

Data Visualization Effectiveness Profile

Stephen Few, Perceptual Edge
Visual Business Intelligence Newsletter
January/February/March 2017

We visualize data, initially to make sense of it, and eventually to communicate our findings to others. We do so when we believe that data must be visually represented to understand its meanings. Not all data visualizations are equal in their ability to communicate data effectively. They range from providing no useful understanding whatsoever to providing all of the understanding that's needed for the task at hand. An effective visualization satisfies a particular set of criteria that address several aspects of effectiveness. The purpose of this article is to identify and describe these criteria for assessing the effectiveness of data visualizations and to propose a way to graphically represent how well specific visualizations satisfy them.

Critiquing the effectiveness of our own visualizations using a consistent set of criteria can serve as a simple checklist of sorts that we can use to confirm the quality of our work before releasing it into the wild where it will affect others. Critiquing the effectiveness of visualizations created by others can serve as a great learning exercise to hone our own skills. A consistent set of criteria provides the means to compare the effectiveness of various visualizations and to discuss their relative merits with others using a common language. Shown together, how well a visualization achieves these criteria can serve as its *data visualization effectiveness profile*.

The data visualization effectiveness profile that I'm proposing consists of seven criteria, which fall into two general categories: criteria that address the degree to which a visualization is *informative* (i.e., produces understanding) and criteria that address the degree to which it is *emotive* (i.e., produces a useful emotional response):

- Informative
 - Usefulness
 - Completeness
 - Perceptibility
 - Truthfulness
 - Intuitiveness
- Emotive
 - Aesthetics
 - Engagement

What I mean by each of these criteria requires explanation, which we'll get to in a moment. First, however, we should acknowledge that the effectiveness of a data visualization can only be fully determined in light of its creator's intentions and its audience's needs. When we create a data visualization, we do so with particular objectives in mind—or at least we should. We visualize data to help people understand particular facts and what they mean. We design the visualization to feature those facts in a way that can be understood and to a degree that is useful. Any data set can be visualized in many ways, depending on the meanings contained therein that we choose to communicate. Visualizing data is an act of interpretation. What we choose to include and how we display it transforms data into a message.

Let me take this a little further by expressing a value-based commitment that I always bring to data visualization: we should inform people to the degree that will serve their needs as well as possible. We should not provide partial understanding when full understanding would benefit people most, nor should we attempt to explain everything when less is needed. As in all forms of communication, understanding the needs of our audience is essential.

Now let's examine each of the criteria in turn for determining a data visualization's effectiveness.

Usefulness

This first criterion is one that can only be determined in light of the audience's needs. A data visualization is of little value if it helps people understand something that doesn't matter to them. In my years of experience helping organizations—mostly businesses—use data effectively, I have observed that most of the data that is routinely provided in the form of reports and dashboards is of little value to them. The knowledge that they provide is not useful. Most organizations still do a poor job of identifying the information that they should understand to do their work effectively. As such, they waste a great deal of time ingesting data consisting of empty calories, for it will never lead them to make decisions or take actions to benefit their organization. It is a sad irony that in the so-called information age most organizations remain malnourished.

We want to provide information that is useful. If it isn't useful, it isn't important. A visualization can be eloquent in every way, but if it communicates information that is of no value, it is ineffective. It's that simple.

Usefulness _____
Useless Very useful

Evaluating the usefulness of a data visualization is subjective, for it is based on an assessment of the needs and values of others, but that does not diminish the relevance of this criterion and our attempts to measure it.

Completeness

An effective data visualization includes all of the information that's needed to produce the intended level of understanding, but not more. This involves the right information and the right amount of it. This also means that all of the context that's needed to understand the information has been provided as well. Context is usually provided in the form of comparisons, such as comparisons to targets, measures of the norm, and historical values.

Completeness _____
No relevant data All relevant data

Perceptibility

The information must be displayed in a manner that the human eye and brain can perceive with minimal effort and appropriate precision. This involves selecting the type of graph that displays the information most effectively and designing it in a way that presents the information as clearly as possible. Forcing people to compare the sizes or color intensities of objects, which is somewhat difficult and imprecise, when the positions or lengths of objects could have been used instead, reduces perceptibility.

