Our Irresistible Fascination with All Things Circular

Stephen Few, Perceptual Edge
Visual Business Intelligence Newsletter
March/April/May 2010

We humans are drawn to circles. From prehistoric gathering places in the round, to Hindu mandalas, to halos that float over the heads of saints, to modern pie charts, we find circles naturally satisfying. I suspect that those parts of us that crave symmetry, wholeness, and closure take comfort in them. A circular archetype must reside deep in our psyches, which swaddles us in a snug, corner-less nest. Despite the undeniable beauty of their perfect form, however, circles often fail to support the functions that we assign to them. This is especially so in the field of data visualization. Elegance of form is undermined when a square-peg function is forced upon a circle.

I've written at length in the past about the failures of pie charts and circular dashboard gauges, so I'll keep my comments brief about them here. Although the evidence of a pie chart's dysfunction is hard to ignore when you actually take a moment to inspect it, nothing that I teach is met with such fierce opposition as my low opinion of them. People cling to them aggressively. But it is almost always difficult to compare the slices of a pie, because visual perception supports only rough comparisons of areas and angles, which are the two primary ways that pie charts encode quantities.

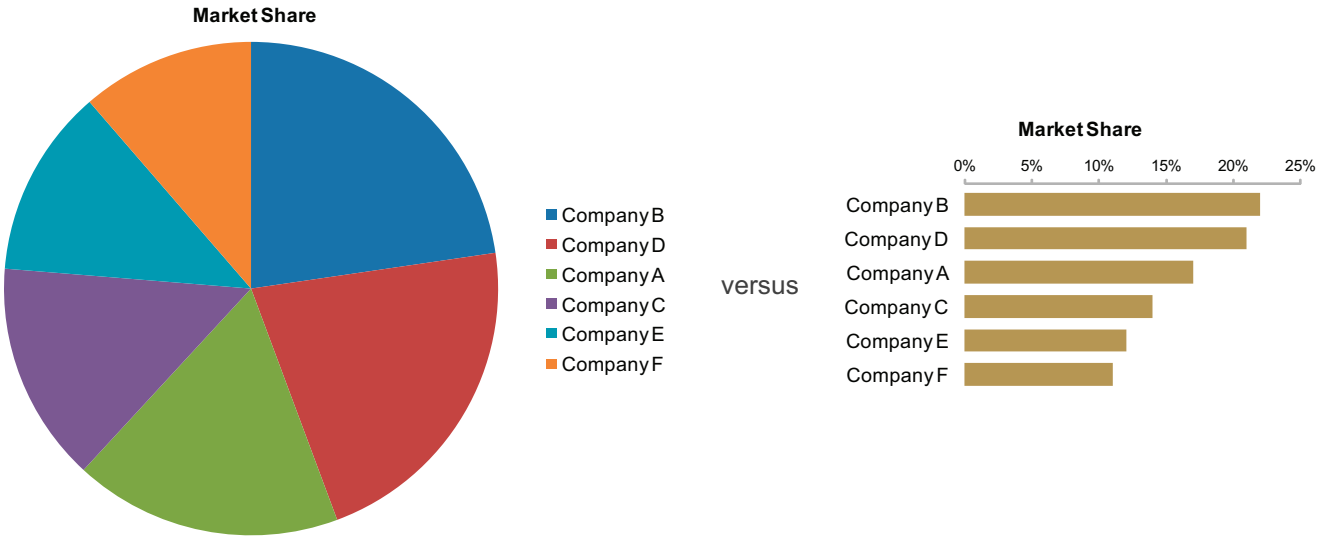

Circular gauges—the darlings of dashboards—waste space, but there is none to waste on a dashboard where a great deal of information must be displayed in a single screen. Nevertheless, software vendors insist on populating their chart libraries with every imaginable variation of cute gauges and speedometers, such as the typical examples from Dundas on the next page.

By their very nature, circular gauges use a great deal of space to say relatively little, and they fail spectacularly when intended for comparison, such as the following set of three.

