Benefitting InfoVis with Visual Difficulties?

Provocation Without a Cause

Stephen Few, Perceptual Edge
Visual Business Intelligence Newsletter
October/November/December 2011

A research paper was presented last year at CHI, the premier conference on human-computer interation, which caused quite a stir in the information visualization community: “Useful Chartjunk? The Effects of Visual Embellishment on Comprehension and Memorability of Charts” by Bateman, et al. It was immediately taken up as a rallying cry by advocates of graphical embellishment. Upon examination, however, the paper demonstrated nothing surprising, and much less than it claimed. You can read my review in the April/May/June 2011 *Visual Business Intelligence Newsletter* article titled “The Chartjunk Debate: A Close Examination of Recent Findings.” During this year’s VisWeek 2011 conference, a related research paper by Jessica Hullman, Eytan Adar, and Priti Shah of the University of Michigan, titled “Benefitting InfoVis with Visual Difficulties,” made the provocative claim that chartjunk can function as a form of visual difficulty that benefits the user. The authors summarized their claims in the first few lines of the paper’s abstract:

> *Many well-cited theories for visualization design state that a visual representation should be optimized for quick and immediate interpretation by a user. Distracting elements like decorative “chartjunk” or extraneous information are avoided so as not to slow comprehension. Yet several recent studies in visualization research provide evidence that non-efficient visual elements may benefit comprehension and recall on the part of users.*

We can’t ignore this claim. If the authors are right, we may need to reverse many current practices. If they’re wrong, however, which indeed they are, their claim could do great harm. An extensive body of research from several disciplines has long directed us to display data as simply as possible, eliminating gratuitous, distracting, and difficult-to-perceive visual content in an effort to reduce unnecessary and unproductive cognitive load. “Benefitting InfoVis with Visual Difficulties” challenges those guidelines based solely on an illusion of evidence and credible research, which dissolves into an ephemeral mist upon examination.

An Opportunity Missed

This paper is the result of ill-conceived and poorly-conducted research, yet it was awarded one of only two honorable mentions by the InfoVis judges. This is embarrassing. It reveals a deep-seated problem that undermines much of our work.

What happened during the course of this research is a travesty because the authors began with an important observation:

> *With very few exceptions,* [information visualization research] *evaluation measures tend to assess how a user interacts with a visualization and not how well they learn the important concepts or patterns represented. It has been noted that it is somewhat unrealistic to hold visualizations to an efficiency-based standard of minimal response times and high response accuracy as in practice these measures tend to be in a trade-off with higher accuracy requiring more time examining the representation.*

This is indeed a problem that deserves a provocative challenge. Information visualization research should test relevant outcomes, not merely measure decoding accuracy and efficiency. Unfortunately, rather than focusing on ways to engage users with visualizations in more meaningful and productive ways, resulting in greater insights, the authors argue that we should introduce visual difficulties—features that hinder perception—to promote “higher accuracy requiring more time examining the representation.” Had they focused solely on meaningful engagement with data, proposed ways to encourage it, subjected these means of engagement to rigorous tests, and then reported the results, either positive or negative, I might have found this research useful.

The Inspiration

Based on several citations to last year's paper by Bateman, et al, I suspect that the authors were inspired by this flawed and ultimately hollow endorsement of chartjunk. At most, Bateman and his colleagues demonstrated that extremely simple charts, when embellished with thematically-relevant cartoon images created by a talented graphic artist, are easier to remember weeks later. Let's bear in mind that long-term recall is rarely the purpose of information visualization. When we interact with visualizations to explore and make sense of information, insights arise and lead to new questions, which in turn lead to further insights, and so on. Long-term recall is seldom needed.

Visual Difficulties Defined

What exactly are these "visual difficulties" that we should add to visualizations for the benefit of the user? The authors refer to visual difficulties variously as "non-efficient visual elements," "distracting visual elements and irrelevant information," and "extraneous elements (e.g., chartjunk)." All of these descriptive phrases refer to visual attributes that would hinder visual processing required to perform the task at hand. Several features of the chart below—the 3-D effects of depth, angle, light, and shadow; the hard-to-read font; the use of pie slices (areas and angles) to encode values rather than an easily perceived means such as the lengths of bars—are examples of visual difficulties that this paper seems to endorse as potentially beneficial.