Perceptibility _____
Unclear and difficult Clear and easy

Truthfulness

By truthfulness, I mean the degree to which a data visualization is accurate and valid. Accuracy is a measure of reliability and appropriate precision. Validity indicates how well something represents what it claims. Data visualizations should exhibit both attributes to the fullest. If a bar in a bar graph that is twice the height of another represents anything other than twice the value, it is inaccurate. If a graph presents a country's median household income as a measure of its people's happiness, it is invalid.




Intuitiveness

A data visualization is intuitive to the degree that it is familiar and easily understood. There are times when an unfamiliar form of display is preferable because it communicates the intended message more clearly than any familiar form of display could do. In these cases, the unfamiliar form of display should not be more difficult than necessary to understand and should be accompanied by simple instructions that enable its use with relatively little learning. Whether or not a particular form of display is intuitive depends entirely on the audience. A parallel coordinates plot would be unfamiliar to most members of the general public, but familiar to select groups of scientists and statisticians.



Aesthetics

The aesthetic quality of a data visualization can range from ugly (a display that no one would choose to view for long) to beautiful (one that would invite even those who care little for the information to take perhaps delight in learning about it). In between ugly and beautiful lie data visualizations that are merely pleasing to the eye, which is usually adequate, except on those occasions when the audience must be given an extra incentive beyond their own natural interests to examine what they might otherwise ignore.

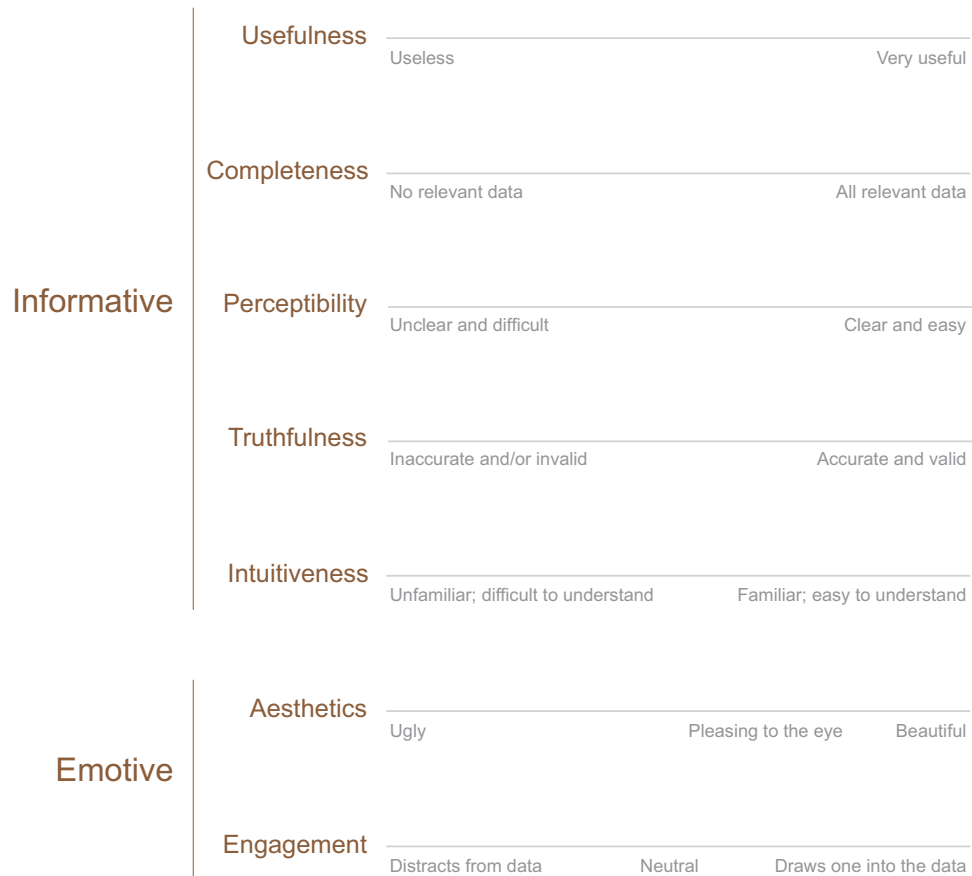


Engagement

By engagement, I mean a quality that can be achieved by various means, including but not limited to aesthetics, that invites the audience to examine the information. I do not mean visual effects or ornamentation that entice people to examine those qualities only without becoming engaged with the information itself in a way that leads to understanding.



