Designing Graphs for Decision-Makers

Jeffrey M. Zacks¹ & Steven L. Franconeri²

¹Washington University in Saint Louis

²Northwestern University

Jeffrey M. Zacks

Department of Psychological & Brain Sciences

Washington University in Saint Louis

1 Brookings Drive

Saint Louis, MO 63130

314-8935-8454

jzacks@wustl.edu

Abstract

Data graphics can be a powerful aid to decision-making—if they are designed to mesh well with human vision and understanding. Perceiving data values can be more precise for some graphical types, such as a scatterplot, and less precise for others, such as a heatmap. The eye can extract some types of statistics from large arrays in an eyeblink, as quickly as recognizing an object or face. But perceiving some patterns in visualized numbers—particularly comparisons within a dataset—is slow and effortful, unfolding over a series of operations that are guided by attention and previous experience. Effective data graphics map important messages onto visual patterns that are easily extracted, likely to be attended, and as consistent as possible with the audience’s previous experience. User-centered design methods, which rely on iteration and experimentation to improve a design, are critical tools for creating effective data visualizations.

Keywords: graphs, vision, data visualization, user-centered design

Bulleted Highlights

- Human vision is astonishingly powerful, able to extract complex statistical information quickly from large arrays.
- Data graphics can capitalize on this computing horsepower to facilitate effective decision-making and communication.
- However, some kinds of visual processing—particularly comparisons—are slow and difficult. Best to minimize these.
- When people read graphs, they bring to bear knowledge about the world, along with specific knowledge about graph formats. Graphics work best when they respect this knowledge.
- User-centered design offers efficient and effective tools to improve communication through data graphics.
- By leveraging the science of data visualization, policy-makers can improve communication and decision-making.

Tweet

Your eyes are a powerful tool for extracting patterns in data. The science of visualization uncovers how to optimize communication and decision-making using visual displays.

Designing Graphs for Decision-Makers

On January 13, 2009, the Securities and Exchange Commission issued new rules for how mutual funds are described to consumers in a “Summary Prospectus.” Comparing two investments is a demanding cognitive task, and the rules were intended to make this easier. It appears that considerable thought was devoted to deciding what information to include, and in what order. However, it does not appear that much thought was given to how consumers might process this information. Most of the data are presented in tabular format; the sole exception is that a bar chart of annual returns over the previous 7 years is mandated—and the specifications for those charts are poorly designed and may not depict the most helpful information (Figure 1). When researchers assessed how people used this chart, they found that viewers focused on unhelpful information and relied on irrelevant data, leading to poor decision-making (Hüsser & Wirth, 2014).

Figure 1. Examples of SEC-mandated summary prospectus performance graphs for Vanguard S&P 500 fund and BlackRock iShares Core S&P 500 ETF (bottom). The two funds had very similar performance, which was expected because they are both broad index funds. However, the two figures look quite different. More important, neither allows consumers to easily figure out what an investment of, say, \$10,000 made in 2009 would be worth today. Retrieved from the Vanguard and BlackRock corporate web sites, August 17, 2019.

In Figure 2, an important analysis of the economic impact of proposed climate change legislation is extremely difficult to understand. The graph is cluttered; it requires constant looks back and forth to a color-lookup legend; and the bullets at right require sluggish integration of high-level conclusions with patterns among the data values.

The display shown in Figure 3 is used by the National Hurricane Center to visualize forecasted paths of a storm. The blue area shows a counter-intuitive collection of forecasted paths, such that a majority of weather models predict that the storm will travel through that region. That region grows larger as the storm approaches, because of the increasing uncertainty in the forecasts. But many viewers mistakenly interpret that the storm that is growing in size as it approaches (Padilla, Creem-Regehr, Hegarty, & Stefanucci, 2018).

Is this the best we can do? In many situations, people make better decisions from graphical presentations than from verbal or tabular formats. But these visualizations must first be designed in a way

that taps the power—and respects the limitations—of visual processing. When data visualizations are constructed well, analysis is powerful and communication is clear. But when designed ineffectively, visualizations leave critical patterns opaque or leave viewers confused about how to navigate cluttered or unfamiliar displays. The fact that data visualizations can have massive effects on a viewer's ability to analyze and understand patterns means that it is crucial for decision-makers and communicators to understand the fundamentals of data design.