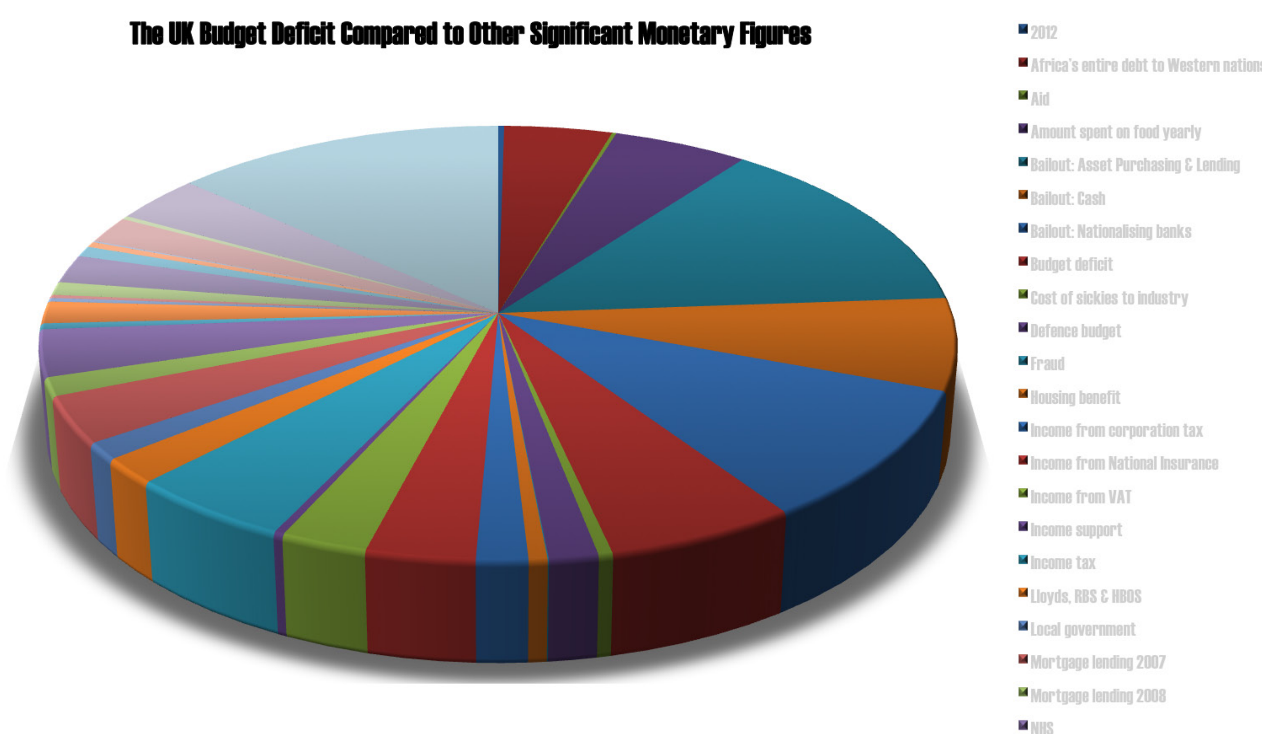

Perhaps developers of software products such as *Microsoft Excel* have had it right all along, but for the wrong reasons. According to this paper, gratuitous embellishments and hard-to-decode graphs (e.g., pie charts) are not desirable merely because they're entertaining and pretty, but because they increase cognitive load, which leads to greater comprehension.

From descriptive phrases such as "non-efficient visual elements," "distracting visual elements and irrelevant information," and "extraneous elements (e.g., chartjunk)," the authors go on to formalize their definition of "visual difficulties" in ways that stray from the expected path.

> *We conceptualize visual difficulties as the set of beneficial learning and cognitive processing obstructions that may be applied to static and interactive information visualizations. This include* [sic] *but are not limited to visual techniques that manipulate the visual representation, also including manipulations of accompanying information and task and contextual characteristics of visualization interaction.*

Notice that the definition begins to drift from the realm of chartjunk to a broader realm of "beneficial learning and cognitive processing difficulties," which is not necessarily visual in nature. In fact, you'll see as we proceed that none of the research examples of empirically-demonstrated learning and cognitive obstructions are actually visual in nature, with one exception, and in none of the cited studies were cognitive obstructions applied to data visualization. Techniques that "manipulate the visual representation" (filtering, sorting, brushing, switching chart types, etc.) don't qualify as visual difficulties. We enable interaction with visualizations, not to increase cognitive load for some indirect benefit, but to reduce it by producing meaningful and insight-inducing views that could not be otherwise fathomed without great difficulty. As the paper continues, the authors increasingly focus on ways to engage people with information that are not fundamentally visual in nature, don't qualify as cognitive difficulties, don't relate to information visualization, and certainly don't resemble chartjunk.

