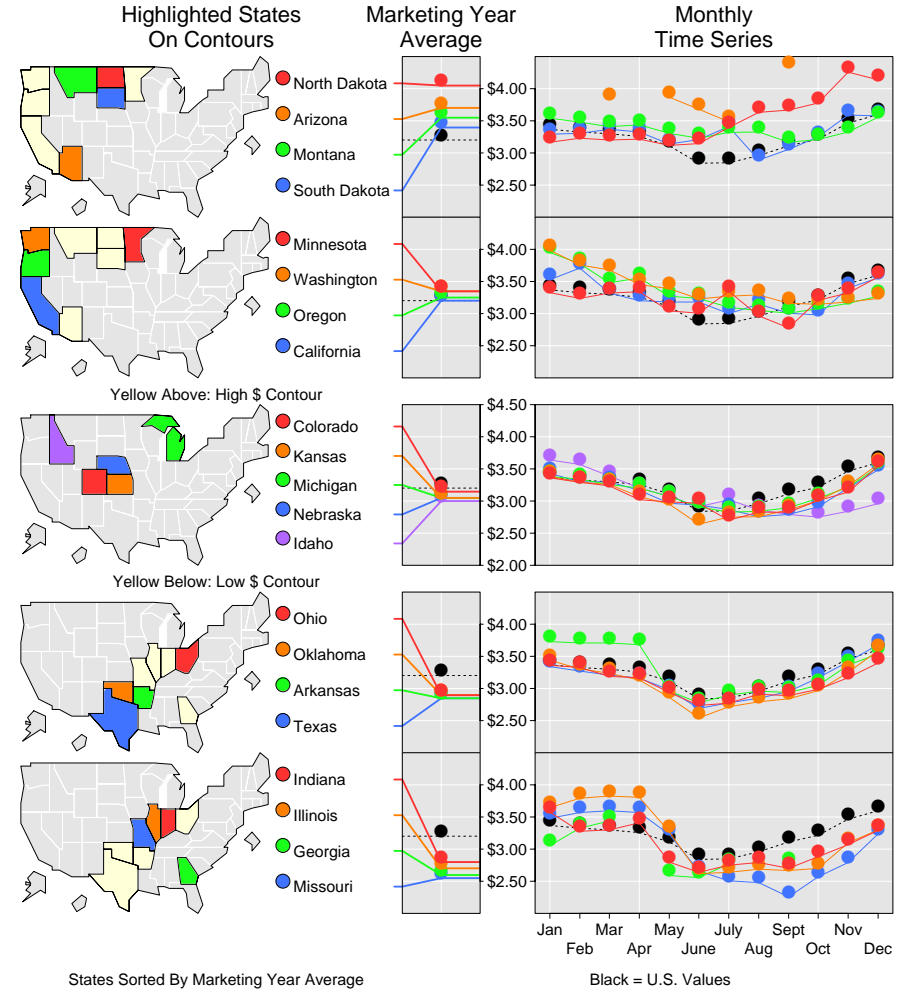
We claim that LM template is new, but there are, of course, many connections to previous graphics and conceptualizations. While we were intrigued by the thumbnail images of Eddie and Mockus (1996), a stronger connection is to the work of Edward Tufte. The LM plots belong to class of graphics that Tufte (1983, 1993, 1997) calls small multiples. In "The Visual Display of Quantitative Data," his eloquent description of well-designed small multiples include phrases such as "inevitably comparative", "deftly multivariate", "efficient in interpretation", and "often narrative in content." We designed LM plots with the hope that such phrases would apply. In "Visual Explanations," Tufte calls particular attention to explanatory power of parallelism. While our use of parallelism precedes this book, Tufte's