Figure 3: A graph that violates conventional expectations about the meaning of shaded regions. Viewers interpret the size of the region as the anticipated size of the hurricane, whereas it is meant to indicate the degree of uncertainty of the path forecast (Padilla et al., 2018; reproduced per CC-BY license, <https://creativecommons.org/licenses/by/4.0/>).

Who studies the visual communication of quantitative information?

Several intersecting fields study how people translate visual displays into understanding and help to develop practical guidelines for making that process more effective. In the social sciences, psychologists study how people extract data from visual images (e.g., Pinker, 1990; Shah & Carpenter, 1995), how this is affected by knowledge and expertise (e.g., Parsons, 2018; Smallman & Hegarty, 2007; Tversky, 2001). Education researchers study how to improve how students learn how to read and reason with graphs (e.g., Friel & Bright, 1996; Shah & Hoeffner, 2002). In communications, political science, and public policy, work on the persuasive power of visualized data, for example, explores whether visually depicted information is more likely to persuade a climate change skeptic, compared to quoting numbers and statistics (Nyhan & Reifler, 2019). Health communication researchers study how to best communicate risk to patients who need to understand probabilities of success across treatment options

(e.g., Ancker, Senathirajah, Kukafka, & Starren, 2006). Last but not least, one academic field, primarily made up of computer scientists and designers, specifically studies information visualization. This work ranges from developing tools that solve particular challenges for scientific collaborators or the general public, psychological studies of visualization perception (Cleveland & McGill, 1985), understanding (Padilla et al., 2018), intuitive interactivity (Elmqvist et al., 2011), and engagement (Naps et al., 2002). The information visualization community relies increasingly on empirical evaluation—both through qualitative and quantitative data—to validate guidelines for effective visualization design (Kosara, 2016).

Many practitioners also create, illustrate, and curate guidelines for effective presentation of data. In the past two decades, dozens of books have been published, with increasing frequency, on effective data visualization (see the References at the end of this article). While many of these are oriented toward the business community, their lessons are equally relevant across science, education, and public policy communication. Also, a rapidly growing community of journalists specialize in communicating complex stories that rely on quantitative data about public policy, demographic trends, or election forecasts (for example, www.nytimes.com/section/upshot, fivethirtyeight.com, www.washingtonpost.com/news/wonk/wp/category/data-visualization).

Graphs in Human Vision

The Power of Vision

Visualizations can leverage the massive processing power of the human visual system (around half of the human brain; Van Essen, Anderson, & Felleman, 1992), allowing rapid foraging through patterns in data and intuitive communication of those patterns to the human eye. In the table of numbers in the left of Figure 4, pattern-seeking requires processing text-based numeric symbols, presenting a processing bottleneck that slows you down. Simple tasks, like finding the highest or lowest numbers, or uncovering a broader spatial pattern across the high numbers, are tedious. But in the color-coded version next to it, these tasks are trivial and immediate. That is because the visual system can analyze patterns of

a set of basic features—such as position, length, area, or pixel intensity—in a single parallel analysis across and entire visual display.

Figure 4: Leveraging the visual system to detect patterns.

Some research focuses on how to best tap this visual processing power. One bedrock finding is that visual features vary in how precisely they convey data values to the human eye. These studies ask viewers to guess the ratio between two values, manipulating the features that code those values, and measure the average error in those estimates (Cleveland & McGill, 1985; Heer & Bostock, 2010). Figure 5 organizes a subset of the typically tested features from most precise (position) at the top, to the least (intensity) at the bottom; each depicts a 1:7 ratio between the two values. The high precision of position information is one driver of the ubiquity of position-based visual representations, such as dot plots, scatterplots, bar graphs, and line graphs. Although area and intensity (e.g., greyscale value, or color saturation) are far less precise, they are still immensely useful for depicting ‘big picture’ information, as in the color-coded number grid above (Albers, Correll, Gleicher, & Franconeri, 2014).

