Visual Prompts and Graphical Design: A Framework for Exploring the Design Space of 2-D Charts and Graphs

Vibhu O. Mittal†
Computer Science Department and LRDC
University of Pittsburgh
Pittsburgh, PA 15260

Abstract

Graphical presentations can be very effective in communicating large datasets and patterns, trends and relationships in them. Charts and graphs used in reporting data usually tend to highlight certain aspects and suppress others. In fact, a recent study of several hundred annual reports found that more than 30% of charts in these reports were designed to facilitate inferences favorable to the companies while hindering others. Unfortunately, many of the techniques used to achieve these effects may not be obvious to the average user. One solution to this problem is to make design choices explicit to the user. This paper presents a data analysis interface that educates users by enabling them to explore the visualization space and modifying chart design parameters. This interface is based on an analysis of a corpus of charts and graphs and uses knowledge about a variety of techniques for emphasizing specific trends and/or values shown in 2-D charts and graphs.

Introduction

Graphical presentations can be an effective means of succinctly communicating information about multiple, diverse data attributes and their inter-relationships. Thus, it should come as no surprise that more than eighty percent of all business reports contain graphs, charts or some other form of information graphics (Beattie & Jones 1994; Schmid 1983). However, what is surprising is that a recent study (Beattie & Jones 1994) of the financial reports of 240 companies in Britain–randomly selected from amongst the 500 largest companies–found that almost a third of these reports contained graphs or charts that suffered from *measurement distortion* (Tufte 1983), i.e., they deliberately violated the principle that the representation of data values, as physically measured on the graph's surface, should be directly proportional to the numerical values being presented. Furthermore, three times as many companies manipulated perceptual factors to present an exaggerated favorable impression as the ones that did not.[1]

Quite clearly, reports generated by organizations with a stake in highlighting specific aspects or skewing certain perceptions make it difficult for users to (literally) get the full picture. A typical chart discussed by Beattie and Jones (1994) was found to contain a combination of techniques that make quick, visual interpretations difficult, if not impossible: positionally ambiguous 3-D bars, origin displacements that magnified differences, the use of different colors without any underlying domain mappings, etc. Unfortunately, evidence suggests that only a small fraction of the general population is graphically expert, and few people realize the number of decisions implicit in the design of a graphic presentation (Shah 1995). Since each design decision can influence user perception in a variety of ways, it is essential that as datasets grow in size and complexity that users understand and explore the implications of the various parameters involved in the visualization of these datasets.

This paper presents a library of techniques that can be used in designing information graphics to either highlight (or downplay) aspects of the dataset. These techniques were culled from an analysis of several thousand charts in newspapers, magazines and business/governmental reports. One possible categorization of these techniques is presented in this paper. We discuss some of these techniques, using examples to illustrate the points. We describe our implementation and discuss its utility in helping users understand presentations in general and chart designs in particular. Finally, we conclude with a discussion of future areas of research.

Chart Designs: A Corpus Analysis

There are a number of systems that can design complex, effective and novel graphic presentations, but this paper discusses only 2-D charts and graphs. Systems like

†Current address: JPRC, 4616 Henry Street, Pittsburgh, PA 15213, USA. e-mail: mittal@jprc.com

[1] Among companies that did not have positive results to report, a number of them used the converse tactic of presenting information in a fashion that made it more difficult to compare performance across years or with other companies.

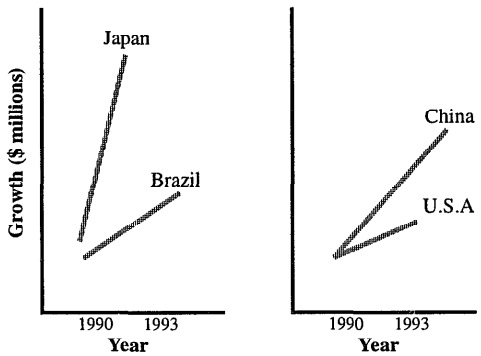

Figure 1: Domains in the same chart are easier to compare.