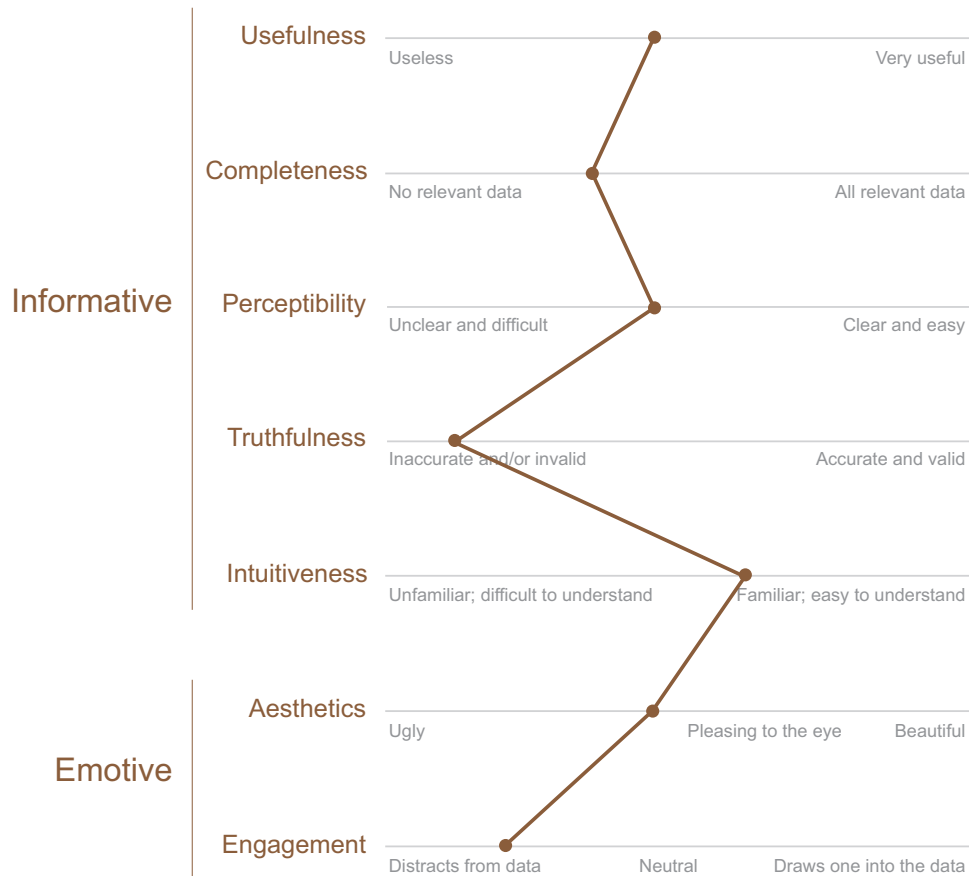
Now that we've covered all seven criteria, here's the full set shown together:



This arrangement can be used to display the effectiveness profile of a particular visualization. To illustrate, let me begin by profiling a typical business graph created with Excel. Here's the graph:

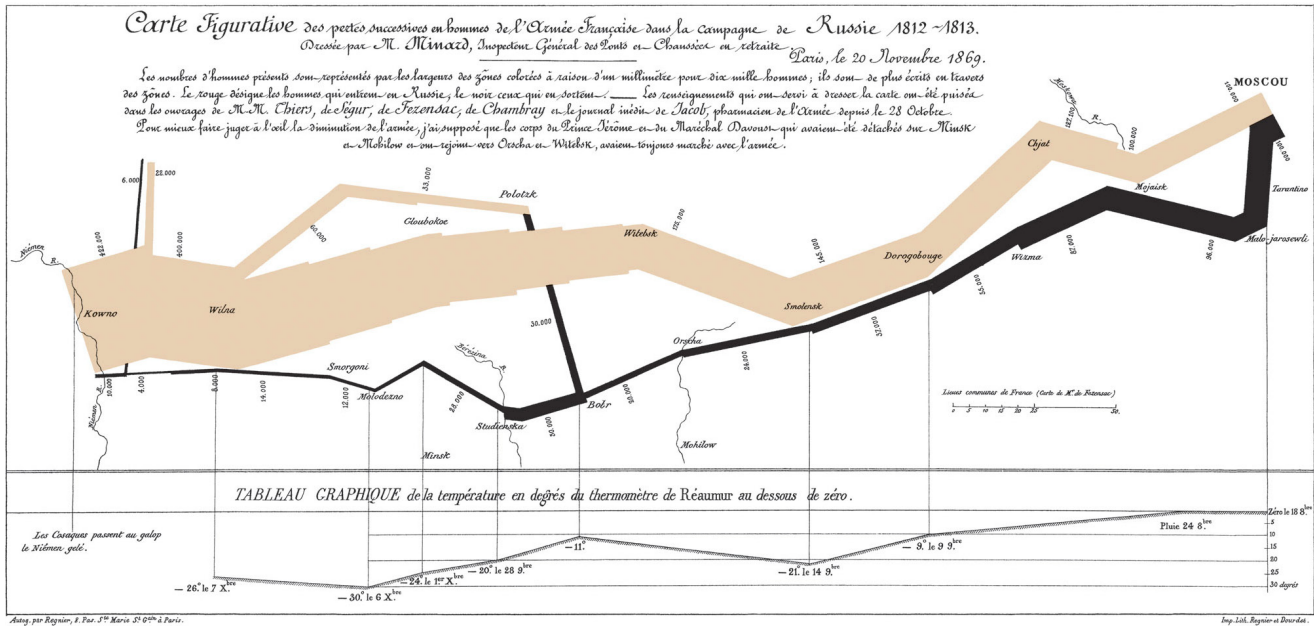


And here's an assessment of the graph's effectiveness:

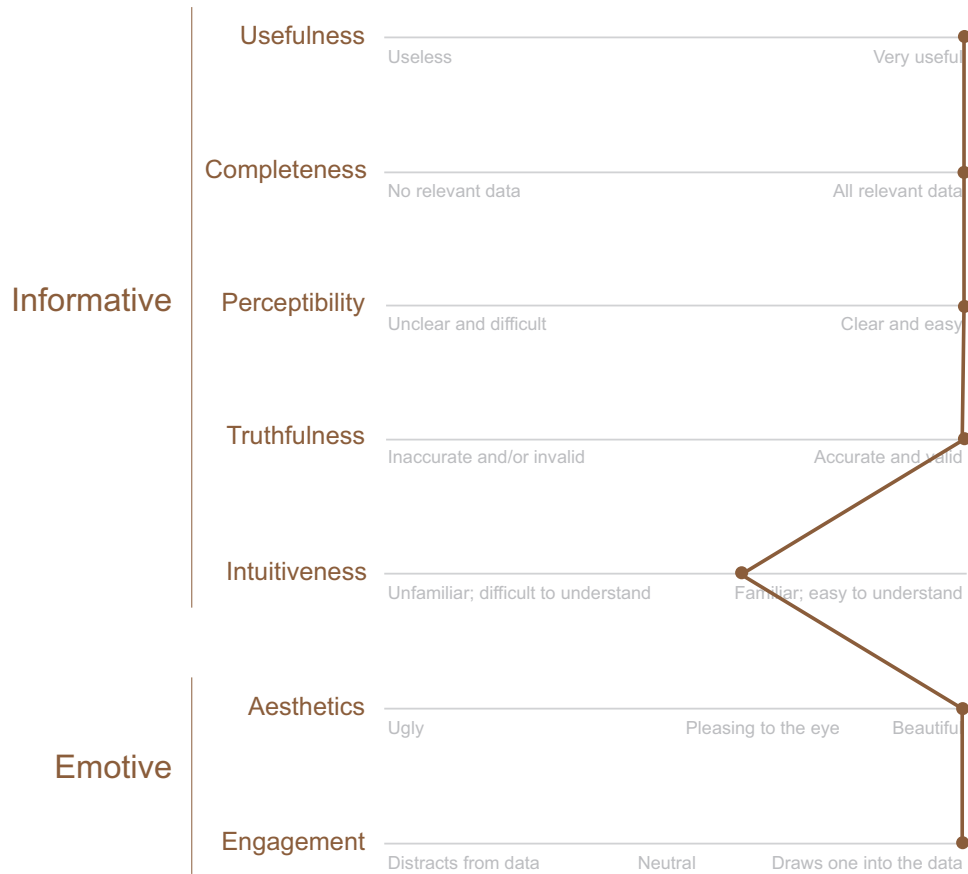


Sales revenues are certainly important, but revenues alone don't tell the story and are therefore of limited usefulness by themselves. In fact, these monthly revenues alone give us no way of determining if sales are performing well, so the data is significantly incomplete. Perceptibility suffers due to the distracting 3-D effects and the fact that bars don't show the pattern of change as clearly as a line would. We don't know if the title of this graph, "SALES ARE IMPROVING," is truthful because no insufficient historical context has been provided and we definitely know that the bars don't accurately represent the values because the scale doesn't start at zero. This bar graph would be quite intuitive if the scale started at zero, but the truncated scale could easily lead viewers to read it incorrectly. Aesthetically, it isn't quite pleasing to the eye, in part because the values at the tops of the bars create a cluttered appearance. And finally, nothing about this design invites the viewer to examine the data and, in fact, the 3-D effects distract from the data. Excel can be used to create highly effective graphs, but seldom is.

Now that I've illustrated the process, let's profile a well-known historical example: Minard's visualization of Napoleon's march on Moscow.

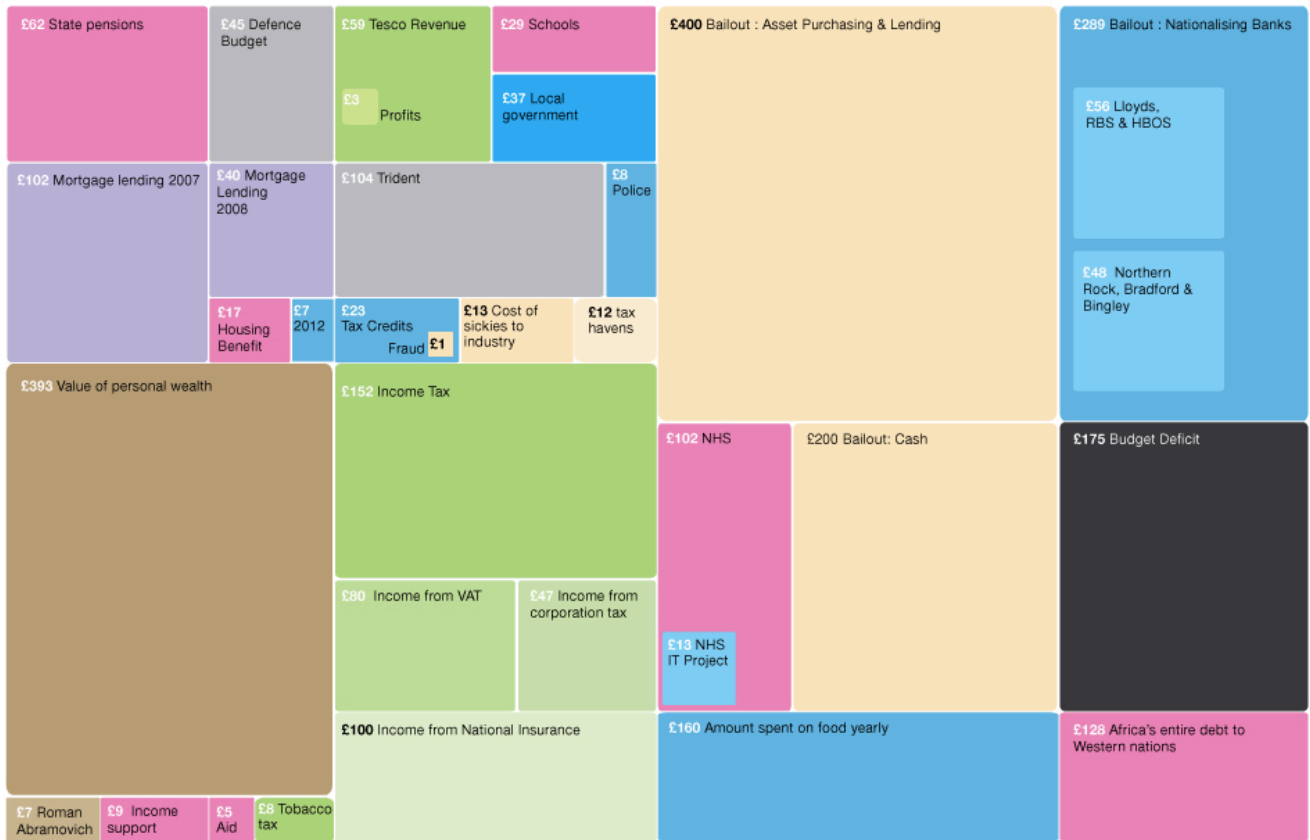


Some believe this to be the finest data visualization ever made. I wouldn't go this far, but it is certainly a brilliant piece of work. Here's how I would profile its effectiveness:



According to my assessment, Minard’s visualization is optimal in all respects except for a lower score in intuitiveness, because it does take some time, when first encountered, to figure out how to read it. While it is true that the varying thickness of the path taken by the army cannot be easily perceived except as a rough approximation of the army’s size, there is no better way to incorporate these values along the path itself, so I haven’t lowered the perceptibility score.

Let’s consider another fairly well-known example, this time one that doesn’t fare so well: David McCandless’ “The Billion Pound-O-Gram.”



The Billion Pound-O-Gram

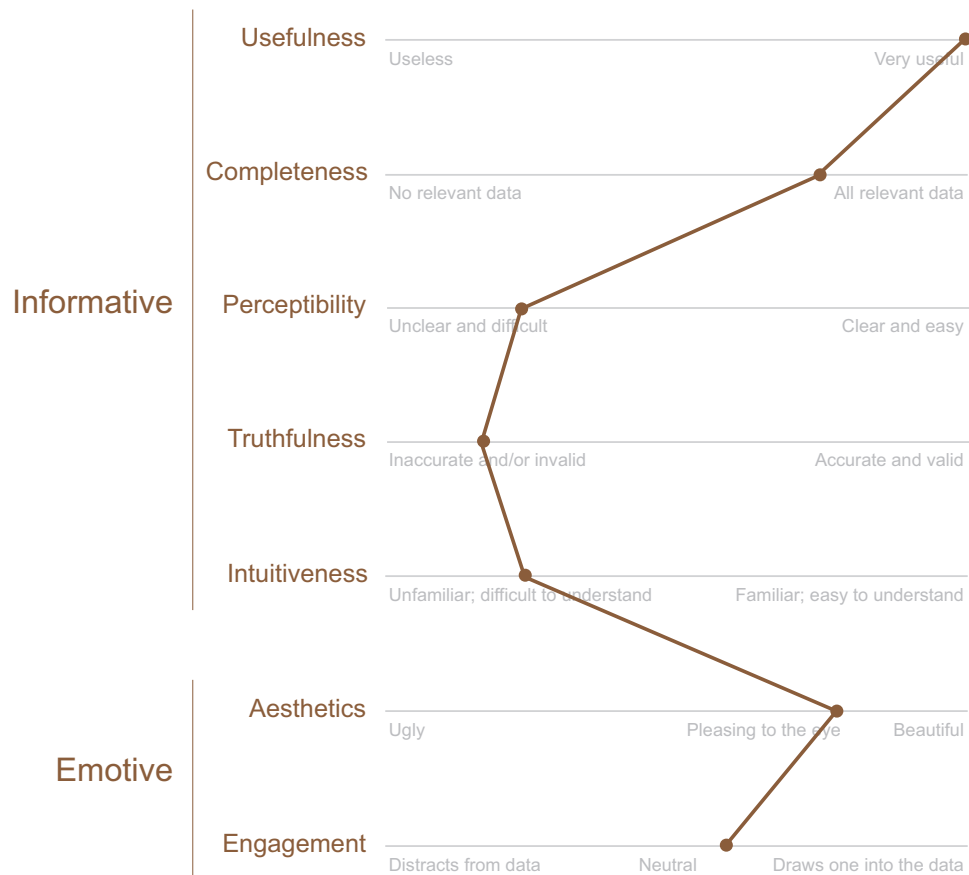
Legend: Giving (pink), Spending (blue), Fighting (grey), Hoarding (brown), Lending (purple), Bailing (yellow), Earning (green)

David McCandless / InformationIsBeautiful.net

Source: UK Treasury, Guardian


The purpose of this chart was to show how great the U.K.’s budget deficit was. Unfortunately, this treemap made comparisons unnecessarily difficult and the values were not accurately represented (e.g., notice the size of “Africa’s entire debt” compared to the “Budget Deficit” in the bottom-right corner).