These features do not only convey individual values or simple comparisons between pairs of values. The human visual system also can analyze statistics across large collections of these values. In the natural world, compiling such statistics on distributions of visual features can be extremely useful. If there is a lot of green and brown, you are likely in a park; if there is a lot of light brown and blue, and a salient horizontal horizon line, you are likely to be on a beach. A plateful of cookies makes it easy to pick out which has larger area, and has the most chips. Figure 6 (reproduced from Szafir, Haroz, Gleicher, & Franconeri, 2016) shows four types of patterns that might be pulled from large collections, across four types of visual features.

The Limits of Vision

Visual processing evolved and developed to interact with the natural world, not to extract numbers and statistics from artificial displays. This mismatch can cause visual illusions that bias our view of data (Szaifir et al., 2016). The map at left in Figure 7 simulates a problem that arises when using intensity to plot two variables at once, in the same spatial area. Although the two circles are the same intensity, the bottom one appears darker because of its lighter background. This illusion likely stems from the visual system's reliance on background contrast to judge the brightness of an object. In the natural world, this helps calibrate brightness judgments to discount ambient illumination (Purves, Lotto, & Nundy, 2002), but in artificial data displays this causes systematic bias (Ware, 2013).

In the line graph at right in Figure 7, the two lines form exactly the same shape, but one is vertically offset from the other such that the three red dotted lines are exactly the same length. Yet, when calculating the difference between the lines by eye, that difference appears far larger for the region on the left relative to the region on the right. This illusion likely stems from the visual system's focus on the

length and width of objects in the world, seeing the space between the lines as an object that is ‘thick’ at bottom left, and ‘thin’ on the top right—effectively measuring the shortest distance between the curves rather than the distance measured vertically.

Figure 7: Visual illusions can distort graphical perception. Left: When identical gray circles are superimposed on a light or dark region, they appear dark or light due to luminance contrast. In this case, viewers would be misled into seeing the United States circle as darker than the Canada Circle. The simplified version at center plots the same circle twice, on a continuously changing background. Right: Parallel curves often do not appear parallel due to perception of the space between them as an object.

Finally, seeing in the natural world typically does not require a substantial amount of short-term memory. If you need to know the color or shape of an object, you can simply look at it, which might create an illusion that the information was already stored in your head (O'Regan, 1992; Rensink, 2000). This works when the information is close to where you are already looking—it's the reason that your car's GPS map display is on the dashboard in front of your eyes. If it were on the ceiling, your limited memory would force you to make more effortful glances back and forth between the map and the view outside. In the natural environment, related things are usually nearby—tomatoes hang on the vines that grow them, and ducklings follow a parent. This is not *always* true, however; you might see a tomato on the ground and look for the plant from which it fell. Good visualizations minimize the degree to which viewers need to glance around between different parts of a visualization. Figure 8 shows a visualization that violates this principle; such legends cause viewers unnecessary pain.

Comparison Limitations

The visual system excels at processing many data values at once when those values encode basic features such as position, area, or color saturation. As described previously, the visual system can even pull statistics from this analysis, such as the average or maximum value from a collection. For example, Figure 9 allows immediately locating the longest bar among the 20 total bars. The ease of extracting those statistics gives us confidence that we see the data in a powerful way—and for many operations we do. But for one critically important data analysis operation, our visual system slows to a crawl.

That operation is comparison, and in many cases, people can only do one at a time (Franconeri, Scimeca, Roth, Helseth, & Kahn, 2012). Comparing one bar's length to another (Is the left or the right bar longer?) takes a split second to process. That sounds fast, but doing dozens or hundreds of those comparison steps will take anywhere from several seconds to several minutes (Logan, 1994; Wolfe, 1998), including in data visualizations (Nothelfer & Franconeri, 2019). In Figure 9, notice how it takes a few seconds to find the two pairs of bars with a unique arrangement (tall on the left, short on the right). Imagine making all of the possible comparisons from a simple bar graph with 7 bars in it. How many pairwise comparisons (bar 1 to 2, 2 to 3, 1 to 3...) are possible? Even this small dataset allows 21 unique pairwise comparisons; with 20 bars the number is 190. And that's only pairwise comparisons—most data

graphics have a large number of potential comparisons, and the viewer doesn't always know which ones to prioritize.

