



Beyond Bertin: Seeing the Forest despite the Trees

Caroline Ziemkiewicz and Robert Kosara

University of North Carolina at Charlotte

Visualization is at a point in its development where its practitioners frequently find themselves grappling with big questions about its nature and purpose. These include fundamental questions about how visualization works—that is, how do people interpret visual forms as information? However, the answers to this question haven't evolved greatly since visualization's early days. The classical view is that visualization is a process of encoding numerical or categorical values as visual (or *retinal*) variables such as size, distance, or color, which the viewer then decodes to reconstruct the original information. This variable-encoding model is the simplified essence of Jacques Bertin's *Semiology of Graphics*¹ and the years of visualization theory that have built upon it, including Jock Mackinlay's work on automated presentation design² and Leland Wilkinson's *The Grammar of Graphics*.³

The variable-encoding model has indeed contributed much to visualization design and practice. Nevertheless, there's much to suggest that it fails to capture the whole process of visualization use. Visualization is meant to facilitate pattern recognition, insight, and grasping the gist of a large data set, yet the variable-encoding model doesn't sufficiently explain any of these purposes. They all arise instead from combinations of visual objects or from the viewer's impression of the visualization's overall structure.

In addition, people traditionally display deep-seated preferences for certain information layouts that a reductionist view such as the variable-encoding model can't explain. If visualization use is just a matter of decoding visual variables, then a bar chart in which zero is on top and the bars point down should be exactly as easy to read as one in which zero is at the bottom and the bars point up. But the former case is almost never used and would be confusing if it were. Such prefer-


ences might be learned conventions and in some cases culturally determined, but some are so universal that they suggest the arrangement of data isn't wholly arbitrary in practice.

We argue that visual forms and structures used in visualization carry meaning of their own and can affect how people perceive visualized information. Studies we've conducted on visual structure support this position.⁴⁻⁶ Although our evidence isn't yet comprehensive or conclusive, it does challenge many assumptions in visualization. It also leads us to conclude that to explain how a visualization's overall structure affects a viewer, we need a holistic model in addition to the variable-encoding model.

The Meaning of Visual Forms

In a recent study on how design choices affect a user's perception of visualized data, we showed participants a series of simple charts we described as representing the budgets of departments in fictional companies.⁵ We varied five types of charts according to four design dimensions. The participants rated each company on a number of semantic variables (for example, stable, structured, unified, and controlled) on the basis of the budget distribution. In reality, all the charts had the same data proportions, but we randomly varied the sections' arrangement and color to conceal this fact. Although this task is particularly artificial and not necessarily suited to the data we presented, the study didn't aim to show that structural effects can take place in a realistic task environment (as was the case in our earlier studies on metaphor conflicts^{4,6}). Rather, we wanted to identify patterns in semantic interpretations of structural elements when users were prompted to provide them.

To this end, we asked participants to describe and compare selected charts, focusing on similarities or differences between the companies they



"The different shapes seem to indicate different types of organizations—the round seems more organic and cooperative, whereas the squares demonstrate isolation, as do the similar-sized lines between each cube."

Figure 1. Our participants felt that circular charts such as pie charts and bubble charts (top) indicate wholeness and a sense of freedom, whereas rectangular shapes such as treemaps and bar charts (bottom) suggest rigidity and organization.

represented (Figures 1–3 provide representative quotes from the participants). We wanted to find patterns in responses to help explain our quantitative data. However, the responses also pointed to several intriguing patterns in how people interpreted these basic visualizations.

For example, we found strong semantic responses to the shapes in a visualization. Two types of charts were rectangular overall (bar chart and treemap); three were circular (bubble chart, pie chart, and donut chart). The participants felt that circles in the form of pie and donut charts indicate closeness and a sense of wholeness, whereas rectangles suggest a more regimented or compartmentalized structure (see Figure 1). Circles in the form of bubble charts suggest freedom and a lack of structure altogether, perhaps because viewers perceived each circle as its own "whole." We saw no clear pattern of positive or negative feelings toward either of these impressions. Some people seemed to like structure better; others liked the sense of freedom and creativity that circles suggested.

Shape also interacted with gender in unexpected ways. Men found circular charts more controlled, whereas women felt that way about rectangular ones. This is particularly interesting in light of the

traditional association of round shapes with femininity and angular shapes with masculinity, a common rule of thumb in advertising and other design areas.⁷ Perhaps the reason for this gender difference lies in the respective appeal of the two shape types to men and women, which resulted in more negative assessments of the "opposing" shapes.