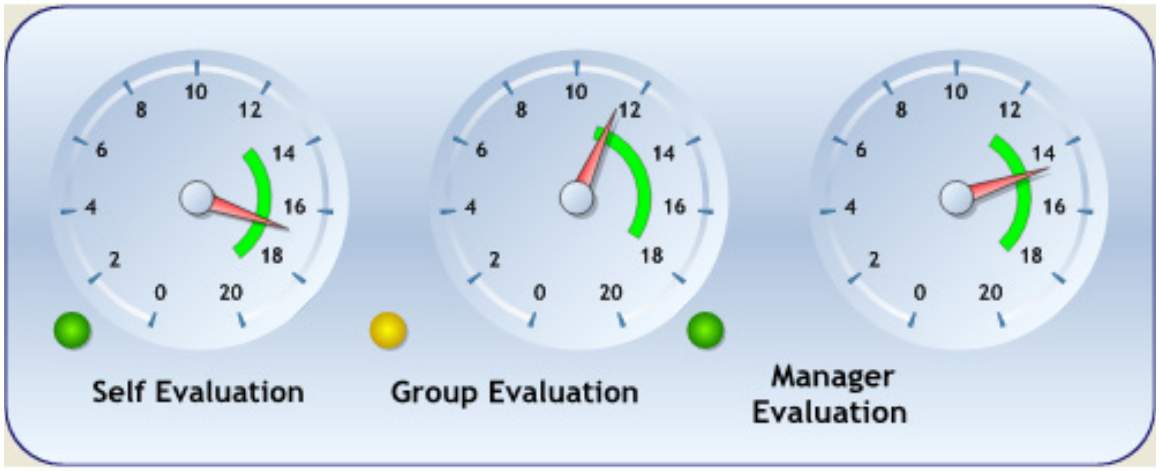

In the following display of the same three variables, I've used bullet graphs to show the evaluations of four individuals, rather than one, in less space and in a way that makes comparisons much easier and faster. The greater effectiveness of the bullet graphs is due in large part to their linear design.

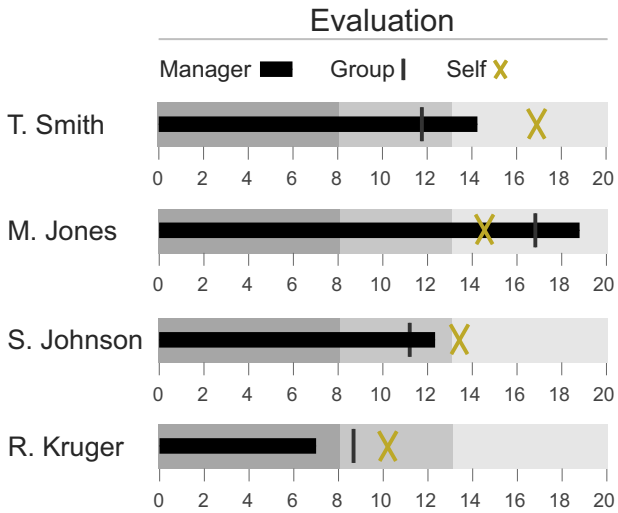

When we intend to inform, which is the essential purpose of data visualization, we can't select the form of the display arbitrarily: the form must be selected to support the intended function.

The chart below was published in Bloomberg in January of 2009. The title and legend make it clear that the market values of banks in quarter 2 of 2007 are being compared to their values in early 2009, after the credit crash had wreaked its havoc. It isn't obvious, however, that the specific purpose of this particular visualization was to show that J.P. Morgan suffered a lesser decline than all but one bank: Santander.

