DISCUSSION ARTICLE

Why Tables Are Really Much Better Than Graphs

Andrew GELMAN

The statistical community is divided when it comes to graphical methods and models. Graphics researchers tend to disparage models and to focus on direct representations of data, mediated perhaps by research on perceptions but certainly not by probability distributions. From the other side, modelers tend to think of graphics as a cute toy for exploring raw data but not much help when it comes to the serious business of modeling. In order to better understand the benefits and limitations of graphs in statistical analysis, this article presents a series of criticisms of graphical methods in the voice of a hypothetical old-school analytical statistician or social scientist. We hope to elicit elaborations and extensions of these and other arguments on the limitations of graphics, along with responses from graphical researchers who might have different perceptions of these issues.

Key Words: Statistical communication.

THE BENEFITS AND LIMITATIONS OF STATISTICAL GRAPHICS

My purpose in writing this article is to elicit lively discussion of the uses of graphical methods in statistical analysis.

Graphs tend to be ignored or underused in much of the literature of statistics and applied fields (see, e.g., Gelman, Dodhia, and Pasarica 2002; Kastellec and Leoni 2007), and the literature on graphical methods is small and is mostly separate from the rest of statistics. The related field of data visualization has become increasingly prominent in digital communication, and the arts, but there the focus is typically on eye-catching design rather than on conveying statistical information.[1]

Here I would like to stimulate considerations of the connections between graphics and more formal statistical analysis, along with a serious discussion of the drawbacks of visual

---

Andrew Gelman is Professor, Department of Statistics and Department of Political Science, Columbia University, New York, NY 10027 (E-mail: *gelman@stat.columbia.edu*; url: *http://www.stat.columbia.edu/~gelman/*).

[1]For example, when the influential (and interesting) Flowing Data blog published a list of the "5 Best Data Visualization Projects of the Year" (Yao 2008), a debate ensued over whether data visualizations should aim for transparent communication of information (Gelman 2009b, 2009c) or for visual novelty and beauty (Yao 2009).

presentation of quantitative information: if graphs really are so wonderful and underused in applied statistics (as I in fact believe), what is holding people back from integrating them much more into data analysis?

Following the revolution begun by Tukey (1972, 1977) and continued by Chambers et al. (1983), Cleveland (1985), and Tufte (1983, 1990), graphical methods for exploratory data analysis have generally been recognized to be a useful first step in any statistical study. Beyond this, though, there is disagreement, with the dominant strain of applied researchers (at least in social science) feeling that, when the serious models come out, it is time to put the graphical toys away. From the other direction, researchers in statistical graphics often disparage models and focus on direct representations of data.

There is a lot of valuable work combining analytical modeling and graphical display (for a classical example, consider the article by Daniel (1959)), but in much of the published work in political science, economics, sociology, and other areas, graphics have little if any serious role, being used to display some simple data summaries and never seen again, with important findings displayed in tabular form.

Those of us who believe graphing to be important and even essential to research would be well advised to think hard about why visual displays are not used more extensively in serious applied research. To this end, this article presents a series of attacks on graphical methods in the voice of a hypothetical old-school analytical statistician or social scientist. Although this originated as an April Fool's blog entry (Gelman 2009a), I believe these are strong arguments to be taken seriously—and ultimately accepted in some settings and refuted in others. I welcome elaboration and discussion of these points by statisticians and statistically-minded researchers in applied fields. I have my own answers to some of these objections but do not present them here, in the interest of presenting an open forum for discussion.

The arguments I lay out are, briefly, that graphs are a distraction from more serious analysis; that graphs can mislead in displaying compelling patterns that are not statistically significant and that could easily enough be consistent with chance variation; that diagnostic plots could be useful in the development of a model but do not belong in final reports; that, when they take the place of tables, graphs place the careful reader one step further away from the numerical inferences that are the essence of rigorous scientific inquiry; and that the effort spent making flashy graphics would be better spent on the substance of the problem being studied.