Another strong cue to meaning in the charts was color arrangement. People assigned emotion to color frequently, but not consistently. Some participants found bright colors distracting, whereas others found them cheerful; some found darker colors soothing, whereas others found them grim. One participant associated specific semantic variables with certain colors, so that a chart in which orange dominated meant flexibility, whereas red meant the company was a bad place to work.

Color also affected the charts' dynamics overall. Participants cited lopsided placements of dark and light colors as making a company seem disorganized or imbalanced. A more even color distribution had a positive effect: "The blue and brown reduces the agitative effect of the reds, yellows, and orange." The combination of color and size was also a frequent topic, as when participants saw a large, brightly colored segment as dominating the entire image (see Figure 2).

Although color is a well-studied topic in visualization at a perceptual level and has received attention beyond perception (such as consideration of cultural differences in color symbolism), researchers have paid little attention to how colors interact across a visualization as a whole. For our participants, however, color arrangement powerfully affected the visualization's apparent dynamics, and the participants readily applied these dynamics to assessments of the company itself.

The participants' comments have the flavor of design critiques, like those in the writings of color theorists such as Josef Albers⁸ or Johannes Itten.⁹ Such theories point to the expressive power of color combinations and arrangements in suggesting rhythm, agitation, harmony, separation, and innumerable other semantic properties. In many cases, this is neither unexpected nor problematic for visualization. For example, in a scatterplot in which color indicates a point's category, the fact that the eye is swiftly drawn to imbalances in color distribution will likely be a good thing. In a pie chart in which the spatial distribution of colors is less meaningful, however, this might produce unintended responses.

Structure and Visual Dynamics

Although these effects of shape, color, and size

on participants' perceptions of the companies are interesting in themselves, the participants' descriptions were most often driven by combinations of these elements. Their responses showed that they had paid close attention to a chart's overall composition, particularly how the interactions of shape, color, size, and arrangement led to inferences about the dynamics at work in the visualization. These cases show how much overall visual structure can create meaning for a viewer.

Throughout our participants' comments, visual dynamics such as balance, stability, freedom of movement, and cohesiveness were all readily applied from visual structure to conceptual structure. Of course, the idea that visual structure embodies conceptual structure isn't new; it's a fundamental assumption in graphic design and the visual arts. In this context, our participants' responses are expected and in many cases reflect what researchers in other fields already know.

Many comments by our participants along these lines are reminiscent of Rudolf Arnheim's writings on the psychology of art,¹⁰ which call attention to how the apparent weight and forces acting on objects in a painting or sculpture can be interpreted as artistic expression. Understanding how such visual dynamics affect a user's perception of information is precisely the sort of thing a reductionist model of visualization can't easily do. Viewing visualization from a more holistic perspective might let us explain these findings and incorporate useful knowledge from those domains that deal in visual dynamics more regularly.

The need for this understanding is especially apparent in several cases in which these visual dynamics prompted extensive storytelling from our participants (see Figure 3). Whereas most participants tended to jump from visual descriptions to descriptions of the company itself (as in, "The first company's design looks like a fun place to work"), some participants went further than others. These users vividly imagined the companies' working styles and atmospheres on the basis of next to no information. Here's one example:

The company that is symbolized by the square-looking figure [treemap] is more structured, more organized, goes by the rules all of the time. Going by the rules is the most important thing in this company, and to violate them can get you in serious trouble. It is well organized in that everyone understands where the company stands, but it is not a friendly place to work. People play politics and turn on each other.



"[It] seems massive, and stacking the two large sections on the left of the form makes it appear unbalanced. The pink and orange dominate the brown and purple, making their size seem even bigger in comparison. It looks like a company that's out of kilter."

Figure 2. Participants frequently analyzed the overall composition of color as a cue to interpretations about the data, as in this treemap.

These stories weren't wholly arbitrary flights of fancy. To the contrary, they tend to come across as more specific versions of the more usual type of comment that described a company with a few general adjectives. Participants frequently rated the treemap as highly structured and organized; the participant we just quoted merely fleshed out that assessment's details.

Our participants' ability to derive meaning and even stories from simple arrangements of shapes would probably not come as a surprise to most artists, who rely on such effects to express meaning through composition. In visualization, however, we haven't seriously considered how the meaning of visual forms and their composition might affect a user's interpretation of visualized information. Although we found that people derive fairly consistent meaning from shapes and visual structure, this doesn't necessarily show that this meaning has any effect on the lower-level process of decoding visual variables. Yet, it seems more reckless to assume that no such interference exists than to entertain the possibility that it does. If we admit to that possibility, what are the implications for visualization in practice?