The Evidence

Let's consider the empirical foundation on which these claims are built. The authors state: "Several recent studies in visualization research provide evidence that non-efficient visual elements may benefit comprehension." They also promise to "summarize a large body of psychological research on learning from graphs and diagrams that suggests that introducing visual difficulties to visualizations can be an effective way to stimulate important ingredients like active processing of information and engagement." Do they deliver? Let's find out.

What is this evidence from psychological and learning research "that finds that [using] more difficult to interpret graphs and other visual materials yields better learning"? The authors whet our appetite before revealing the proffered evidence: "Specifically, we note that difficult displays require mentally manipulating internal representation, forming self explanations, and appraising and noting disfluency, all of which promote *active* processing on the part of the user." Actually, we will find that the research they cite does not indicate that *difficult displays* involve any of these cognitive operations. Making visualizations of data unnecessarily difficult to read and interpret leads to unnecessary cognitive effort, without benefit. What the authors cite from psychological and learning research primarily suggests that it is *engagement* with information that leads to reflective thinking, which actually produces the positive outcomes that they mistakenly attribute to visual difficulties.

For example, they refer to the work of psychologists Elizabeth and Robert Bjork of UCLA who have researched the use of "desirable difficulties" to promote learning. What are these desirable difficulties? They include:

- "Varying the conditions of practice" (e.g., studying the same material in two different rooms rather than twice in the same room)
- "Spacing study or practice sessions" (e.g., not trying to learn by cramming for the test)
- "Interleaving versus blocking instruction on separate to-be-learned tasks" (i.e., to produce longer-term retention of information, use random and interleaved practice sessions, rather than sessions that block content into rigid topical chunks
- "Generation effects and using tests (rather than presentations) as learning events" (i.e., involve students in activities that force them to think about the material rather than simply giving them the answers).

These tactics for engaging students in thinking reflectively about information are not in the least related to the visual difficulties that the authors recommend. They have appropriated irrelevant research to support their claim. In an imaginative leap, when they begin to present the so-called evidence from psychology and education research, they revise their definition of visual difficulties to broadly refer to the act of "stimulating more intense cognitive activity with a visualization." However, none of the evidence that they present, with only one exception, pertains to visual difficulties.

For instance, the authors cite self-explanation—an activity that forces someone to think the information through more thoroughly—as an example of a visual difficulty. But requiring someone to fill in that portion of a bar that they believe represents some value's proportion of the whole (e.g., sales in the east region as a proportion of total sales), which is what subjects in this study were required to do, is not the same as adding 3-D effects to a bar graph. Self-explanation engages the user in thinking about the data, but visual difficulties merely force the user to work around and attempt to ignore something like a third dimension of depth while comparing the heights of the bars.

They further cite novelty, tailoring and personalization, challenge and game-play—characteristics that have been shown in particular contexts to promote engagement—as evidence for the benefits of visual difficulties. None of these means of engagement necessarily involve what can legitimately be called visual difficulties. It's useful to introduce a novel form of visualization when it will invite the user to think about the data in a new way that's meaningful, not arbitrarily for the sake of novelty. The ability to tailor a visualization to display data in ways that make the story more accessible to your eyes works because it provides what's needed, not because it introduces visual difficulties. The authors provide no examples of challenge and game play that were introduced by means of visual difficulties. People indeed love to be challenged, but not in ways that waste their time by presenting obstacles that aren't useful. When game play is incorporated into interaction with a visualization in a way that draws people into thinking about the data, it can yield benefits. It's hard to imagine any useful form of game play, however, that merely involves visual difficulties.