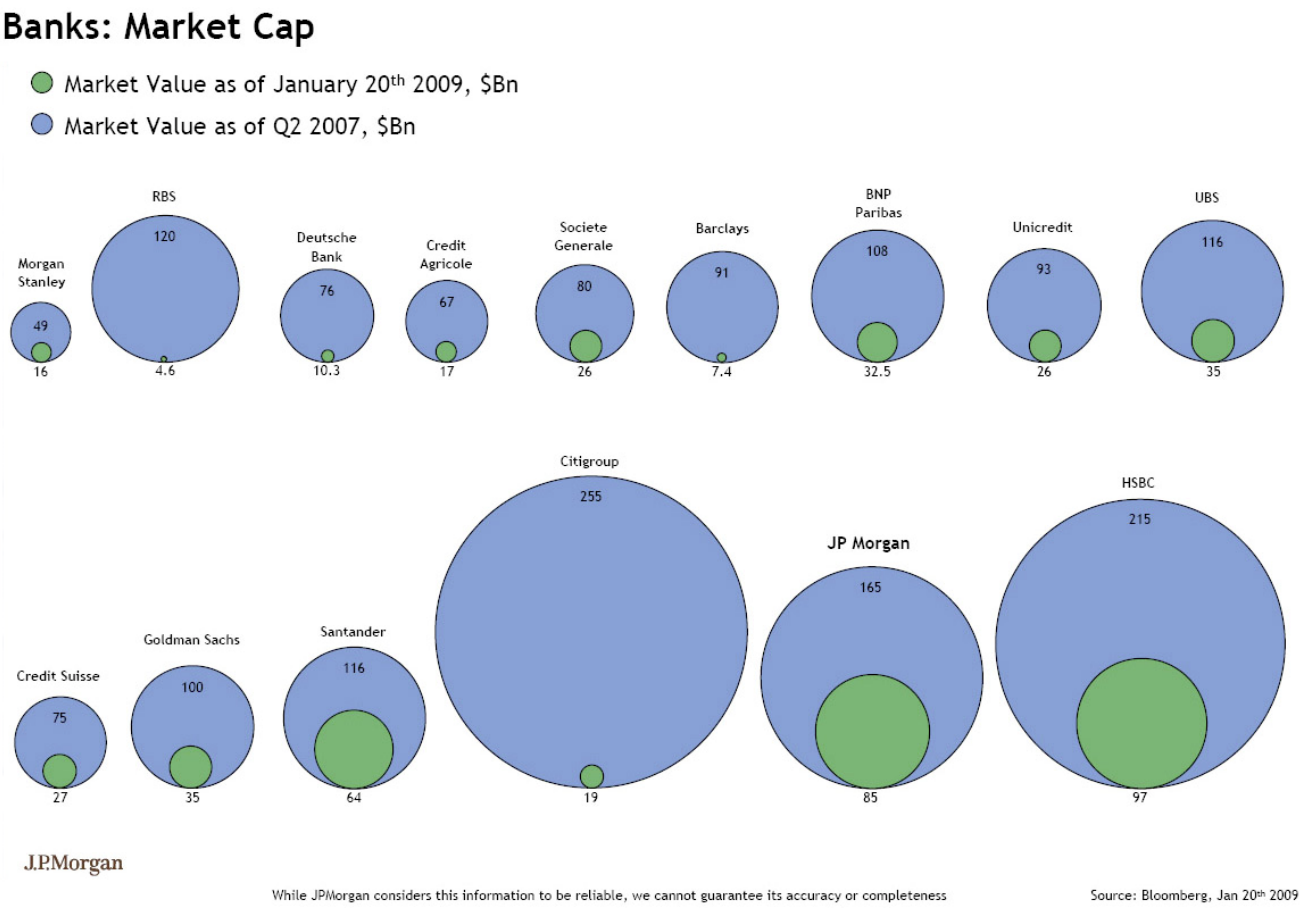

The primary failure that I want to point out, however, is not its failure to clearly support its purpose, but the fact that the designer responsible for it couldn't help using circles to represent the values, despite their inability to deliver. We simply can't compare the circles and get a good sense of their differences. Taking J.P. Morgan as an example, would you ever guess that its market value before the crash was slightly less than double its value afterwards? Actually, in this particular case if you did perceive it in this way, your perception would be way off, because the size of the green circle is actually much less than half the size of the blue circle, which is due to the fact that the designer mistakenly used the diameters of the circles rather than their areas to encode their relative values. Even knowing this, we still can't easily compare the diameters, because our eyes are naturally inclined to compare their areas, which are almost impossible to ignore.

Why did the designer of this display use circles when a bar graph would have done made the comparisons extremely easy to make with great accuracy, as I've shown in my redesign of the graph on the following page? Because people love circles. People love sugary candies as well, but we wouldn't rely on them for a nutritious meal, and in like manner we shouldn't rely on circles when we wish to compare quantities.