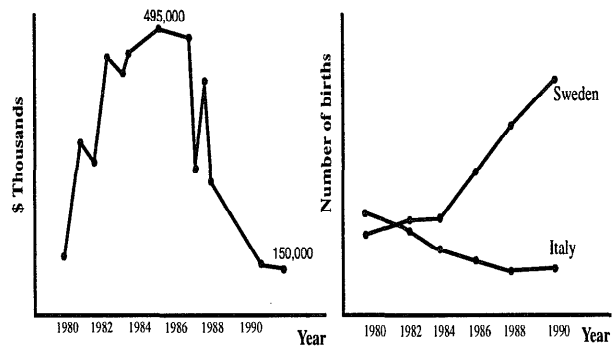

Figure 2: Labels can highlight values or specific points.

SAGE (Roth *et al.* 1994) can automatically design chart based presentations, but all of them are based, in part, on the principle of certain licensed mappings between data types (coordinates, nominals, etc) and graphical attributes (position, lexical items, etc).[2] While some of these systems do reason about certain aspects of the presentation, e.g., ordering items along an axis in lexical (or otherwise sorted) order in the case of certain nominal or quantitative types, none of these systems makes such *design decisions explicit*. Furthermore, none of these systems (as yet) allow the user to reason about, and explore, the gamut of ways in which graphical rhetorical devices such as visual prompts (and combinations of these) could be used to facilitate certain inferences and suppress others in the case of charts and graphs.[3]

To find a set of techniques that can be used to emphasize certain inferences, we analyzed a large number of charts in a variety of reports, books and periodicals. The charts in our corpus used a combination of several techniques to convey their preferred message; these could be categorized into two distinct classes:

- **data set partitioning:** if the point to be made required comparisons within smaller subsets of the domains to be visualized, partitioning the data-set displayed into separate groupings, clusters or regions was a commonly used technique. Consider the charts shown in Figure 1.[4] The charts allow the user to lookup and compare growths in the four countries between 1990 and 1993. However, the allocation of countries to the charts suggests specific groupings, or pairwise comparisons between Japan/Brazil and China/USA. If the dataset had been partitioned differently (e.g., Japan/USA and Brazil/China), the comparisons suggested would be different.

- **visual prompts:** graphical annotations were often used to highlight specific points.

Both types of techniques are often used together to achieve complex effects. Visual prompts or annotations are frequently used as rhetorical devices and include, among others: changes in color, fonts, background, orientation, location, as well as the use of pointing devices such as arrows. A combination of these rhetorical devices within a single presentation can be used with great effect to convey a desired perspective. This section describes some of the prompting techniques and their effects. These techniques are then categorized on the basis of their integration with the original graphic design:

- **planned as part of original design:** this set of techniques includes those that can be used only if planned as part of the original presentation (e.g., a modification to an axis origin)

- **usable as a post-design overlay:** this set includes those techniques and annotations that can be designed as part of an overlay of the original picture.

Note that this is just one possible classification. Another, and possibly more useful one, would be one that categorized these prompts in terms of the *communicative goals* that they help achieve: magnifying small differences, facilitating specific comparisons (differences, ratios, progressions, etc), highlighting trends and specific subsets, suppressing specific data readings, etc. For lack of space, we will not discuss some of the issues that arise in determining that categorization.[5] Our goal here was to build a framework that allowed users to familiarize themselves with the more lower-level design parameters.

[2]For a detailed discussion on the issues in designing graphic presentations and problems arising from violating these mapping rules, see (Bertin 1983; Kosslyn 1989; 1994; Shah 1995).

[3]Marks and his colleagues have also looked at graphical implicatures in the context of network diagrams (Marks & Reiter 1990).

[4]Chart source: (Kosslyn 1994), p. 57.