Figure 9: In the top panel, which bar is highest? That task is easy. Now, which pair is decreasing? That task is harder. The comparison task can be made simple by turning the differences into visual objects, as in the bottom panel.

This means that seeing a graph is not like the immediate and effortless recognition of a face, car, or Pokemon. The term “reading” is more apt than seeing, because reading a graph is more like reading a paragraph. You construct an understanding based on the structure of the graph, your previous knowledge, and the questions you are trying to answer. Good graph design can help.

Graphs in Human Understanding

Visualization designers help viewers navigate the large number of potential visual comparisons afforded by a graph. One strategy depends on general-purpose mechanisms of visual attention to guide viewers; another strategy leverages previous graph experience by following a conventional format, so that viewers can use learned associations to extract the relevant comparisons. Most good designs do both of these. Sometimes there is not an obvious way to guide attention effectively while sticking to convention; these situations often pose the toughest design challenges.

Guiding Attention for Effective Comprehension

One important aspect of visual processing for guiding graph comprehension is that vision forms *groups* or collections of objects. Patterns are automatically clustered based on proximity, shared color or shape, shared orientation, falling on a smooth line, or forming a familiar pattern such as a square (Brooks, 2015). This can enable powerful visual comparisons; for example, we can effortlessly compare two large groups of plot symbols based on their having different colors or being close in space. In the top panel of Figure 10, comparing the populations of different regions in Africa is difficult because the text labels and arrangement do not facilitate visual grouping. The middle panel uses color to group them, which helps somewhat. The bottom panel further strengthens the visual grouping, by using both spatial proximity and color to guide attention.

Visual grouping can effectively guide viewers' attention, facilitating comprehension. However, it also can lead to distortions; for example, if one data point from a group falls far from the others in the group, it may be missed, mis-categorized, or never compared to the rest of its group. This can present a challenge to the designer, particularly when the data include outliers.

Another aspect of vision for guiding graph comprehension is that visually *salient* objects attract attention (Itti, Koch, & Niebur, 1998). Saliency is the degree to which an object or region of space stands out from others. Salient things differ from their neighbors--for example, a vertical tree in a field of felled horizontal ones, a square candy in a pile of oval ones, or a red poppy in field of green. For example, the bar for Nigeria is salient in all three panels of Figure 10 because it is the biggest. In the middle and bottom panels, it is even more salient because it is a blue object surrounded by yellow and orange objects. Designers can use visual saliency to guide viewers to the comparisons of interest, by highlighting them. Highlighting can also create new visual objects that make a comparison of interest directly visible, as illustrated in Figure 11.

Figure 10 Population of the ten largest countries in Africa. In the top panel it is nearly impossible to compare regions. The second panel allows grouping by color, but the fact that countries from the different regions are not contiguous still makes the groups hard to see. The bottom panel uses color and spatial proximity to facilitate grouping by region.

A final, crucial, way to guide attention is by using language to tell viewers what to see. This is also illustrated in Figure 11. Annotating figure elements can be a great help in guiding viewers to the visual comparisons of interest while preserving their ability to verify the designer's claims for themselves by inspecting the data. However, two caveats are worth bearing in mind. First, locate the verbal information near the visual information of interest to avoid forcing viewers to look back and forth

between the words and the visual information (Moreno & Mayer, 1999). Second, only a very limited number of things can be called out with language without overwhelming the viewer.

Figure 11: Adding visually salient highlighting and annotation effectively guides comprehension.

Graph Schemas

Even if a viewer has the visual machinery available to process patterns in data, they need to have learned how to use it to extract and understand values from a particular graph design. Take the bar graph in Figure 12. You immediately recognize that those two bars depict the height of two separate groups of people. The grey rectangles matter, not the surrounding box. Their distance from each other, and the brightness of each bar, is irrelevant. Each rectangle presents some summary (average) statistic that stands for a whole group. For each rectangle, the height matters but not the width; the top matters but not the bottom. The numbers along the vertical axis represent values.