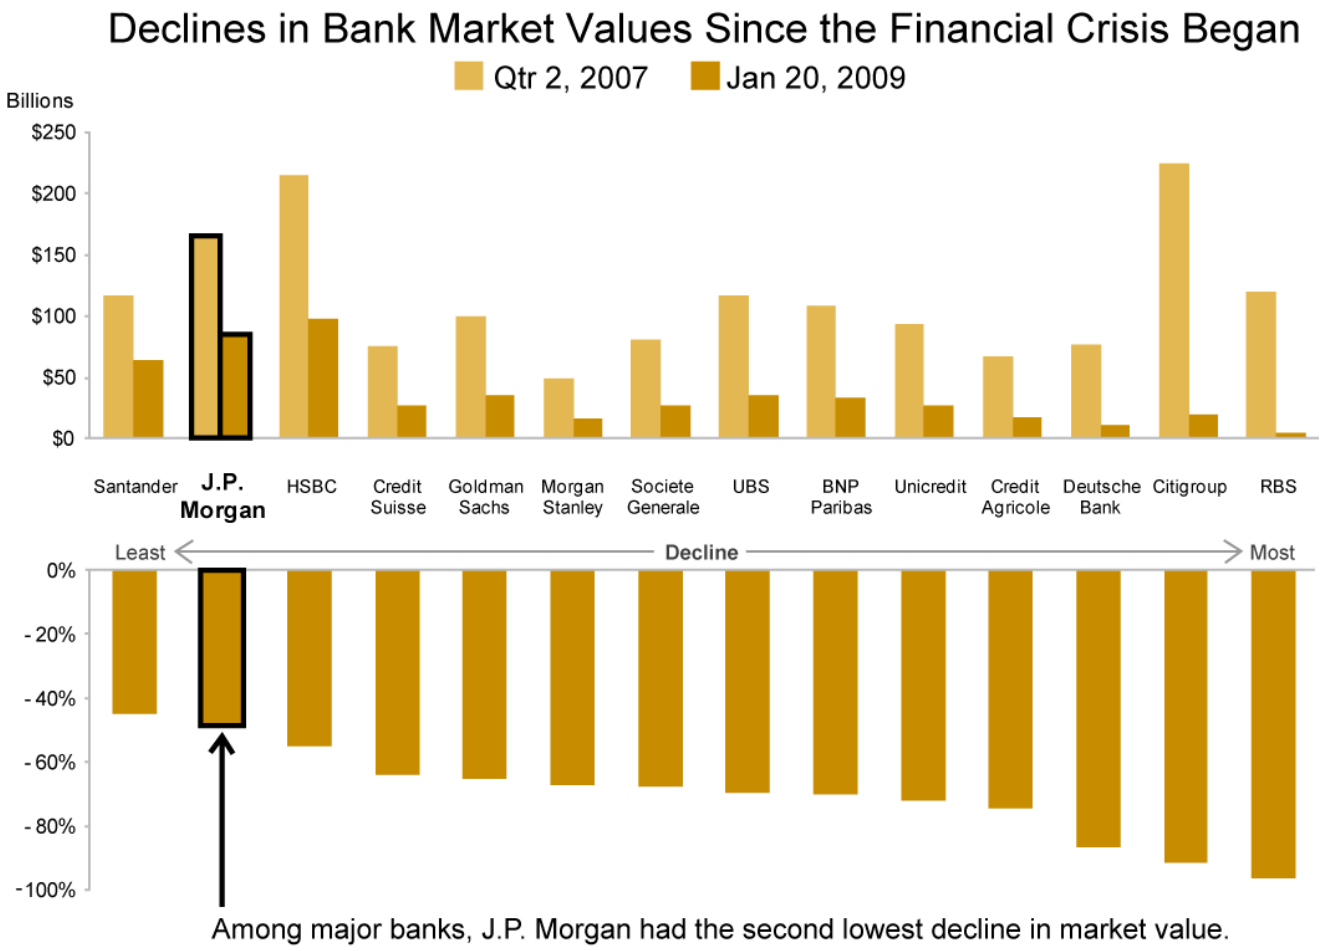

Notice that in the upper portion of my visual version of the story of J.P. Morgan's loss in market value, the side-by-side bars in the upper graph make the differences easy to discern and the lower graph directly and clearly shows that its loss was second least by sequencing each bank's percentage loss from least to most.

Designers of infographics often seem to have forgotten, or perhaps never learned, that circles rarely visualize quantitative data effectively. Shame on them. The following "Colours In Cultures" infographic appears on the cover of David McCandless' book *Information Is Beautiful*. I was reminded of it a few days ago when it was featured on a website as the "Infographic of the Day."