[5]Kerpedjiev et al (1997) discuss some mappings between communicative goals and graphical techniques in a related project.

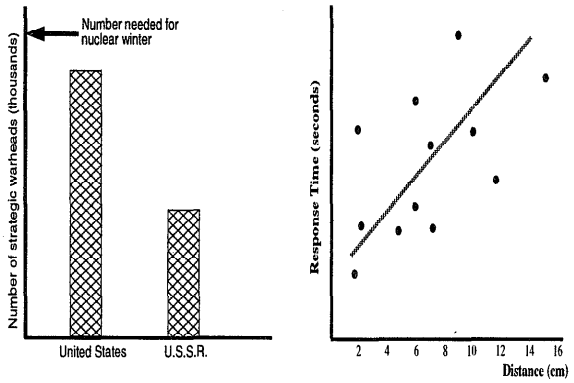

Figure 3: Arrows and trend lines highlight values/patterns.

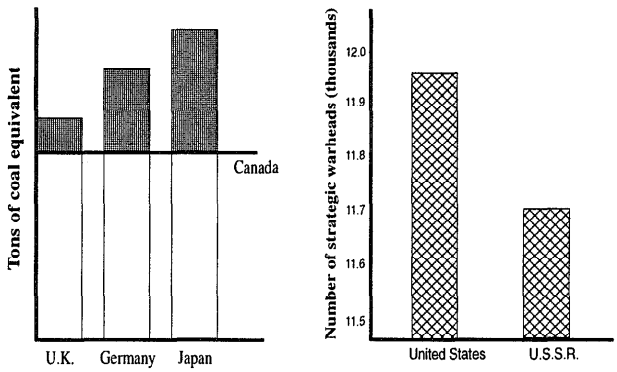

Figure 4: Origin changes facilitate specific inferences.

Using Labels as Visual Prompts

Studies have shown that labelling points in a chart can be both an effective way of drawing attention to that point, as well as communicating the exact value of that point (Culbertson & Powers 1959). Consider the charts shown in Figure 2.[6] The one on the left uses labels to show values corresponding to 1985 and 1990. The implicit assumption in this case is that these values are important for some reason. The chart on the right in Figure 2 shows a case where the labels serve a slightly different purpose: not only do they help identify the countries the lines represent, but they also draw attention to the point where the difference between the two trends is maximal.

Using Arrows or Other Pointing Devices

Arrows, and other similar pointing devices, strongly attract attention (Bertin 1983) and can used to:

- draw attention to specific points in a line or an axis. This is illustrated in Figure 3,[7] where the arrow is used to highlight an important limit value along the Y axis. A secondary goal achieved by using such devices is to imply that the values/objects pointed to are important or critical in some way.
- indicate graphemes or objects being referred to: by selectively pointing to some objects, attention is drawn to those objects relative to the other objects in the graphic.

Using Lines as Visual Prompts

Studies suggest that people believe that lines imposed on scatter-plots show underlying trends in the data (Kosslyn 1994). Clearly, such "trend" lines should not be presented in a chart/graph unless such a trend actually exists in the data. The second chart in Figure 3 shows how a line inserted in a scatter-plot without an actual trend can sometimes lead to incorrect implicatures about the existence of a trend. Line segments are also often used to connect points together, as is done in graphs. Depending on how the segments connecting any two points are drawn, the graph can change its overall appearance significantly even without any changes in data.

Another way in which lines can be used to influence perception of various values is by generating a *reference* line. Consider, for instance, the first chart in Figure 4.[8] It shows energy consumption for four countries. The flat line displaces the origin and draws attention to the portions above it. Both visual comparisons as well as perceptual ratio calculations for the absolute lengths of the bars are hindered by the line. Reference lines inserted in charts can serve many different goals; in general, they are used for comparison with some predetermined level.