Here's my assessment of The Billion Pound-O-Gram's effectiveness:

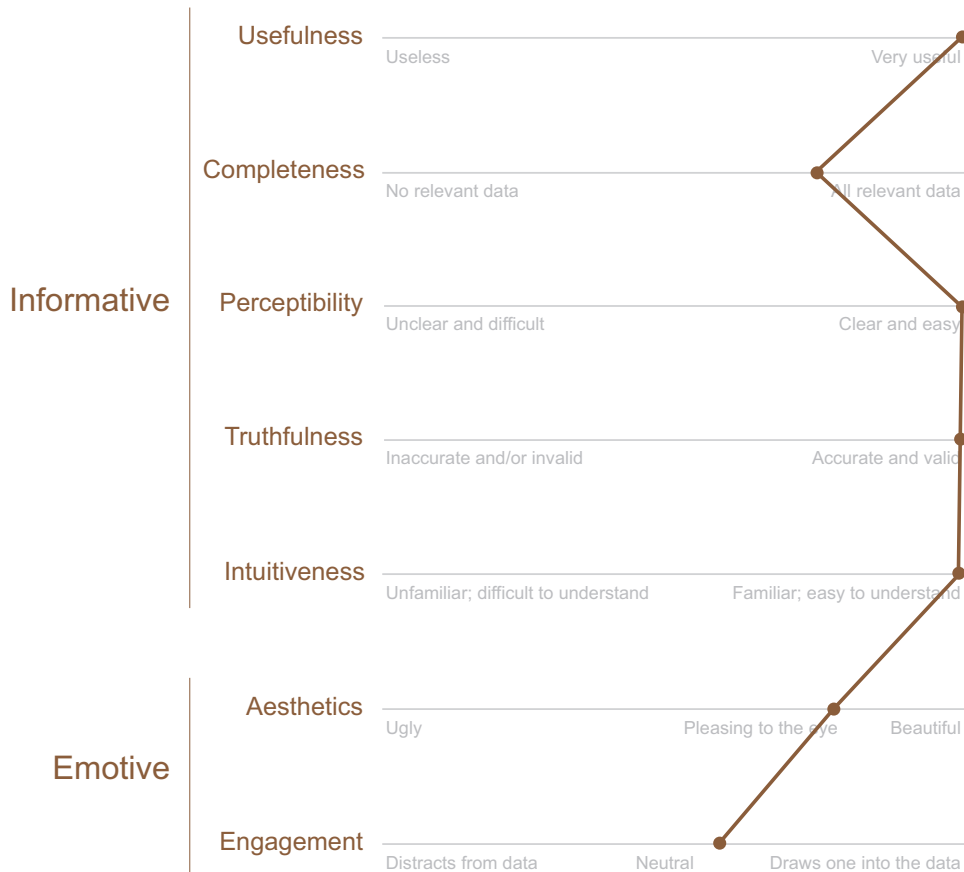


When this originally appeared in the *Guardian*, it was probably the first example of a treemap that many people had seen, so its novelty provided a fair amount of engagement but relatively little meaningful engagement with the data. You might assess the effectiveness of this chart differently than I have, and that's fine. Let's not get distracted by these particulars, but focus instead on the potential usefulness of the data visualization effectiveness profile. Years ago, I critiqued The Billion Pound-O-Gram and proposed the following redesign that solved most of its problems.

The UK Budget Deficit Compared to Other Big Amounts of Money (Billions of Pounds)