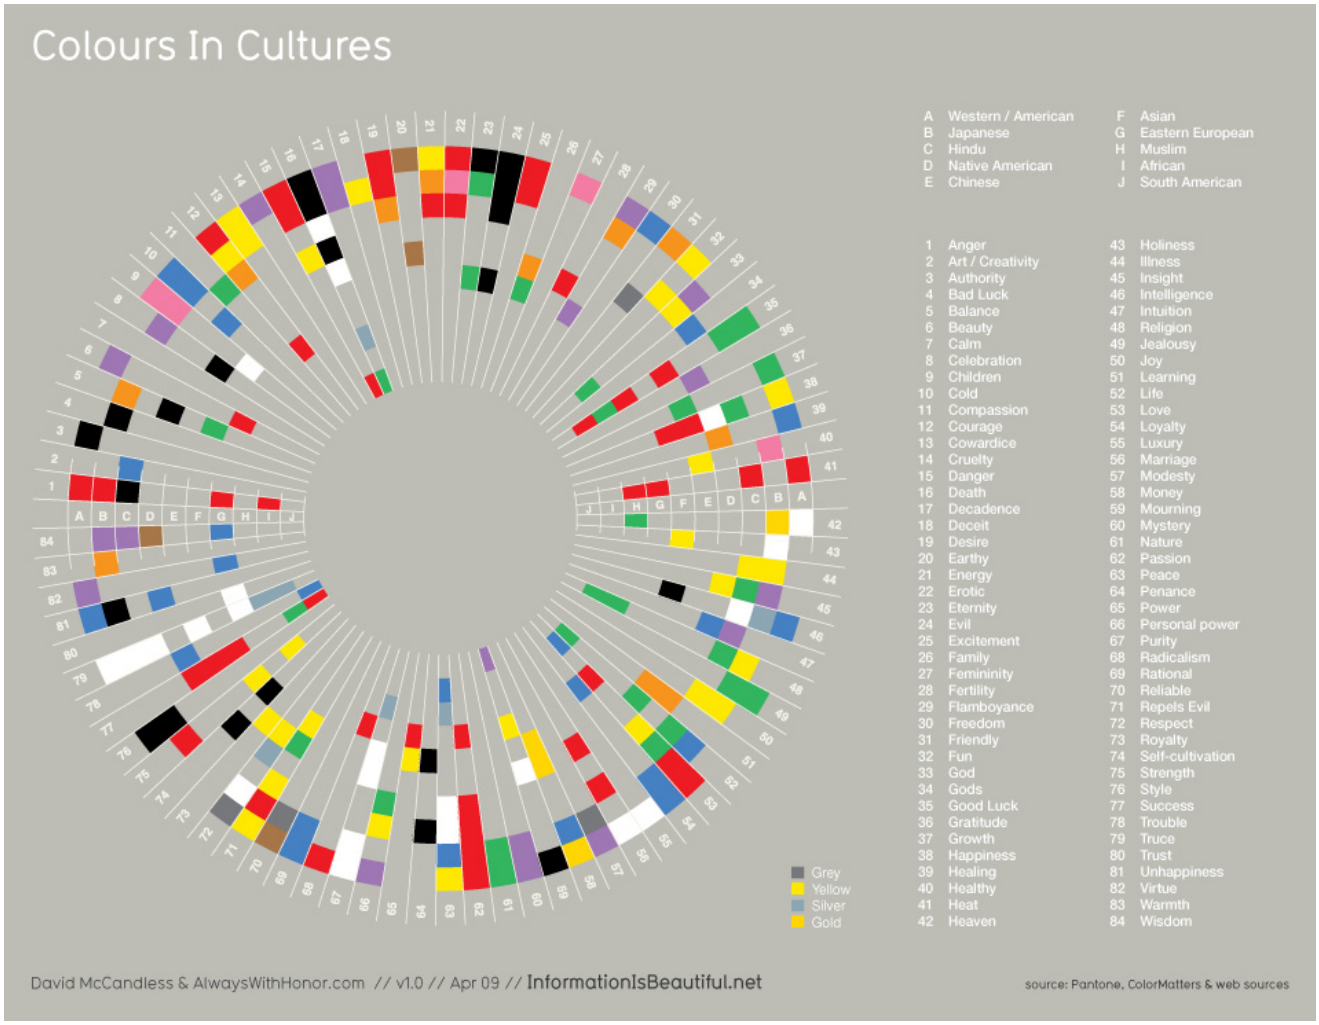

For the moment, let's assume that the information it contains is reliable and useful. (We'll question this later.) What are we supposed to learn about the meanings of colors in various cultures? What perceptual and intellectual processes are required to get this information from the chart? One apparent task involves looking up particular colors. For example, someone designing an information display for a particular culture might want to look up the color that's associated with a particular emotion or concept. Try it out. Take a moment (or longer) to look up the color associated with "mourning" in among Native Americans.

.....

Did you find the color white? If so, you goofed. According to this chart, the color that Native Americans associate with mourning is unknown or perhaps doesn't exist. It was difficult to find the intersection of "Mourning" (59) and "Native American" (D), in part due to the fact that you were forced to rely on a legend because the emotions/concepts and cultures are not labeled in the chart itself, and also because it's hard to trace the "D" position for Native American culture in a circle without losing orientation.

Try doing it again, but this time use my version of the chart on the next page.