Only one psychology or education research study that they cite actually relates to visual difficulties: "Overcoming Intuition: Metacognitive difficulty activates analytic reasoning," by A.L. Alter, D.M. Oppenheimer, N. Epley, and R.N. Eyre, published in the *Journal of Experiential Psychology-General*, vol. 136, no. 4, 2007. This study was built on the premise that people use two modes of thinking: System 1 processing (aka, automatic), such as intuition and heuristics, and System 2 processing (aka, reflective), which relies on slower, more careful, analytical thinking. Subjects were asked to read written content and then make judgments about it. Half of the subjects read content that was printed in a difficult-to-read typeface (*light gray, small, and italicized*) rather than one that was easy to read. The questions were designed to give subjects the false impression that they could be easily answered using an intuitive System 1 response alone. The purpose of visual difficulty was to force subjects to use System 2 processing instead, which was required to respond correctly. They found that the difficult-to-read typeface often seemed to prompt this behavior, producing more accurate answers.

When people use visualizations to explore and make sense of data for the tasks that require thinking, they usually adopt a System 2 processing strategy, at least in part, so the need to trigger a shift from System 1 processing rarely occurs. There are certainly situations when a visualization might appear deceivingly simple when in fact a closer examination is required to interpret it accurately, but no one would seriously argue that we should therefore make all visualizations difficult to read to make sure that System 2 processing is always triggered. Skilled data analysts learn to view data from many perspectives to prevent knee-jerk conclusions based on inappropriate heuristics. Even if there were reason to believe that information disfluency in some form might be useful for triggering System 2 processing when viewing a visualization, we shouldn't assume that features such as gratuitous 3-D effects or poor color choices would trigger this response. Ideally, when a catalyst is needed to initiate System 2 processing, a means that produces negative effects (e.g., the extra time required to process a visually difficult visualization), should be avoided by using a catalyst that has no negative effects. Although I wouldn't seriously suggest it, you could instruct users to frown while viewing the visualization. Just as the degree of concentration that goes with diligent System 2 processing causes one's brows to furrow, the association usually works in the opposite direction as well. When someone frowns, which naturally induces a furrowing of the brows, System 2 processing kicks in as a result (Daniel Kahneman, *Thinking Fast and Slow*, 2011, Farrar, Straus, and Giroux, p. 152). Unlike gratuitous visual difficulties, frowning won't impede progress.

The authors ignored another study involving the use of a difficult-to-read typeface, which is far more relevant to information visualization than the one that they cited. This other study found that people view difficult-to-read text as less trustworthy than easy-to-read text (ibid, p. 63). We don't ordinarily want data analysts or other readers of graphs to distrust the data, so poorly legible text would work against our intentions. We might want to use less-fluent (e.g., lighter) representations of data, however, to indicate uncertainty when it exists. That would be meaningful.

The authors bemoan the fact that "little research synthesizes the empirical evidence so as to derive the underlying forces that allow some effective graphs to contradict cognitive efficiency assumptions." What empirical evidence? They cite only one study related to graphs that even hints at this: research done by one of the authors, Priti Shah. Unfortunately, the study that they cite in the paper did not in fact address visual difficulties at all. What it did was mention another more recent, as-yet-unpublished study, which was also done by Shah. Without access to this study to see for ourselves what it claims and assess its merits, we can't count it as evidence.

So far we've found that, of the "large body" of psychology and education research cited to by the authors to support their claims, none seem to do so.

More Erroneous Validation

After misappropriating psychology and education research, the authors attempt other validation strategies that involve equal mismatches. For example, they cite "small multiples" as a form of visual difficulty when they're used as an alternative to animated charts, when in fact small multiples work because they reduce cognitive load by placing everything that must be compared in front of the user's eyes at once, thus reducing reliance on working memory. They also cite aesthetic aspects of graph design as an example of visual difficulty. Displaying information in aesthetically pleasing ways, however, does not represent visual difficulty. Quite the opposite, when done with skill. Aesthetics only produce visual difficulty when applied to extraneous visual elements that add no value to the visualization or in ways that suggest meanings that don't exist.

In an attempted counterpoint to Tufte's data-ink ratio, the authors demonstrate by reversing the meaning of the term that they don't understand the concept. They state that a chart with a high data-ink ratio is one with decoration or embellishment. In fact, a high data-ink ratio indicates that most of the ink represents data (i.e., the ratio of data ink to the total amount of ink is high), which is beneficial according to Tufte. They then embrace the effectiveness of 3D, stating that, according to Levy, et al, graphs that include 3-D effects are sometimes preferred. Preferred? Such as in the way that many people prefer pie charts? This study by Levy, et al, tested the graph preferences of Stanford students, not the effectiveness of particular graph types. As we all know, people too often prefer things that aren't good for them.