Varying The Axes

Axis modification is often exploited by writers to convey a specific impression: almost 70% of the charts studied in (Beattie & Jones 1994) that suffered from presentation biasing involved the use of axis modification in some way. This is particularly insidious, since studies suggest that people tend to overlook changes in the axes more than changes in other factors (such as color or shape of graphical objects in the chart) (Tufte 1983).

There are at least four ways that a presentation can be changed by modifying the axis design:

- by changing the *origin* or by having a discontinuous or truncated axis. Consider the second chart in Figure 4.[9] By truncating the Y axis, the relative difference between the USA and the USSR is magnified and visual computation of ratios is no longer possible. Beattie and Jones (1994) found some companies using this technique to seemingly magnify earnings by factors of as much as 700% compared to their previous year.

[6]Chart sources: (Kosslyn 1994), p. 151, and *The Economist*, 13th April 1991, p. 47.

[7]Chart source: (Kosslyn 1994), p. 211.

[8]Chart source: *United Nations Energy Statistics Yearbook* (1985 figures).

[9]Chart source: (Kosslyn 1994), p. 210.

- by changing the *scales* of the axes.
- by changing the *range* shown; e.g., increasing the range displayed can make scatter-plots appear highly correlated (Cleveland, Diaconis, & McGill 1982).
- by changing the items' *order of presentation*.

Other Visual Rhetorical Strategies

Among other, less frequently used visual rhetorical strategies are the use of color and saturation, line thickness, physical separation or positional prominence, zooming/magnification and the use of titles and/or captions to accompany the charts. Taken together, this set of techniques accounted for about 30% of the biasing techniques seen in our corpus.

It should be noted that our emphasis in this work was not as much on valid or invalid mapping techniques (for instance, whether the length attribute can be used to map coordinate data types–it should not), but on explicating the fact that the various design choices from the set of *valid mappings* are not equal: certain design choices tend to facilitate/suppress certain types of inferences. For instance, color is commonly used to differentiate nominal types. However, coloring a selected subset distinctively from the others tends to distinguish the selected objects and attract additional attention to them. Color can also be used more subtly: *chromostereopsis* is a physiological phenomenon that causes warm colors (such as red or orange) to appear closer than cooler colors (such as green, blue or black).[10] This fact is often exploited by designers to foreground certain domains. A chart, for instance,[11] with two lines in it, one black and the other red, does *not* convey both trends equally effectively: users tend to perceive the red line as being "in front of" the black line and give it more prominence.

Other techniques less frequently seen in our study were the use of positional prominence or physical separation, zooming/magnification, using iconic graphemes (as is done in isotypes), etc. For lack of space, these techniques and their effectiveness are not further discussed here. However we list all the ones we found and attempt to categorize them in the following section.

A Categorization of Focusing Techniques

The set of techniques we came up with based on the analysis of charts in our corpus is shown in Figure 5. As can be seen, there are two major classes of strategies for facilitating certain user inferences, and we have concentrated on using visual prompts. These are further subdivided into those that need to be planned for as part of the original presentation and those that can be generated as part of an overlay *after* the charts have been designed. The overlays are then further divided into those that require significant amounts of coordination with the presentation (and therefore, extra planning) and those that do not.

As mentioned earlier, these techniques could also have been classified under alternative (and depending on one's goals, more useful) categorizations. For instance, if one's goal were to design an automatic presentation system, one could categorize these techniques in terms of the communicative goals that they can help achieve. The problem with that approach lies in the fact that most of these techniques can achieve several different communicative goals, and evaluating the effectiveness of either specific techniques (or combinations of these) for specific goals can be highly context dependent. Since we did not wish to try and characterize the degree of effectiveness of a particular technique for a goal such as `contrast`, we did not attempt to develop such a classification. Furthermore, since most of these techniques are used in combination with other techniques, it is difficult to reason about the relative effects of each of these techniques on a picture.