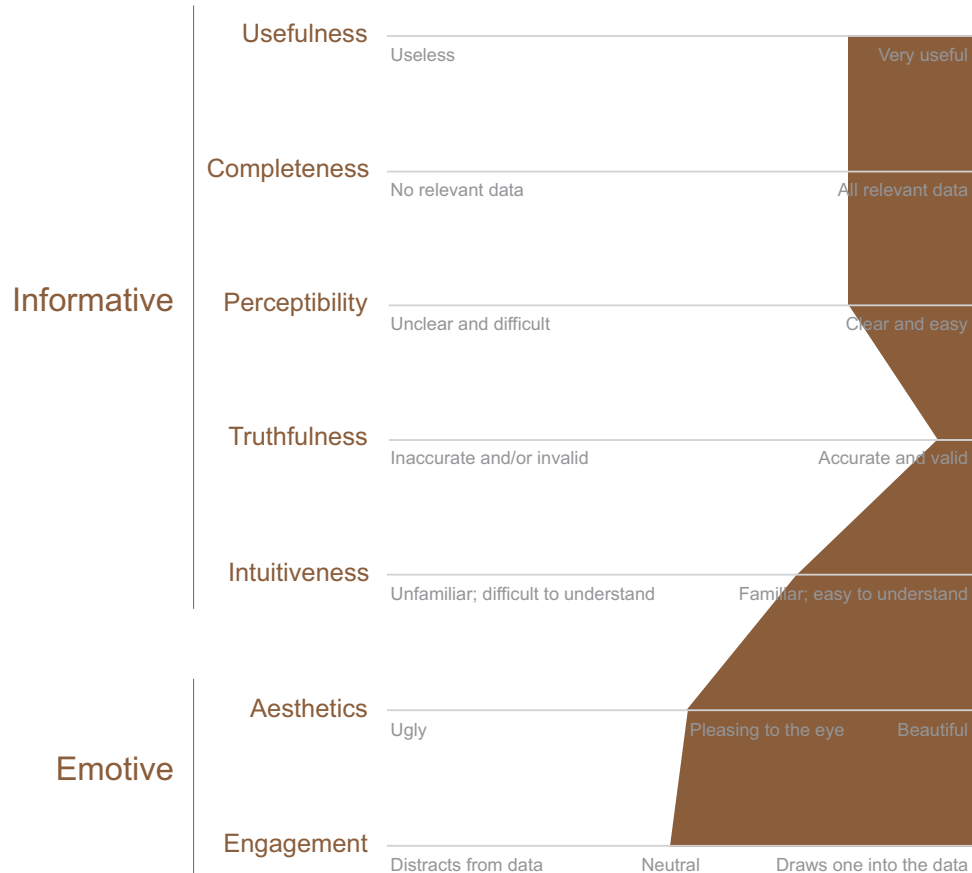


Here's my own assessment of the redesigned display:



By sticking with the same set of values that McCandless used in his original chart, I was not able to improve the completeness of the data, but I was able to improve it in most other respects. Although my redesign didn't exhibit the novelty of the original treemap, I don't believe engagement suffered because it makes the data much more accessible.

In general, we should shoot for data visualizations that fall within the following ranges:



You'll notice that I've allowed little latitude in the realm of truthfulness, but a great deal of latitude along both of the emotive criteria. This is because, in general, if the visualization presents data that is important to the audience, they will be naturally interested in it, so extra effort to catch their eyes or draw them into the data isn't necessary. This isn't always the case, however. A journalistic infographic that appears on the web might need something extra to grab readers' attention and get them to spend a few seconds deciphering information that isn't of great interest to them. In these cases, the acceptable ranges for aesthetics and engagement would be much narrower.

I haven't included a quantitative scale along the X-axis of the data visualization effectiveness profile, but one could be shown, such as a 0 to 10-point scale or one that extends from 0% to 100%. I have a slight preference for omitting a quantitative scale to avoid suggesting a level of quantitative precision that does not exist in these subjective measures. Notice also that I have not included an overall effectiveness score. It could be useful to have one when comparing the merits of several data visualizations, but an overall effectiveness score would require that the individual criteria be weighted to reflect the varying degrees to which they contribute to effectiveness. And finally, I'm not suggesting that these criteria provide scientific rigor, but instead function merely as a consistent guide for meaningful assessment. This data visualization effectiveness profile can no doubt be improved, and I invite you to suggest ways that it can. Our efforts to improve the effectiveness of the data visualizations that we produce are worthwhile, even if they're imperfect. This is my attempt to contribute to these efforts.

Discuss this Article

Share your thoughts about this article by visiting the [Data Visualization Effectiveness Profile](#) thread in our discussion forum.

About the Author

Stephen Few has worked for over 30 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 four books: *Show Me the Numbers: Designing Tables and Graphs to Enlighten*, Second Edition, *Information Dashboard Design: Displaying Data for at-a-Glance Monitoring*, Second Edition, *Now You See It: Simple Visualization Techniques for Quantitative Analysis*, and *Signal: Understanding What Matters in a World of Noise*. 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](#).