The authors also enlist visual complexity as an example of beneficial visual difficulty. They state, "Yet under some conditions, using a complex visual organization may be preferable, such as when the complex visual more accurately represents the target relationship and is presented with a prediction task to insure that users will reflectively engage with the content." When relevant complexity exists in data, we must indeed find a way to represent that complexity visually to encourage comprehensive thinking. This is hardly an example of gratuitous visual difficulty.

The Deeper Problem

My intention in this article is not to revile the authors of this paper. Point out their errors, yes, but not demean or discourage them. It isn't fair to Jessica Hullman, the student author of this paper, that her work is being scrutinized publicly. This paper should have never seen the light of day. With proper academic oversight and guidance, Ms. Hullman might have been directed into a worthwhile research project, conducted with scientific rigor. Had that happened, rather than pointing out her errors I might be showcasing her work. What bothers me is not the failure of this one study, but the excessive prevalence of poorly done research in the field of information visualization. How do these efforts get past the watchful eye of an experienced advisor? How do they get past VisWeek's review process? How do they get chosen as examples of our best work? Once a paper is presented at a major conference and published in an academic journal, it lives forever, part of our venerable body of knowledge. Others will cite it to support their causes. Others will draw inspiration from it and extend its errors through their own misguided efforts. Even worse, from my perspective, data visualization practitioners will follow its advice.

This concerns me. I think it should concern all of us. With papers such as this as benchmarks of information visualization research, it is no wonder that we feel insecure about our status among academic disciplines. We have much to offer, but those out there in the world who stand to benefit from our efforts shouldn't have to dig through piles of mediocrity and downright error to find it. We can do better. Let's get our house in order.

The nature of my work gives me a different perspective than most people in the information visualization research community. I spend much of my time addressing the needs of information visualization practitioners—people who apply what researchers come up with to solve real problems in the world. Most of the people that I support rely on tools from large software vendors such as Microsoft, SAP Business Objects, IBM Cognos, Oracle, and MicroStrategy. Few practitioners today have access to good commercial products for data exploration and sensemaking, such as Tableau, Spotfire, Panopticon, or SAS JMP. Most commercial tools do a horrible job of visualizing information and supporting the kind of interactions that are needed to explore and make sense of it. I spend my time trying to help practitioners and vendors understand what really works. In concrete terms, this means that I must fight the popular notion that data visualization ought to look like the following dashboard:

It shouldn't be hard for you to imagine why I respond with concern when an academic paper claims that visual difficulties such as the meaningless embellishments in this example are beneficial. When claims like this are granted credibility by the research community, they carry weight and have an effect. Researchers should be free to investigate anything that seems potentially worthwhile, but should do so responsibly.

Where is the system of checks and balances that promotes good work and corrects research that goes astray? Many colleagues in the information visualization research community share my concerns, but are happy to let me voice them while they remain silent. I promote good work and expose the bad, not because it gives me pleasure—I hate conflict and dwelling on the negative as much as anyone—but because it matters and needs to be done. Where are the rest of you, especially you who are more directly involved in academic research and have greater opportunities than I to fix this? If what I'm saying is wrong, then tell me so, but if it's right, please lend your voice and your heart in tangible ways to the cause of better information visualization research.

Discuss this Article

Share your thoughts about this article by visiting the Benefitting InfoVis with Visual Difficulties? thread in our discussion forum.

About the Author

Stephen Few has worked for over 25 years as an IT innovator, consultant, and teacher. Today, as Principal of the consultancy Perceptual Edge, Stephen focuses on data visualization for analyzing and communicating quantitative business information. He provides training and consulting services, writes the quarterly *Visual Business Intelligence Newsletter*, and speaks frequently at conferences. He is the author of three books: *Show Me the Numbers: Designing Tables and Graphs to Enlighten*, *Information Dashboard Design: The Effective Visual Communication of Data*, and *Now You See It: Simple Visualization Techniques for Quantitative Analysis*. You can learn more about Stephen's work and access an entire library of articles at www.perceptualedge.com. Between articles, you can read Stephen's thoughts on the industry in his blog.