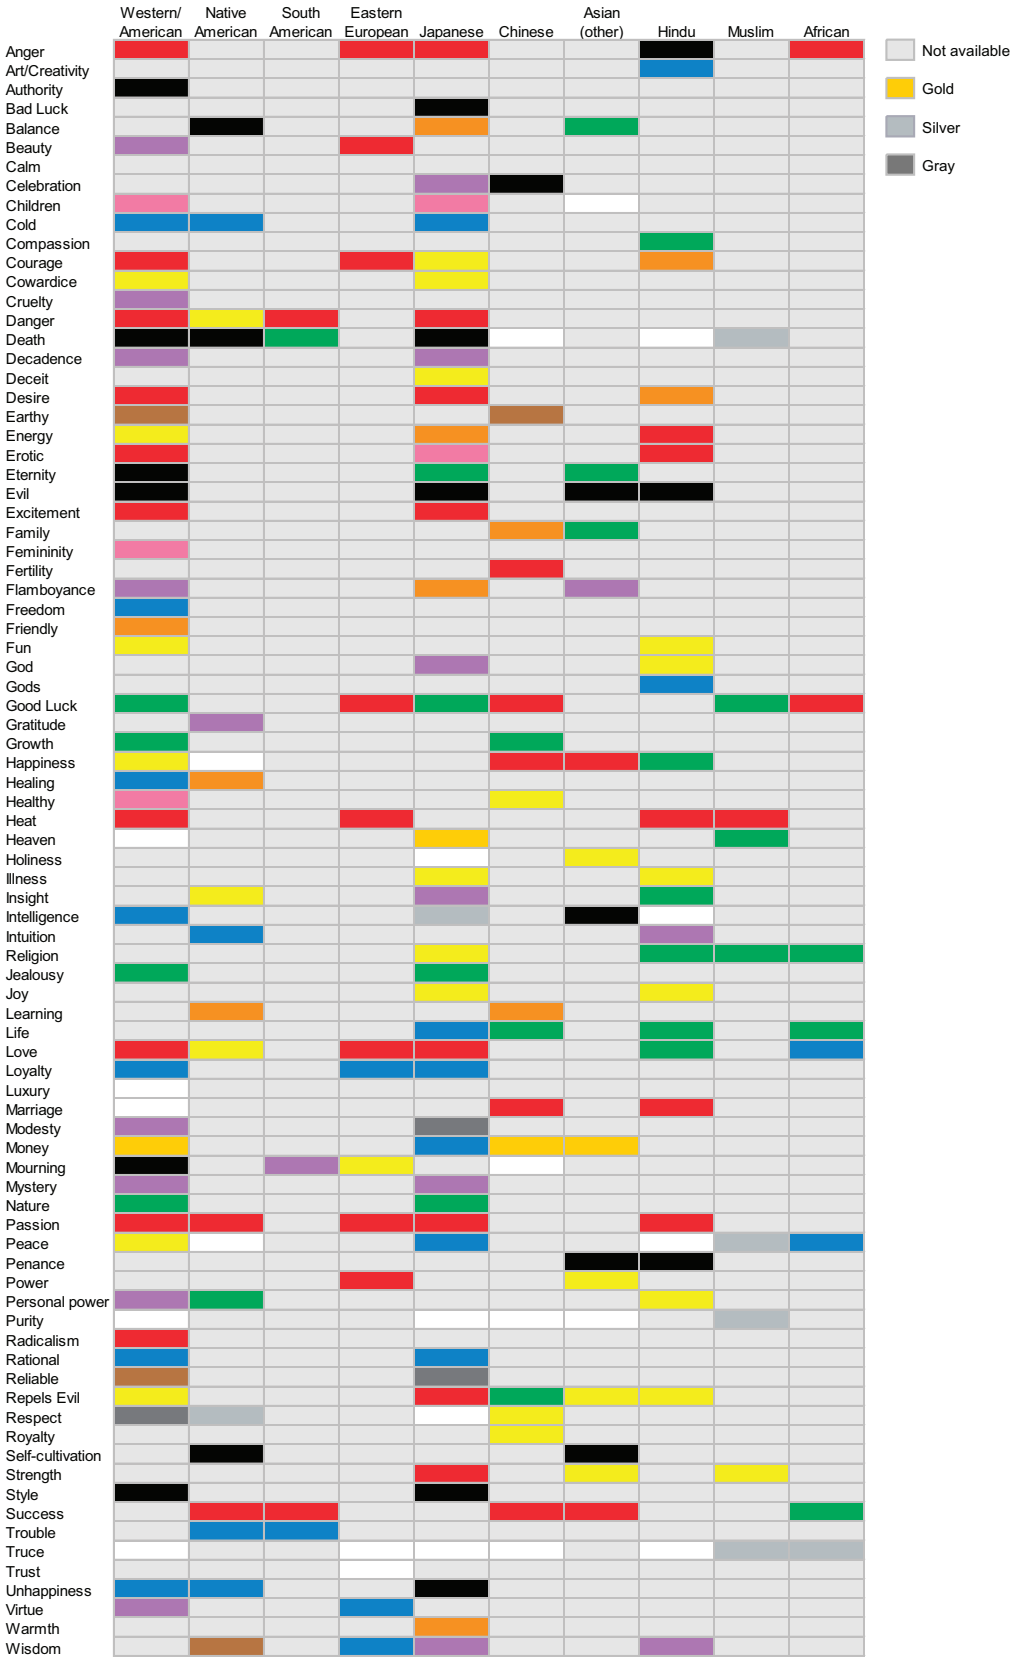

It was incredibly easy this time, wasn't it? That's because a linear, tabular arrangement of the data into columns and rows is a form that effectively supports this lookup task. You might have noticed that I changed the order of the cultures to group them geographically. This makes it easier to see similarities that are perhaps influenced by location.

Another purpose for which this chart might be used that interested me involved the different meanings that a culture associates with a particular color. Let's attempt an investigation of this type. Using McCandless' version of the chart, find the color that has more meanings than any other in Africa. Take your time.

…….

Did you come up with the answer "green," which appears three times for Africa? Had you used my version of the chart, what probably took you awhile could have been accomplished in a moment. Every purpose for which this chart might be used that I can imagine is better supported by a linear, tabular arrangement of the data.