SOME PROBLEMS WITH GRAPHS

Graphs are gimmicks, substituting fancy displays for careful analysis and rigorous reasoning. It is basically a trade-off: the snazzier your display, the more you can get away with a crappy underlying analysis. Conversely, a good analysis does not need a fancy graph to sell itself. The best quantitative research has an underlying clarity and a substantive importance whose results are best presented in a sober, serious tabular display. And the best quantitative researchers trust their peers enough to present their estimates and standard errors directly, with no tricks, for all to see and evaluate. Let us leave the dot plots, pie

charts, moving zip charts, and all the rest to the folks in the marketing department and the art directors of *Newsweek* and *USA Today*. As scientists we are doing actual research and we want to see, and present, the hard numbers.

To get a sense of what is at stake here, consider two sorts of analyses. At one extreme are controlled experiments with clean estimate and p-value, and well-specified regressions with robust standard errors, where the p-values really mean something. At the other extreme are descriptive data summaries—often augmented with models such as multilevel regressions chock full of probability distributions that are not actually justified by any randomization, either in treatment assignment or data collection—with displays of all sorts of cross-classified model estimates. The problem with this latter analysis is not really the modeling—if you state your assumptions carefully, models are fine—but the display of all sorts of numbers and comparisons that in no way are statistically significant.

For example, consider a research article with a graph showing three lines with different slopes. It is natural for the reader to assume, if such a graph is featured prominently in the article, that the three slopes are statistically significantly different from each other. But what if no p-value is given? Worse, what if there are no point estimates or standard errors to be found, let alone the sort of multiple comparisons correction that might be needed, considering all the graphs that might have been displayed? Now, I am not implying any scientific misconduct—and, to keep personalities out of this, I have refrained from referring to the article that I am thinking about here—but it is sloppy at best and statistical malpractice at worst to foreground a comparison that has been presented with no rigorous—or even approximately rigorous—measure of uncertainty. And, no, it is not an excuse that the researchers actually "believe" their claim. Sincerity is no defense. There is a reason our forefathers developed p-values and all the rest, and let us remember those reasons.

To flip this discussion around, what sorts of graphs do we commonly see in statistics textbooks? Residual plots, influence diagrams, quantile-quantile plots. . . a bunch of cookbook routines that have little to do with exploratory data analysis as it might be practiced. In addition, it does not usually make sense to include diagnostic plots in published articles: if the plot reveals a problem, the model should be fixed, whereas if no problem is found, the graph is typically not very informative and so there is no point in displaying it.

THE POSITIVE CASE FOR TABLES

So far I have explained my aversion to graphs as an adornment to, or really a substitute for, scientific research. I have been bothered for a while by the trend of graphical displays in journal articles, but only in writing this piece right here have I realized the real problem, which is not so much that graphs are imprecise, or hard to read, or even that they encourage us to evaluate research by its "production values" (as embodied in fancy models and graphs) rather than its fundamental contributions, but rather that graphs are inherently a way of implying results that are often not statistically significant. (And all but the simplest graphs allow so many different visual comparisons, that even if certain main effects actually do pass the p-value test, many many more inevitably will not. Some techniques have been developed to display multiple-comparisons-corrected uncertainty bounds, but

these are rarely included in graphs for the understandable reason that they magnify visual clutter.)

But enough about graphs. Now I would like to talk a bit about why tables are not merely a necessary evil but are actually a positive good.

A table lays down your results, unadorned, for the readers—and, most importantly, scientific peers—to judge. Good tables often have lots of numbers. That is fine—different readers may be interested in different things. A table is not meant to be read as a narrative, so do not obsess about clarity. It is much more important to put in the exact numbers, as these represent the most important summary of your results, estimated local average treatment effects, and all the rest.

It is also helpful in a table to have a *minimum* of four significant digits. A good choice is often to use the default provided by whatever software you have used to fit the model. Software designers have chosen their defaults for a good reason, and I would go with that. Unnecessary rounding is risky; who knows what information might be lost in the foolish pursuit of a "clean"-looking table?

There is also the question of what words should be used for the rows and columns of the table. In tables of regressions, most of the rows represent independent variables. Here, I recommend using the variable names provided by the computer program, which are typically abbreviations in all caps. Using these abbreviations gets the reader a little closer to your actual analysis and also has the benefit that, if he or she wants to replicate your study with the same dataset, it will be clearer how to do it. In addition, using these raw variable names makes it more clear that you did not do anything shifty such as transforming or combining your variables before putting them in your regression.