System Implementation

We have implemented these techniques as part of a small, experimental framework implemented in JAVA.[12] The presentation design constraints were specified in a constraint satisfaction system based on SCREAMER (Siskind & McAllester 1993), but could also have been specified declaratively in a system such as the Functional Unification Formalism (FUF) (Elhadad 1992).[13] FUF is a freely available, widely used Lisp implementation of a constraint based unifier that takes as input a grammar and a partial specification to generate a fully specified form. Unlike a traditional grammar, however, our system collects *all* possible valid realizations of the (partial) input specifications. Input to the system consists of a data set, and the associated data characterization. In response, the system generates a visualization that satisfies the characterization constraints. Furthermore, the system generates *a set of all the alternative choices possible* for all of the design decisions that were involved in the design of the picture. The unused choices are used to generate an interface that can be used to *explore the space of other possible perspectives* on the data set. This is the crucial point, because our goal was to design a system that can generate the list of decisions, along with other valid alternatives, used in designing the presentation in a form that can be used to explore the space of other designs as well.

[10] It is hypothesized to be a consequence of the fact that the lens of the eye refracts different wavelengths of light differently (Allen & Rubin 1981; Travis 1991).

[11] This chart can be found in "Development of moral judgment: A longitudinal study of males and females" by C. Holstein, 1976, *Child Development*, vol 47, pp. 51–61.

[12] URL: *http://www.jprc.com/users/mittal/*.

[13] This is in fact the approach taken in the AUTOBRIEF system (Kerpedjiev *et al.* 1997).

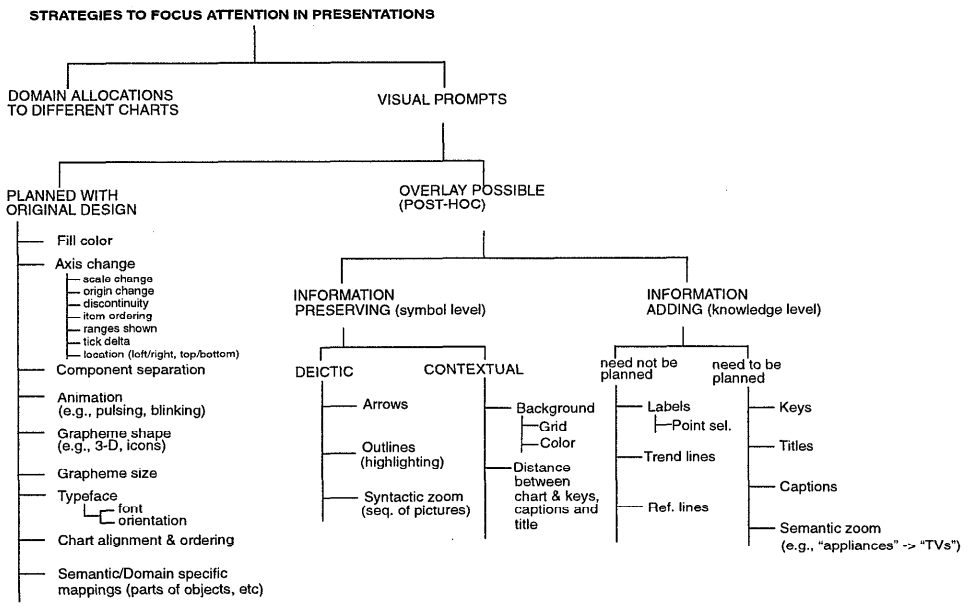

Figure 5: A possible categorization of techniques to focus attention in informational graphics.

Most first-time users, for instance, are unaware of the fact that even simple presentations consisting of a single chart with only a single type of a mark (used to visualize the relationship between two domains) requires over twenty design decisions to be made before it can be rendered on the screen (there are over six decisions just pertaining to the X-axis: x-origin, x-range, x-delta, x-domain, x-truncated?, x-scale, and a number of others pertaining to the framework as well as the grapheme used to present the points). The system generates an interface with radio button menus (for discrete choices) and sliders (for continuous parameters, such as the choice for a numeric axis delta value) to allow the user to explicitly identify the choices used and modify them to see their implications. This framework makes it clear to the user that even a seemingly trivial presentation can have a large number of design decisions implicit in the visualization, many of which can have significant effects on which trends, patterns or clusters in the data set are visible, highlighted or inferred by the users.