Some of this knowledge has been adapted by designers from our general knowledge about objects in the world. When bricks or logs or cans are stacked up, stacks of more are generally taller, so graphs usually use “up” rather than “down” to represent more. Some graph knowledge is broadly shared by people with formal education and can largely be assumed. For example, at least one spatial dimension is usually quantitative (the other may be categorical); the size of each wedge in a pie graph represents a percentage of a whole; error bars represent variability; time usually runs left to right. Such knowledge is

sometimes taught, but also can be picked up through exposure. It can exert surprisingly strong influence on the messages that we take from graphs, without our necessarily being aware of it.

In one study (Zacks & Tversky, 1999), viewers were asked to describe simple datasets that were depicted with a line graph or a bar graph. Their implicit knowledge told them that bar graphs are generally designed to highlight comparisons between data points, whereas line graphs are generally designed to highlight trends across multiple data points. This affected the descriptions they gave: When presented with a bar graph as in the left panel of Figure 12, they were likely to say things like “12-year-olds are taller than 10-year-olds,” rather than “Height increases with age.” More surprisingly, viewers presented with a line graph as in the right panel of Figure 12 sometimes said things like “The more male you are the taller you are.” (In this study, 15% of respondents did so.)

Figure 12: Two graphs that mismatch common expectations for bar and line graphs.

To understand a graph, a viewer needs to know how to map a visual feature to a real-world referent. To understand the right panel of Figure 12, viewers need to know how to map people’s height to the height of the points. In addition, they need to know how to guide attention over the graph to extract relevant information. To read a bar graph such as the one on the left in Figure 12, viewers need to know how to guide attention to the y axis to identify the real-world variable being plotted, and to guide attention to the x axis to identify what function is being plotted. If attention goes to those locations before going to the points themselves, interpretation will be more accurate and more efficient. Viewers learn these aspects of how to use graphs through repeated experience, and sometimes through instruction.

Such organized prior knowledge goes by the label *schema* (Pinker, 1990). When you have a schema for a graph type, mapping visual features to real-world referents is usually easy. Mapping height in a bar graph to the value of the variable named on the y axis, or mapping redness onto heat in a weather map, is easy because we have extensive experience with these graph types. Graphs that are new to us are tougher, but some leverage schemas we have for other domains.

In short, *seeing* the relevant features in a graph is not sufficient to *understand* it. Understanding is a process that unfolds over multiple views, guided by attention and by knowledge—including knowledge about particular graph formats. Good designs effectively guide attention and respect viewers’ knowledge.

Designing for Vision and Understanding

Effective graphs cater to humans’ visual capacities and to their conceptual understanding. Based on the theory and data reviewed in the previous sections, we can offer the following principles for effective graph design.

Table 1: Some principles of effective graph design
1. When precise quantitative judgments are needed, use position or length and avoid area or intensity.
2. Organize patterns in data onto patterns in visual features so that visual grouping by proximity, shape, size, and intensity reflects true grouping in the data.
3. Avoid creating visual illusions that distort data.
4. Minimize visual comparisons.
5. Minimize working memory demands.
6. Respect common conventions when you can.
7. Other things being equal, avoid obscure graph formats.
8. When you need to break with convention, try to respect broader knowledge about how values map onto position, length, and area.
9. Experiment, gather feedback, and iterate.

Principle number 9 in Table 1 merits special comment, because it introduces a new aspect of the design process. The first eight principles summarize general patterns in graph design. They are well supported by data, and they make theoretical sense. However, every dataset is different, and so is every situation that may need to communicate a message about a dataset. Often, two principles may conflict, and too many potential conflicts prevent sort them all out in advance. Because design principles have so many potential interactions, good designs usually require *iteration* and *experimentation*, paired with

feedback from people representative of the eventual audience (Kosara, Drury, Holmquist, & Laidlaw, 2008). Your own eyes and visual brain can tell you a lot about whether a graph “works.” So try a few variants and look carefully at them.