Implications

The idea that visual structure carries information can be initially alarming to the visualization researcher. On one hand, the ability to turn visual

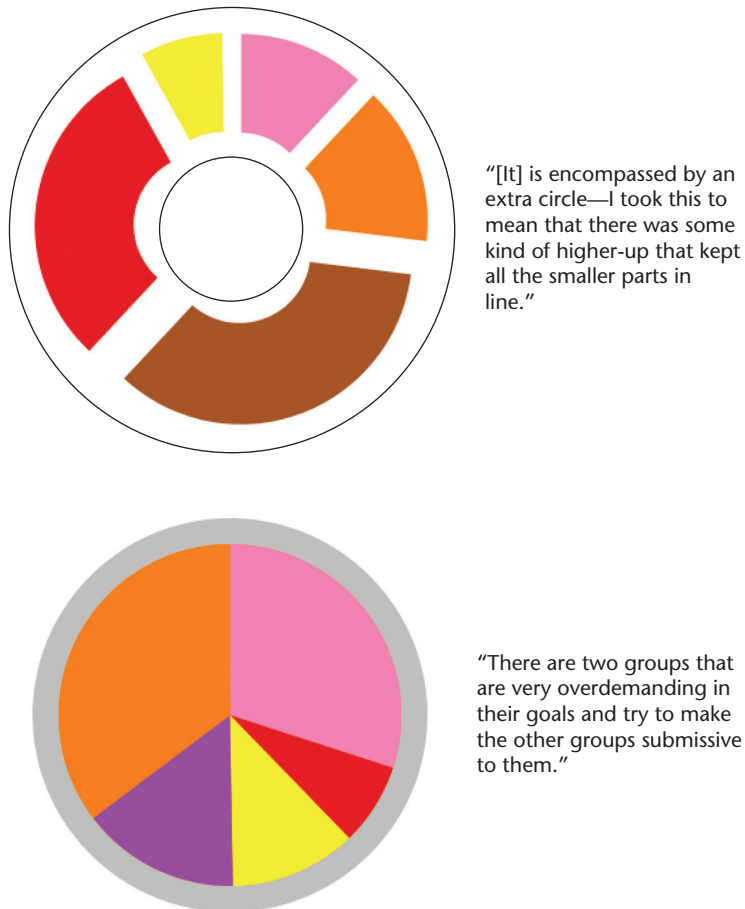


Figure 3. Several participants used their assessments of visual structure (here, a donut chart and pie chart) as the basis for imaginative stories about a company's corporate atmosphere and behavior.

forms into a meaningful story about data is one of visualization's goals. On the other hand, the fact that these stories can arise purely from differences in shape and arrangement, rather than from real differences in data values, seems highly problematic. However, the real problem isn't that these effects happen but that we lack a model of visualization that can satisfyingly explain them. Poorly designed color scales cause unintended effects all the time, but because we have perceptual models of visualization that provide hints on how to use color intelligently, the problem is manageable.

Even without a clear model for using shapes and visual structure, the basic understanding of visualization would benefit from the acknowledgment that these things have power. Taking into account the semantic effects of shape and composition could make it easier to interpret evaluation studies, especially those that compare informationally equivalent methods. Simply paying more attention to the effect of minor design elements, structuring, and shape choices during design would probably improve visualization research and practice on its own.

Design's influence on users' interpretation of data also raises ethical questions about visualization use. The visualization community is quick to call out obviously misleading uses of charts, such as comparisons across mismatched scales. However, our results raise the possibility of more subtle biases introduced by structural elements, arrangement, and color interactions. Design choices used to influence moods or semantic responses could also be used to suggest a certain reading of data—indeed, this is probably already at work in advertising. A better understanding of these structural effects would make it possible for assessments of visualization honesty to take such manipulations into account.

Although acknowledging the power of visual structure calls into question many of our usual assumptions, it would be a mistake to see this only as a problem. Rather, it should be an opportunity to broaden our views and inspire new ways of looking at visualization. If we ignore visual structure, we'll never be able to explain what's going on in visualization use. If we embrace it, we can considerably expand not only our field's capabilities but also its breadth.

Toward a Structural Theory of Visualization

To advance the theory of visualization toward a more holistic perspective, we need a structural model. This requires addressing a number of open questions. The first step is to go beyond the interesting but disorganized results from our participants to a clear set of inferred meanings and interactions people derive from visual structure. Design and aesthetics give us a basis from which to start making predictions about these associations, because descriptions our participants made were often similar to guidelines common in those fields. We should begin determining empirically what associations people make with colors, shapes, and arrangements, and particularly how those associations interact in the context of a visualization as a whole. Because so many of our responses are inspired by a chart's apparent dynamics, it would also be valuable to further study how these dynamics work in the context of animated or interactive visualizations.