Discussion

It is difficult to validate claims or compare systems without a systematic evaluation. In our case, one possible evaluation would be to conduct user studies and see if over time, system users became graphically more sophisticated as compared to a control group exposed to the same set of graphical presentations but without the benefit of this system. Unfortunately, we did not possess the resources needed to conduct such a study.

Instead, we decided to initially focus on using the set of techniques culled from our analysis in building a more explicit or transparent visualization system and help users better understand charts. Using information about these techniques and their applicability conditions, we were able to construct a system capable of displaying the various design decisions that go into chart design. In informal surveys, we found that users were taken aback to find that even supposedly simple presentations required tens of design decisions that potentially affected their perception of the presentation. We also used this system to generate visualizations from different perspectives by varying the parameters in the interface (e.g., look at a data set about house sales from the perspective of the buyer, the seller, the real estate agent, etc.). We found that having an interface that allows the user to interactively pick various design parameters can sometimes be a very useful data navigation and analysis tool, because instead of randomly testing various parameters to see if they generate valid visualizations and then trying to find interesting patterns, this system allows users to see the space of possible design choices and pick extreme values.

We also used the system to generate some of the "classic" alternative presentations discussed in (Huff & Geis 1954; Kosslyn 1994). One of these examples is shown in Figure 6:[14] it shows a data set from two different perspectives, which are conveyed by the titles of the original charts: *"Govt. payrolls up!"* for the first one, and *"Govt. payrolls stable."* for the second one. Other than the titles (which are not generated by the system), the system can be used to

[14]Chart source: a 1938 editorial in the periodical *Dun's Review* discussed in (Huff & Geis 1954), p. 65.