However, you cannot fully trust your own judgment, because as the designer, you are afflicted by the *curse of knowledge*. You know what message a viewer should extract from your visualization, and once you do, it is hard to simulate the viewpoint of someone who does not (Xiong, van Weelden, & Franconeri, 2019). It is impossible to turn off your own expertise, which makes it difficult to see through the eyes of non-experts.

A quick series of experiments can go a long way toward refining a design. The experiments do not have to be elaborate or formal—often, showing two or three potential designs to a few people each and asking them to answer the questions you would like a real viewer to be able to answer easily will do the trick. These features of the design process have been codified as *user-centered design* (Abras, Maloney-Krichmar, & Preece, 2004), and they can be very effective.

This is an exciting time to be thinking about data graphics. The emerging field of data visualization, weaving together empirical disciplines including psychology with technological and esthetic disciplines, is moving fast and is impacting our media and our communications with each other. This context offers great opportunity for policy-makers to facilitate better decision-making with more effective visual communication. Use what we know about visual processing to design to viewers’ strengths. Take advantage of their knowledge about graphs and about the world in general. And don’t hesitate to experiment with your designs.

References

- Abras, C., Maloney-Krichmar, D., & Preece, J. (2004). User-centered design. In W. Bainbridge (Ed.), *BainEncyclopedia of Human-Computer Interaction. Thousand Oaks: Sage Publications* (Vol. 37, pp. 445–456).
- Albers, D., Correll, M., Gleicher, M., & Franconeri, S. L. (2014). Ensemble processing of color and shape: Beyond mean judgments. *Journal of Vision, 14*(10), 1056–1056.
<https://doi.org/10.1167/14.10.1056>
- Ancker, J. S., Senathirajah, Y., Kukafka, R., & Starren, J. B. (2006). Design Features of Graphs in Health Risk Communication: A Systematic Review. *Journal of the American Medical Informatics Association, 13*(6), 608–618. <https://doi.org/10.1197/jamia.M2115>
- Brooks, J. L. (2015). *Traditional and new principles of perceptual grouping*.
- Cleveland, W. S., & McGill, R. (1985). Graphical Perception and Graphical Methods for Analyzing Scientific Data. *Science, 229*(4716), 828–833. <https://doi.org/10.1126/science.229.4716.828>
- Elmqvist, N., Moere, A. V., Jetter, H.-C., Cernea, D., Reiterer, H., & Jankun-Kelly, T. J. (2011). Fluid interaction for information visualization. *Information Visualization, 10*(4), 327–340.
- Franconeri, S. L., Scimeca, J. M., Roth, J. C., Helseth, S. A., & Kahn, L. E. (2012). Flexible visual processing of spatial relationships. *Cognition, 122*(2), 210–227.
<https://doi.org/10.1016/j.cognition.2011.11.002>
- Friel, S. N., & Bright, G. W. (1996). *Building a Theory of Graphicacy: How Do Students Read Graphs?*.
- Heer, J., & Bostock, M. (2010). Crowdsourcing graphical perception: Using mechanical turk to assess visualization design. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 203–212*. ACM.
- Hüsser, A., & Wirth, W. (2014). Do investors show an attentional bias toward past performance? An eye-tracking experiment on visual attention to mutual fund disclosures in simplified fund

