



Graphic Grist for the Mill

Stephen Few

June 2005

I recently received an e-mail from a reader and fellow data visualization consultant named Sue Grist. Although we haven't met, because of this monthly data visualization column, Sue has recognized me as a kindred spirit. Her e-mail was a poignant expression of frustration over the fact that data visualization often gets a bad name because of how poorly it is understood and implemented. Bad data visualization implementations don't get used. They leave a sour taste in the mouth and shape the opinion that data visualization is all hype, all *eye candy* and no substance. Any technology, no matter how promising, is useless if people don't understand it. Data visualization, which is often misleadingly marketed as something that anyone can do without any special skills, is a vivid example of this.

Here's an excerpt of what Sue had to say:

I want clients to be better consumers. I want to educate them that they need to shop carefully for visualization support. They need to focus first on hiring a consultant or firm that is qualified in the area of data visualization. Right now, I think too often clients focus first on purchasing a tool with a known name or large market share. Good tools in unqualified hands result in bad visualizations and a bad name for all of us in the data visualization business.

Working with business data to extract meaningful information and then communicate that information effectively to others is something that many people can do, but not without the proper training. Data visualization tools and techniques provide the most powerful means to bring data to life, but learning to use visualizations for data analysis and communication is a bit like learning a new language with its own set of rules and nuance of expression.

I'd like to spend more and more of my time focusing less on the fundamentals and more on the cutting-edge data visualizations that await an audience of business professionals who are not impressed by superficial dazzle but can discern the immense value of data visualizations that really work. This will only happen when the business intelligence community wakes up to the importance of data visualization and begins to apply standards of excellence to it like those that have forged the dramatic improvements we've seen in most other data warehouse and business intelligence technologies and practices.

We've spent years focusing on the bits and bytes of data warehousing and business intelligence, but none of this matters if people can't understand and make good use of the data. It's now time to shift our focus to the skills, practices, processes and supporting tools needed to transform data into insight. Nothing is more central to this effort than data visualization. Whether the result of evolution or a gift from the divine, our visual connection to the world is like no other in its strength, speed, bandwidth and detail. Even the most

sophisticated statistics fall short in their ability to reveal meaning in data compared to what can be easily seen in a well designed visualization.

Figure 1 provides a quick peek at the type of visualization tool I'd like to see regularly used in business for data analysis. Like many of the exciting developments in data visualization, this tool, called TimeSearcher, results from the innovative efforts of the academic community—in this case, the Human-Computer Interaction Lab (HCIL) at the University of Maryland (<http://www.cs.umd.edu/hcil/>), where a great deal of interesting work is being done under the guidance of Ben Bederson and Ben Shneiderman.