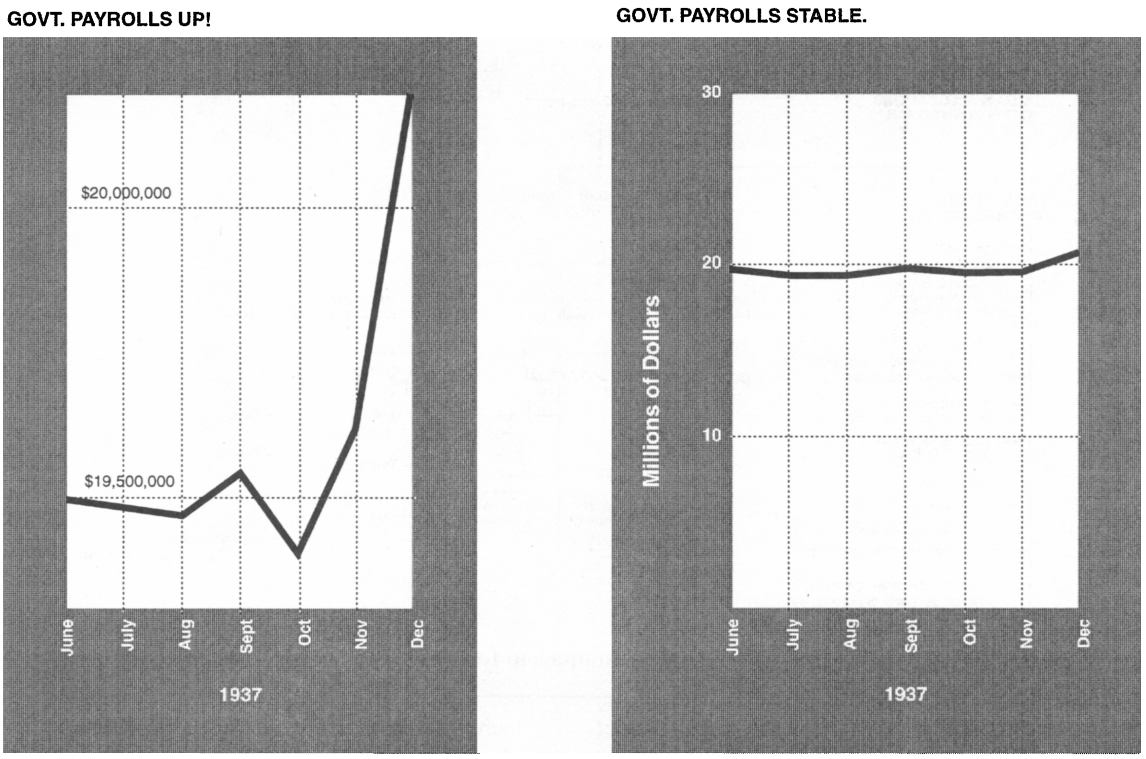

Figure 6: An example of the same data set visualized using different chart design parameters.

set design parameters that will generate the two charts appropriately. While the complete set of design parameters is quite large, some of the ones in which the two charts differ are shown in Figure 7.

We are not making any claims about completeness regarding the techniques presented here; we believe that the current set (as presented in Figure 5) is capable of covering a significant fraction of the variations possible. It is likely to grow as we come across more data. It can be used as a starting point for researchers interested in graph comprehension, perceptual complexity and communicative implicatures. We believe this is an interesting area of research because (unlike natural language or graphics in general), even though charts are so widely used in presenting data, there are only a small number of variations that designers can avail of, which makes the set of techniques amenable to analysis and computational implementation.

Conclusions and Future Work

This paper presents a set of techniques that are often used to facilitate specific inferences in 2-D charts or graphs. This set of techniques was generated by scanning a large number of charts and graphs in a variety of documents, articles and reports. Our goal was to design an interface that made explicit to the reader the various decisions that are necessary in the design of a chart/graph. Our application was motivated by the finding that almost 30% of charts and presentations in financial reports are biased in ways that are not obvious to average users. This work differs from most previous computational approaches to charts/graph design in that the goal here was to make *all* decisions explicit and modifiable by the user; thus, it differs from interfaces such as SAGEBRUSH (Roth *et al.* 1994) in allowing the user to examine and modify very low level details in the design process. Our informal evaluation suggests that this may be an effective approach at making users aware of the numerous design parameters.

Future work may include extending this interface to incorporate results from graph comprehension (Shah 1995) in designing a plan based graphic designer capable of reasoning about goals and effects and using a combination of these techniques to generate sequences of charts and graphs that convey the desired inference. The AUTOBRIEF project is investigating the goal representations and planning processes necessary to generate coherent, effective multi-media documents integrating both graphics and text (Kerpedjiev *et al.* 1997).

Acknowledgments: This work was done while the author was supported under contract DAA-1593K0005 from the Defense Advanced Research Projects Agency (DARPA).

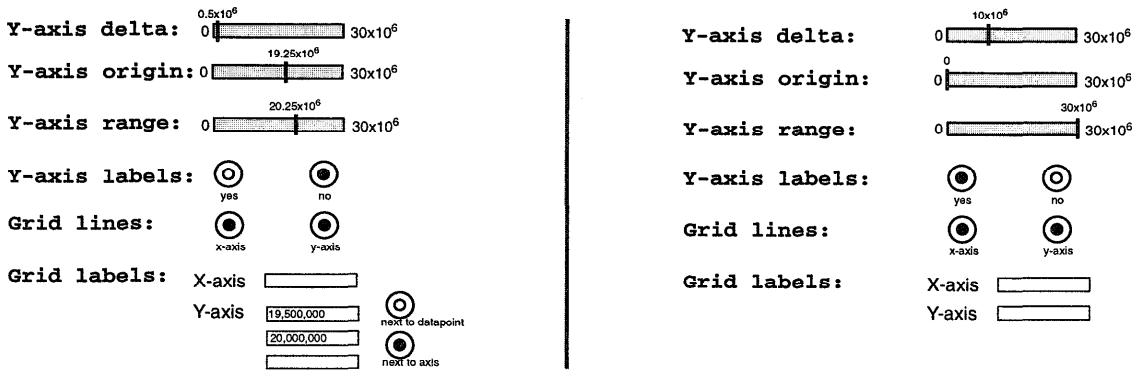

Figure 7: Some design parameters for Figure 6.