Even more fundamental than this chart's design failures, I can't help but question the integrity and usefulness of the information. Is the color of modesty in America really purple and the color of intelligence really blue? Hmmm…this makes me wonder how the data was collected and how the populations were sampled. Also, what's called culture is in fact a mixture of categories: nationalities, geographical regions, and religions. Does all of Africa comprise a single cultural group? Hardly. How about Muslims throughout the world? Is it likely that their color associations are everywhere the same? Not unless all these associations are clearly delineated in the Koran, which I doubt. Because of these concerns, I probably won't be relying on this chart when designing data visualizations for particular cultural groups. For now, I'll stick with usability testing when communicating to an unfamiliar culture to insure that I use colors meaningfully.

Besides infographic designers, many of whom were trained in the graphic arts, another group that seems to favor the circular form consists of business intelligence software companies that attempt to incorporate data visualization functionality into their products without taking time to learn about the field, including an understanding of the human visual system. Of the many examples of this that I've encountered, one that crossed my path again recently is a chart called the "Temporal Super Graphic" from an aspiring BI vendor named Bis[2]. I critiqued the suite of Super Graphics unfavorably shortly after its release in the May/June 2009 issue of the *Visual Business Intelligence Newsletter*, titled "Cartographic Malpractice." If you've heard about this software, and especially if you're thinking about buying it, I strongly recommend that you read my article first. I won't repeat my extensive critique of Bis[2]'s Super Graphics here, but I will show the Temporal Super Graphic to illustrate another circular design failure.

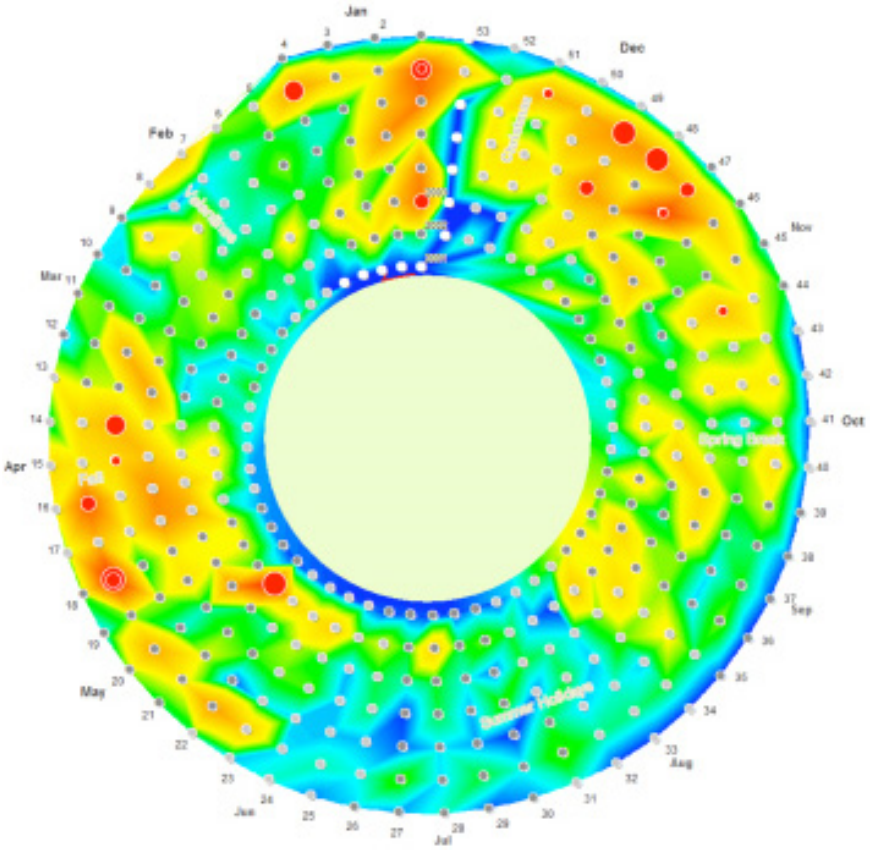

These Super Graphics were designed by Andrew Cardno, who is a cartographer by training. In designing these products, Cardno merely applied cartographic techniques to non-spatial displays, resulting in dysfunctional visualizations. He chose the round form for the Temporal Super Graphic because time is cyclical and we're accustomed to thinking of it in circular terms, based on familiarity with the traditional shape of clock faces. In an attempt to display multiple cycles of time (such as multiple years in the example above), he was forced to modify the perfect circle by switching to a spiral, with the first year beginning near the center (the first white dot at the inner 12:00 o'clock position) and spiraling outward in a counter-clockwise direction into larger and larger circles as if time expands. Of course, the most recent year is not actually longer than the previous years, but inaccurate representation of time apparently doesn't bother Cardno. This particular set of data would be much easier to read and understand if it were displayed as a single line chart with one line per year or as a series of line graphs (small multiples with a consistent quantitative scale), one graph per year arranged vertically, assuming that our purpose is to compare what happens across the years at particular points in time (such as during a particular month). If we want to compare more than a few years (this chart includes only seven), we might consider applying a sequential range of heatmap colors that vary in intensity from light to dark (not the rainbow colors used in the Super Graphic) in the form of a Horizon Graph (see the article "Time on the Horizon"), still sticking with the linear rather than circular design.