Annual CO2 Emissions From Energy Use Units = Tons Per Person



80 84 88 92
Year

Wheat Prices Received By State: 1993 Dollars Per Bushel



Labor Force Statistics By State, 1995 Average

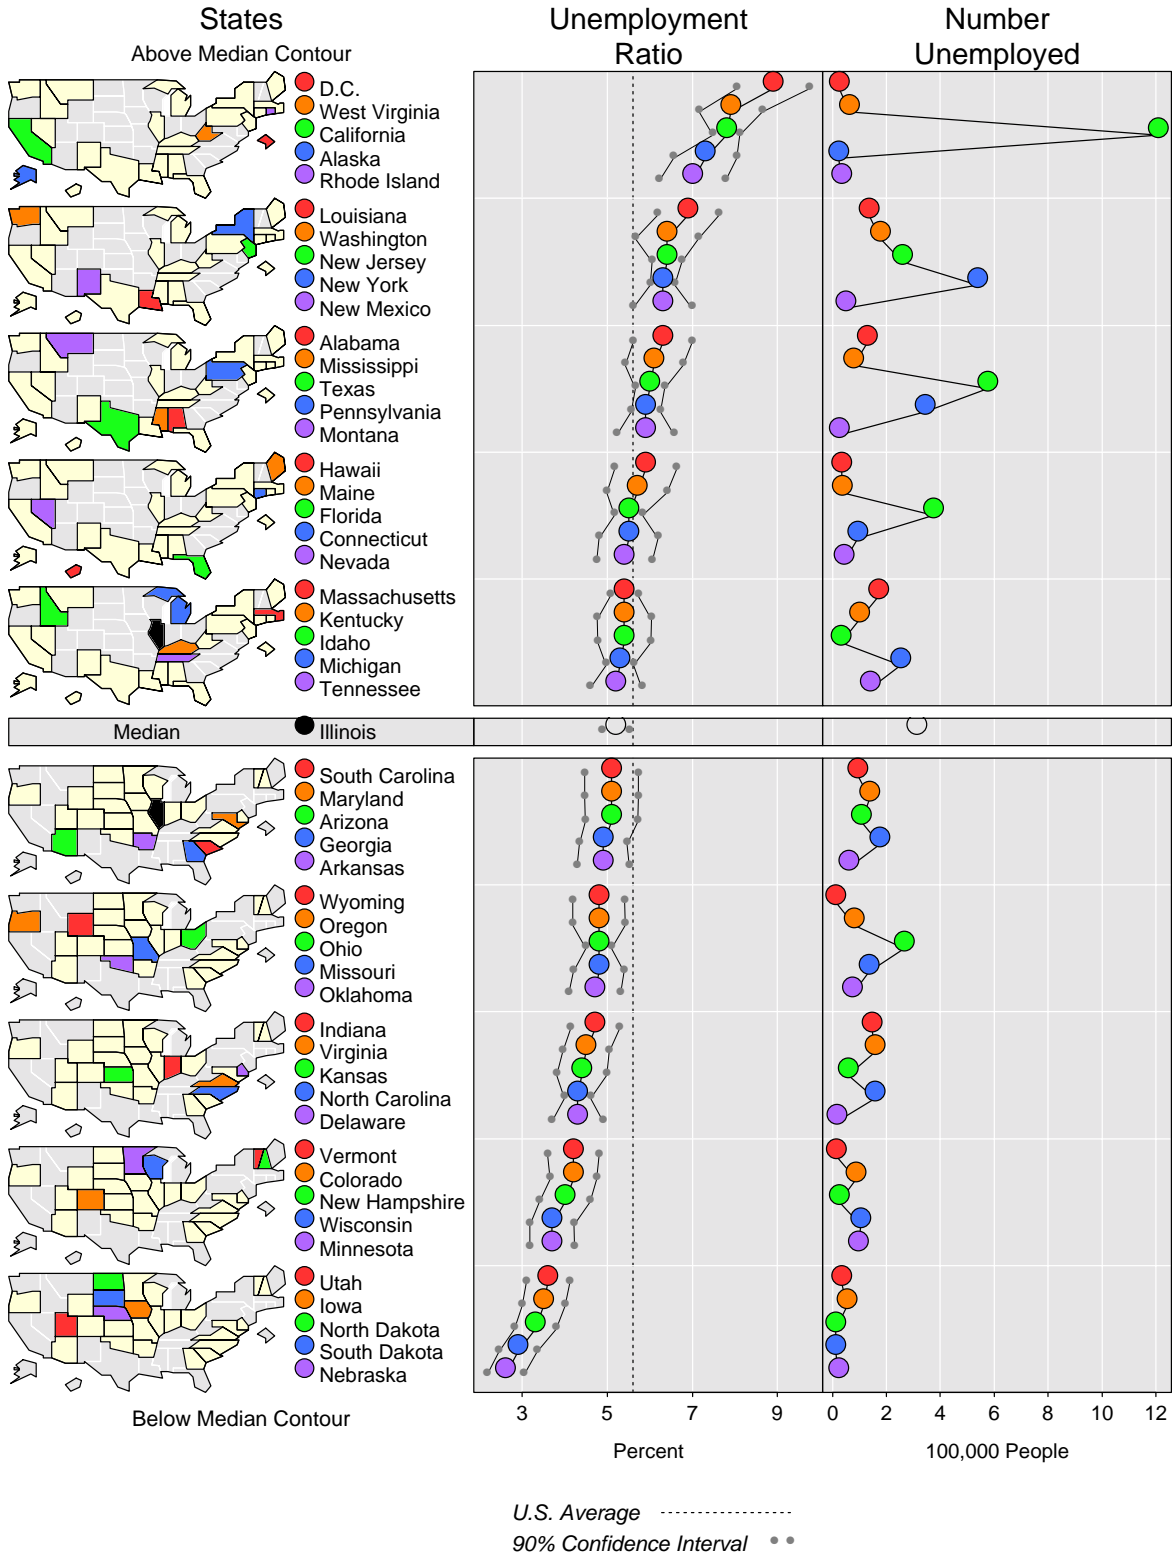


Figure 3. Linked Micromap of 1995 unemployment figures taken from the Bureau of Labor Statistics.

earlier examples may well have guided us to make parallelism a fundamental part of the LM plot design.

As indicated in the introduction, LM plots emerged as a way of linking row-labeled plots to maps. The row-labeled plots in turn build upon Cleveland (1985) and Cleveland and McGill (1984). In fact the development of row-labels plots was part of an effort to encourage EPA staff to use Cleveland's dotplots in EPA graphics. Dissatisfaction with the look of early S-plus dotplots (see Cleveland 1993b for an example) and the promise of multiple panel layouts for expressing complex tables as plots lead to development of new S-plus functions (Carr 1994a and Carr 1997). While the row-labeled plot development was independent of the Trelis graphics development, there are similarities. This is not surprising since Cleveland's design ideas were important in both and S-plus was a common computing environment.

The linking of maps and statistical graphics also builds upon the work of Monmonier (1988) who connected contemporary methodology from cartographic and statistical graphics communities. One of his many interesting examples has a map on the left, labels in the middle and bar plots on the right. This example is a precursor to our LM plots. We have followed Monmonier's work over the years. The statemap caricature in Carr and Pearson (1996) was specifically inspired by a more recent reading of Monmonier (1993) and adapted from coordinates that he graciously supplied.

Despite connections to other research, we claim the LM template is new. Many people have said "There is nothing new under the sun." The truth of this statement always depends on ignoring distinctions. Presumably the speakers of the statement are not identical but if so they collectively get just one vote. We note that the use of parallel sequences of small multiples is relatively uncommon and have called attention to defining features in Section 2. Ultimately others will have to judge the distinctiveness of specific examples and the template in general. What is most important, however, is not the newness of the template, but rather its utility, community awareness of its relative merits, plot production convenience, and statistical graphics literacy.

Plot production convenience remains a big issue. If LM plots are to be used they need to be easily produced. The general S-plus tools we developed (anonymous ftp to `galaxy.gmu.edu` and change directory to `pub/dcarr/newsletter/lmplots`) are flexible building blocks but not easy push button tools. The software also includes a Visual Basic front end that Andrew Carr developed to simplify production of LM plots sim-

ilar to Figure 3. This is a start toward simple production. Much work remains to design micromaps for new applications and to develop software that makes it easy to produce a wide range of LM plots. Much research is appropriate concerning compromises and variations that are motivated by plot purpose, audience, specific data and metadata.