The author gratefully acknowledges discussions with members of the SAGE and AUTOBRIEF projects at CMU and the Univ. of Pittsburgh: G. Carenini, N. Green, S. Kerpedjiev, J. Kolojejchick, J. Mattis, J. Moore and S. Roth. The original suggestion to compile a list of possible graphical prompts–which was extended into the hierarchy presented here–came from Johanna Moore; others suggested specific annotations to add to the hierarchy. Some aspects of the SCREAMER implementation of the graphic designer were based on an earlier FUF design implemented for the AUTOBRIEF project by S. Kerpedjiev. I am indebted to all of the members of the SAGE and AUTOBRIEF projects, both for specific inputs on these ideas, as well as providing the overall intellectual environment. However, it goes without saying that any errors in this paper are due solely to me.

References

Allen, R. C., and Rubin, M. L. 1981. Chromostereopsis. *Survey of Opthalmology* 26:22–27.

Beattie, V., and Jones, M. J. 1994. Information design and manipulation: Financial graphs in corporate annual reports. *Information Design Journal* 7(3):211–226.

Bertin, J. 1983. *Semiology of Graphics*. Madison, WI: University of Wisconsin Press.

Cleveland, W. S.; Diaconis, P.; and McGill, R. 1982. Variables on scatterplots look more highly correlated when the scales are increased. *Science* 216:1138–1141.

Culbertson, H. M., and Powers, R. D. 1959. A study of graph comprehension difficulties. *Audio Visual Communication Review* 7:97–100.

Elhadad, M. 1992. *Using Argumentation to Control Lexical Choice: A Functional Unification Implementation*. Ph.D. Dissertation, Columbia University, New York, NY.

Huff, D., and Geis, I. 1954. *How to lie with statistics*. New York, New York: W. W. Norton.

Kerpedjiev, S.; Carenini, G.; Roth, S. F.; and Moore, J. D. 1997. Integrating planning and task based design for multi-media presentation. In *Proceedings of the International Conference on Intelligent User Interfaces*, 145–152. Orlando, FL: ACM.

Kosslyn, S. M. 1989. Understanding charts and graphs. *Applied Cognitive Psychology* 3:185–226.

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

Marks, J. W., and Reiter, E. 1990. Avoiding unwanted conversational implicatures in text and graphics. In *Proceedings of AAAI-90*, 450–456. Boston, MA: AAAI.

Roth, S. F.; Kolojejchick, J.; Mattis, J.; and Goldstein, J. 1994. Interactive graphic design using automatic presentation knowledge. In *Proceedings of CHI'94*. Boston, MA.

Schmid, C. F. 1983. *Statistical Graphics: Design Principles and Practices*. John Wiley and Sons.

Schutz, H. G. 1961. An evaluation of methods for presentation of graphic multiple trends–experiment III. *Human Factors* 31:108–119.

Shah, P. 1995. *Cognitive Processes in Graph Comprehension*. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA 15213.

Siskind, J. M., and McAllester, D. A. 1993. Nondeterministic Lisp as a substrate for constraint logic programming. In *Proceedings of AAAI-93*.

Travis, D. 1991. *Effective Color Displays: Theory and Practice*. New York: Academic Press.

Tufte, E. R. 1983. *The Visual Display of Quantitative Information*. Cheshire, Conn.: Graphics Press.