If the business intelligence industry hopes to succeed in its use of data visualization, it must discourage silly, ineffective products like this. It pains me to see organizations such as TDWI (The Data Warehousing Institute) promoting abysmal products like this (see "Q&A: Data Visualization Offers Rich Views", written for TDWI by Linda L. Briggs on April 14, 2010). This failure to vet a data visualization product before showcasing it stings because TDWI is an organization that I have taught, spoken, and written for in the past, but more importantly, it is an organization that people trust for good advice.

I would be remiss if I ended this article without showing at least one example of an effective circular data visualization. One of the few things that they can display well are the relationships between a single list of items. The following example, titled "Naming Names", appeared in the New York Times on December 15, 2007 during the early days of the last U.S. presidential campaign when several candidates were in the running. Each line from one candidate to another represents a time when a candidate referred to the other by name during a debate. Arrow heads group all incoming references to a specific candidate from all sources. This display allows us to discern patterns, such as the fact that Hillary Clinton had been mentioned a great deal by candidates in both political parties, while several candidates were only mentioned by others in their own party (for example, Richardson and Huckabee).

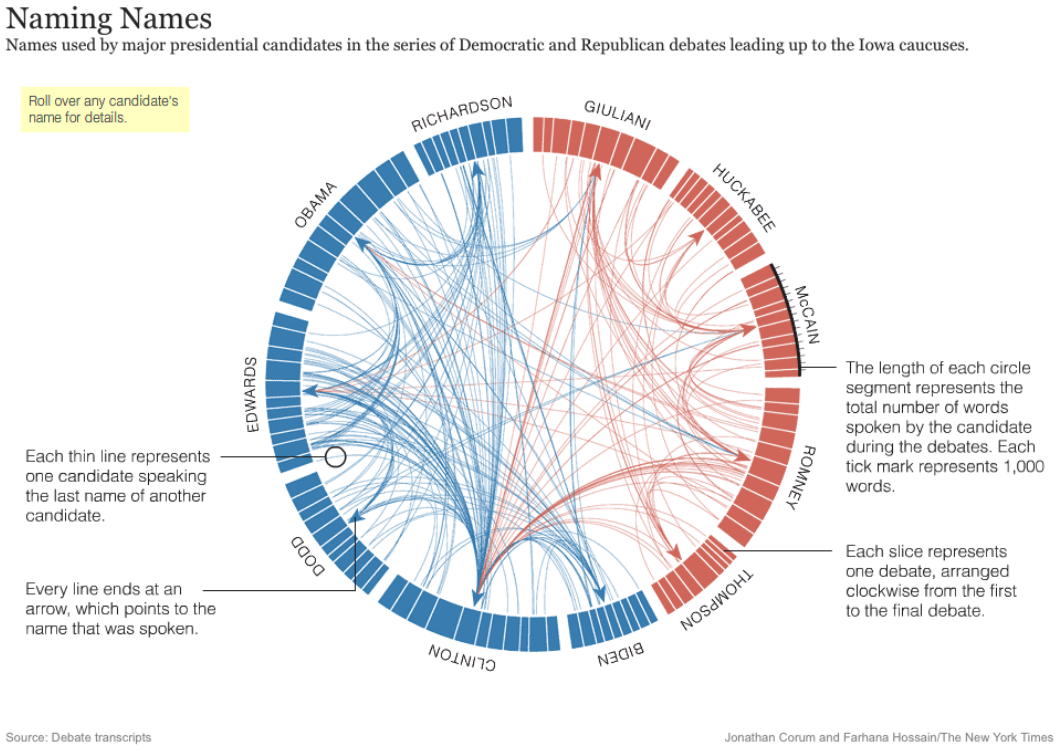

The circular arrangement of a single list of comparable items—presidential candidates in this case—lends itself nicely to the display of relationships between them, represented by lines that connect them. In this case, a linear arrangement would not work as well, as you can probably imagine.

As with most of my articles, the case that I'm trying to make here runs deeper than the failure of circles when quantitative displays are forced to fit into them, despite a mismatch of form and function. More fundamentally, I'm arguing that the primary goal of any data visualization is to inform, and that it should be carefully and knowledgably designed to support this goal based on the principles that we've learned through years of research about graphical perception. We should stick with what works. When we stray from these principles, we should do so because the situation demands it, which is at times the case. If we never learn the principles and understand how and why they work, we'll never know when it's appropriate to stray.

Discuss this Article

Share your thoughts about this article by visiting the Our Irresistible Fascination with All Things Circular 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 bi-monthly *Visual Business Intelligence Newsletter*, speaks frequently at conferences, and teaches in the MBA program at the University of California, Berkeley. 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.