- prospectuses. *Journal of Financial Services Marketing*, 19(3), 169–185.
<https://doi.org/10.1057/fsm.2014.20>
- Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (11), 1254–1259.
- Kosara, R. (2016). An empire built on sand: Reexamining what we think we know about visualization. *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization*, 162–168. <https://doi.org/10.1145/2993901.2993909>
- Kosara, R., Drury, F., Holmquist, L. E., & Laidlaw, D. H. (2008). Visualization Criticism. *IEEE Computer Graphics and Applications*, 28(3), 13–15. <https://doi.org/10.1109/MCG.2008.63>
- Logan, G. D. (1994). Spatial attention and the apprehension of spatial relations. *Journal of Experimental Psychology: Human Perception and Performance*, 20(5), 1015–1036.
<https://doi.org/10.1037/0096-1523.20.5.1015>
- Moreno, R., & Mayer, R. E. (1999). Cognitive principles of multimedia learning: The role of modality and contiguity. *Journal of Educational Psychology*, 91(2), 358–368.
- Naps, T. L., Rößling, G., Almstrum, V., Dann, W., Fleischer, R., Hundhausen, C., ... Velázquez-Iturbide, J. Á. (2002). Exploring the Role of Visualization and Engagement in Computer Science Education. *Working Group Reports from ITiCSE on Innovation and Technology in Computer Science Education*, 131–152. <https://doi.org/10.1145/782941.782998>
- Nothelfer, C., & Franconeri, S. L. (2019). Measures of the benefit of direct encoding of data deltas for data pair relation perception. *Proceedings of IEEE VIS*. Presented at the IEEE VIS, Vancouver, BC.
- Nyhan, B., & Reifler, J. (2019). The roles of information deficits and identity threat in the prevalence of misperceptions. *Journal of Elections, Public Opinion and Parties*, 29(2), 222–244.
- O'Regan, J. K. (1992). Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, 46(3), 461.

- Padilla, L. M., Creem-Regehr, S. H., Hegarty, M., & Stefanucci, J. K. (2018). Decision making with visualizations: A cognitive framework across disciplines. *Cognitive Research: Principles and Implications*, 3(1), 29. <https://doi.org/10.1186/s41235-018-0120-9>
- Parsons, P. (2018). Conceptual metaphor theory as a foundation for communicative visualization design. *To Appear in Visualization for Communication (VisComm)*, 2018.
- Pinker, S. (1990). A theory of graph comprehension. In *Artificial Intelligence and the future of testing: Vol. Artificial Intelligence and the future of testing* (pp. 73–126). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Purves, D., Lotto, R. B., & Nundy, S. (2002). Why we see what we do: A probabilistic strategy based on past experience explains the remarkable difference between what we see and physical reality. *American Scientist*, 90(3), 236–243.
- Rensink, R. A. (2000). The dynamic representation of scenes. *Visual Cognition*, 7(1–3), 17–42.
- Shah, P., & Carpenter, P. A. (1995). Conceptual limitations in comprehending line graphs. *Journal of Experimental Psychology: General*, 124(1), 43–61.
- Shah, P., & Hoeffner, J. (2002). Review of graph comprehension research: Implications for instruction. *Educational Psychology Review*, 14(1), 47–69. <https://doi.org/10.1023/A:1013180410169>
- Smallman, H. S., & Hegarty, M. (2007). Expertise, spatial ability and intuition in the use of complex visual displays. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 51, 200–204. SAGE Publications Sage CA: Los Angeles, CA.
- Szafir, D. A., Haroz, S., Gleicher, M., & Franconeri, S. L. (2016). Four types of ensemble coding in data visualizations. *Journal of Vision*, 16(5), 11–11. <https://doi.org/10.1167/16.5.11>
- Tversky, B. (2001). Spatial schemas in depictions. In *Spatial schemas and abstract thought: Vol. Spatial schemas and abstract thought* (pp. 79–112). 1. Stanford U, Stanford, CA Cambridge, MA: The MIT Press.

Van Essen, D. C., Anderson, C. H., & Felleman, D. J. (1992). Information processing in the primate visual system: An integrated systems perspective. *Science*, 255(5043), 419–423.

<https://doi.org/10.1126/science.1734518>

Ware, C. (2013). *Information visualization: Perception for design* (Third edition). Waltham, MA: Morgan Kaufmann.

Wolfe, J. M. (1998). What can 1 million trials tell us about visual search? *Psychological Science*, 9(1), 33–39. <https://doi.org/10.1111/1467-9280.00006>

Xiong, C., van Weelden, L., & Franconeri, S. L. (2019). The curse of knowledge in visual data communication. *IEEE Transactions on Visualization and Computer Graphics*, 1–1.

<https://doi.org/10.1109/TVCG.2019.2917689>

Zacks, J. M., & Tversky, B. (1999). Bars and lines: A study of graphic communication. *Memory and Cognition*, 27(6), 1073–1079.