Additionally, we need further testing to determine these associations' effects on data perception. A common source of other visualization researchers' skepticism toward our findings is that, although these structural effects might prompt various semantic responses when users are pressed to give them, they don't necessarily have anything to do with a user's reading of data

during an actual task. However, we suspect that they do, because research shows that the presentation of data can significantly affect reasoning and decision-making. In one case, physicians who were shown clinical study data using different types of visual representations that also stressed either the positive or negative made different decisions on whether to continue or abort the studies.¹¹ In a study by Jeff Zacks and Barbara Tversky, the presentation of the same data as either a bar chart or line graph led to predictably different inferences by participants, even when those inferences made no sense in terms of data labels.¹² Nonetheless, this is indeed a question our current research doesn't answer and is well worth studying.

Finally, simply listing structural associations would have limited utility without a model explaining them in terms of general principles. Such a model would help determine visualization design guidelines, interpret evaluation studies, and predict how novel methods will behave. On the basis of our findings, we propose that a good model for explaining these findings would be based on analyzing a scene's visual dynamics. Such analysis would explain what forces seem to be at work in a visualization and how those forces imply data patterns and relationships between objects.

More research is needed to test the viability of this or any other holistic model of visualization, but the effort must be made. The variable-encoding model tells us a great deal but leaves out too many effects that will continue to happen whether we understand them or not. Building this understanding at the structural level will expand not only our knowledge of visualization but also our control of its possibilities. ❑

References

1. J. Bertin, *Semiology of Graphics*, Univ. of Wisconsin Press, 1967.
2. J. Mackinlay, "Automating the Design of Graphical Presentations of Relational Information," *ACM Trans. Graphics*, vol. 5, no. 2, 1986, pp. 110-141.
3. L. Wilkinson, *The Grammar of Graphics*, Springer, 2004.
4. C. Ziemkiewicz and R. Kosara, "The Shaping of Information by Visual Metaphors," *IEEE Trans. Visualization and Computer Graphics*, vol. 14, no. 6, 2008, pp. 1269-1276.
5. C. Ziemkiewicz and R. Kosara, "Implied Dynamics in Information Visualization," *Proc. 10th Int'l Working Conf. Advanced Visual Interfaces (AVI 10)*, ACM Press, 2010, pp. 215-222.
6. C. Ziemkiewicz and R. Kosara, "Preconceptions and Individual Differences in Understanding Visual Metaphors," *Computer Graphics Forum*, vol. 28, no. 3, 2009, pp. 911-918.
7. B. Schmitt and A. Simonson, *Marketing Aesthetics: The Strategic Management of Brands, Identity, and Image*, Simon and Schuster, 1997.
8. J. Albers, *Interaction of Color*, Yale Univ. Press, 1963.
9. J. Itten, *The Elements of Color*, Von Nostrand Reinhold, 1970.
10. R. Arnheim, *The Power of the Center: A Study of Composition in the Visual Arts*, Univ. of California Press, 1982.
11. L. Elting et al., "Influence of Data Display Formats on Physician Investigators' Decisions to Stop Clinical Trials: Prospective Trial with Repeated Measures," *British Medical J.*, 5 June 1999, pp. 1527-1531.
12. J. Zacks and B. Tversky, "Bars and Lines: A Study of Graphic Communication," *Memory & Cognition*, vol. 27, no. 6, 1999, pp. 1073-1079.

Caroline Ziemkiewicz is a postdoctoral researcher at Brown University; she was a doctoral student in the Charlotte Visualization Center at the University of North Carolina at Charlotte. Contact her at cziemki@gmail.com.

Robert Kosara is an assistant professor in the Department of Computer Science at the University of North Carolina at Charlotte. Contact him at rkosara@uncc.edu.

Contact department editor Theresa-Marie Rhyne at theresamarieryne@gmail.com.

Call for Articles

IEEE Pervasive Computing

seeks accessible, useful papers on the latest peer-reviewed developments in pervasive, mobile, and ubiquitous computing. Topics include hardware technology, software infrastructure, real-world sensing and interaction, human-computer interaction, and systems considerations, including deployment, scalability, security, and privacy.

Further details:
pervasive@computer.org
www.computer.org/pervasive

Author guidelines:
www.computer.org/mc/pervasive/author.htm

IEEE pervasive computing
 MOBILE AND UBIGUITOUS SYSTEMS