We would do well to take a lead from our most prominent social science colleagues—the economists—who have, by and large, held the line on graphics and have insisted on tabular presentations of results in their journals. One advantage of these norms is that, when you read an econ article, you can find the numbers that you want; the authors of these articles are laying it on the line and giving you their betas. Beyond this, the standardization is a benefit in itself: a patterned way of presenting results allows the expert readers—who, after all, represent the most important audience for journal articles—to find and evaluate the key results in an article without having to figure out new sorts of displays. Less form, more content: that is what tables are all about. If you have found something great and you want to share it with the world, sure, make a pretty graph and put it on a blog. But please, please, keep these abominations out of our scientific journals.

B-B-B-BUT...

Yes, you might reply, sure, graphics are manipulative tricks and tables are the best. But does not the ambitious researcher *need* to make graphs, just to keep up with everybody else, just to get his or her research noticed? It is the peacock's tail all over again—I do not want to waste my precious time making fancy 3-D color bar charts, but if I do not, my work will get lost in the nation's collective in-box.

To this I say, No! Stand firm! Do not bend your principles for short-term gain. We are all in this together and we all have to be strong, to resist the transformation of serious social science into a set of statistical bells and whistles. Everything up to and including ordered logistic regression is OK, and it is fine—nay, mandatory—to use heteroscedasticity-consistent standard errors. But No to inappropriate models and No to graphical displays that imply research findings where none necessarily exist.

Beyond this, even in the short term I think there are some gains from going graph-free. The time you save not agonizing over details of graphs can instead be used to think more seriously about your research. Undoubtedly there is a time substitution: effort spent tinkering with graphs (or, for that matter, blogging) is effort not spent collecting data, working out models, proving theorems, and all the rest. If you *must* make a graph, try only to graph unadorned raw data, so that you are not implying you have anything you do not. And I recommend using Excel, which has some really nice defaults as well as options such as those 3-D colored bar charts. If you are going to have a graph, you might as well make it pretty. I recommend a separate color for each bar—and if you want to throw in a line as well, use a separate y-axis on the right side of the graph.

I am sure there are a lot of other problems with statistical graphics that I have missed. You can take it from here.

[Received September 2009. Revised September 2010.]

REFERENCES

Chambers, J. M., Cleveland, W. S., Kleiner, B., and Tukey, P. A. (1983), *Graphical Methods for Data Analysis*, Pacific Grove, CA: Wadsworth. [4]

Cleveland, W. S. (1985), *The Elements of Graphing Data*, Pacific Grove, CA: Wadsworth. [4]

Daniel, C. (1959), "Use of Half-Normal Plots in Interpreting Factorial Two-Level Experiments," *Technometrics*, 1, 311–341. [4]

Gelman, A. (2009a), "Why Tables Are Really Much Better Than Graphs," *Statistical Modeling, Causal Inference, and Social Science blog*, 1 April. [4]

———— (2009b), "More on Data Visualization, Beauty, etc.," *Statistical Modeling, Causal Inference, and Social Science blog*, 22 April. [3]

———— (2009c), "Doing Graphics the George Orwell Way," *Statistical Modeling, Causal Inference, and Social Science blog*, 17 July. [3]

Gelman, A., Pasarica, C., and Dodhia, R. (2002), "Let's Practice What We Preach: Turning Tables Into Graphs," *American Statistician*, 56, 121–130. [3]

Kastellec, J., and Leoni, E. (2007), "Using Graphs Instead of Tables in Political Science," *Perspectives on Politics*, 5, 755–771. [3]

Tufte, E. R. (1983), *The Visual Display of Quantitative Information*, Cheshire, CT: Graphics Press. [4]

———— (1990), *Envisioning Information*, Cheshire, CT: Graphics Press. [4]

Tukey, J. W. (1972), "Some Graphic and Semigraphic Displays," in *Statistical Papers in Honor of George W. Snedecor*, ed. T. A. Bancroft, Ames, IA: Iowa State University Press. [4]

———— (1977), *Exploratory Data Analysis*, Reading, MA: Addison-Wesley. [4]

Yao, N. (2008), "5 Best Data Visualization Projects of the Year," *Flowing Data blog*, 19 December. [3]

———— (2009), "Narrow-Minded Data Visualization," *Flowing Data blog*, 22 April. [3]