The other big recurrent issue is statistical literacy. Carr and Pierson (1996) suggest that if federal agencies distribute estimates with confidence boundaries, then the Web literate public will grow comfortable with the general idea. Similarly, medians and other statistics can become familiar. A big challenge is to start the ball rolling with federal statistical graphics distributed on the Web (see Carr, Valliant and Rope 1996), a topic to be revisited in future articles.

Acknowledgements

The majority of the work behind this paper was supported by the EPA under cooperative agreement No. CR8280820-01-0. Some facets of this work have been supported by BLS, NASS, and NCHS. The article has not been subject to review by EPA, BLS, NASS, and NCHS so does not necessarily reflect the view of the agencies, and no official endorsement should be inferred.

We wish to thank Wing K. Chong who helped in developing the OECD micromap and the many people commenting at our past presentations.

References

- Brewer, C. A. (1997), "Spectral Schemes: Controversial Color Use on Maps," *Cartography and Geographic Information Systems*, Vol. 24, No. 4, pp. 203-220.
- Carr, D. B. (1994a), "Converting Plots to Tables," Technical Report No. 101, Center for Computational Statistics, George Mason University, Fairfax, VA.
- Carr, D. B. (1994b), "A Colorful Variation on Boxplots," *Statistical Computing & Graphics Newsletter*, Vol. 5, No. 3, pp. 19-23.
- Carr, D. B. (1997), "Some Simple Splus Tools for Matrix Layouts," Bureau of Statistics Statistical Note Series, No. 42.
- Carr, D. B. and A. R. Olsen (1996), "Simplifying Visual Appearance By Sorting: An Example Using 159 AVHRR Classes," *Statistical Computing & Graphics Newsletter*, Vol. 7 No. 1 pp. 10-16.

Carr, D. B., R. Somogyi and G. Michaels (1997), "Templates for Looking at Gene Expression Clustering," *Statistical Computing & Graphics Newsletter*, Vol. 8, No. 1, pp. 20-29.

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

Carr, D. B., R. Valliant, and D. Rope (1996), "Plot Interpretation and Information Webs: A Time-Series Example From the Bureau of Labor Statistics," *Statistical Computing & Graphics Newsletter*, Vol. 7, No. 2, pp. 19-26.

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

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

Cleveland, W. S. (1993b), "Display Methods of Statistical Graphics," *Journal of Computational and Graphical Statistics*, Vol. 2., No. 4, pp. 327.

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

Eddy, W. F. and A. Mockus (1996), "An Interactive Icon Index: Images of the Outer Planets," *Journal of Computational and Graphical Statistics*, Vol. 5., No. 1, pp. 101-111.

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

Monmonier, M. (1988), "Geographical Representations in Statistical Graphics: A Conceptual Framework," American Statistical Association 1988 Proceedings of the Section on Statistical Graphics, American Statistical Association, Alexandria VA. pp. 1-10.

Monmonier, M. (1993), *Mapping It Out*, The University of Chicago Press, Chicago, IL.

Olsen, A. R., D. B. Carr, J. P. Courbois, and S. M. Pierson (1996), "Presentation of Data in Linked Attribute and Geographic Space," Poster presentation, 1996 ASA Annual Meeting, Chicago, IL.

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

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

Tufte, E. R. (1997), *Visual Explanations*, Graphics Press, Cheshire, CT.

Wood, D. (1992), *The Power of Maps*, The Guilford Press, NY.

Dan Carr
*Institute for Computational Statistics
and Informatics*
George Mason University
dcarr@galaxy.gmu.edu

Tony Olsen
*EPA National Health and
Environmental Effects
Research Laboratory*
tolsen@mail.cor.epa.gov

Pip Courbois
Oregon State University
courbois@stat.orst.edu

Suzanne M. Pierson
OAO Corporation
spierson@mail.cor.epa.gov

D. Andrew Carr
Bureau of Labor Statistics
carr_a@bls.gov