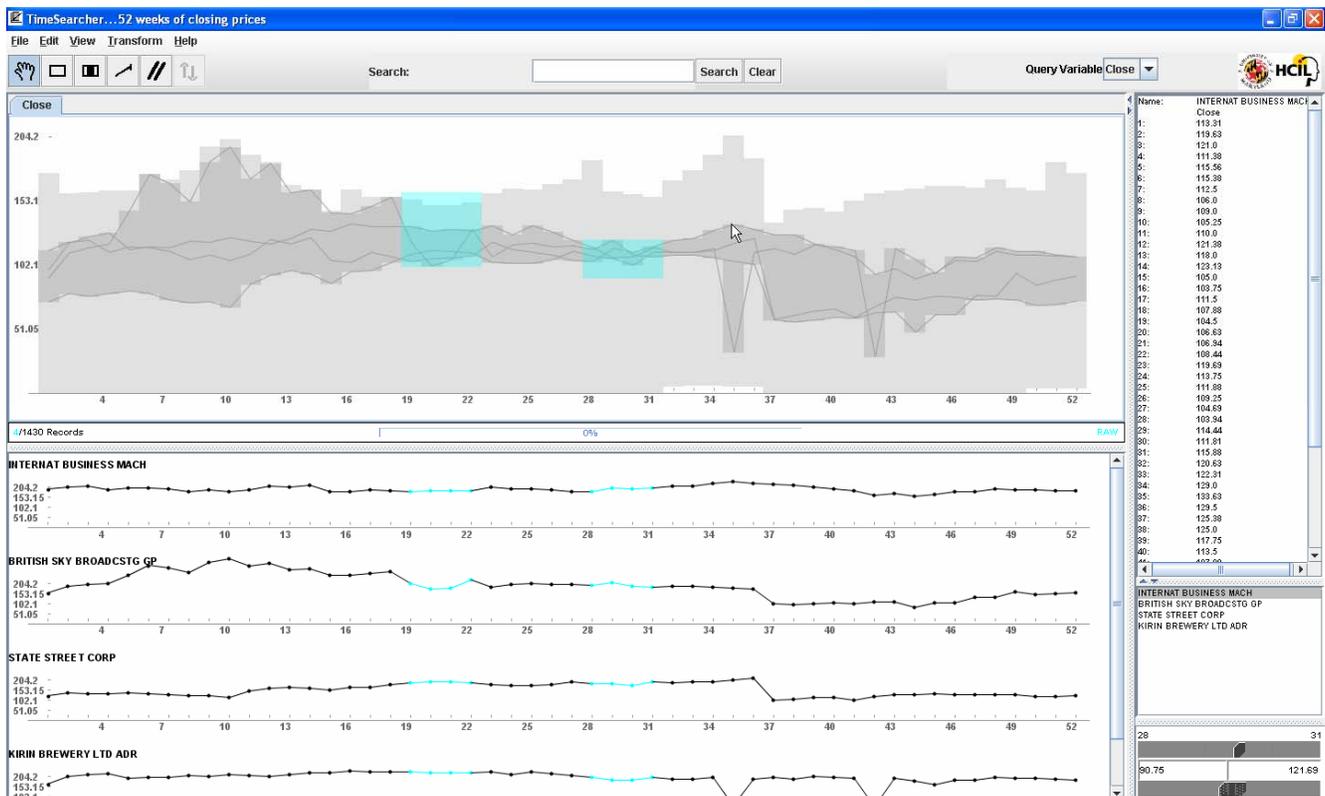


Figure 1: Screenshot of TimeSearcher, a tool for examining time-series data.

Expert data visualization recognizes the fact that different tasks and different types of data require different approaches and designs. The analysis of time-series data demands customized visualizations and modes of interaction. The example in Figure 1 presents 1,430 stock closing prices spread across 52 weeks, in four linked panels: 1) the entire data series in the upper left graph, 2) individual line graphs of the four selected stocks in the lower left, 3) details regarding the highlighted stock presented as tabular text in the upper right, and 4) a list of the four selected stocks in the lower right. The overview in the upper left allowed me to easily select subsets of interest within the context of the whole by drawing "timeboxes" around them (the aqua-colored rectangles), causing them to be highlighted and the line graphs in the lower left to be filtered to include only the selected stocks. What you see in this screenshot resulted from drawing two timeboxes to select only those stocks that closed between \$100 and \$160 during weeks 19 through 22 and then dropped to between \$90 and \$122 during weeks 28 through 31. The time periods that I selected in the upper overview were automatically reflected in the individual line graphs below through aqua color coding to match the timeboxes above. The full range of activity across all 52 weeks for the four

selected stocks is made visible by the dark gray background shading in the upper overview, thereby making the full historical context easy to see.

The way that TimeSearcher allows interaction with the data by directly manipulating the visualization, such as by drawing timeboxes to select subsets and thereby narrow focus, is as important as its effective visual design. Rather than formulating a query and then waiting for the resulting visualization to appear, tools such as TimeSearcher allow seamless interaction with the data, an intimate marriage of thought and action that directly extends cognition. The creators of TimeSearcher have given a great deal of thought to the questions people might ask when examining and trying to make sense of time-series data. For instance, they recognized that sometimes when we see a particular pattern in the data, such as a spike followed by a rapid and lasting decline, we might also want to examine other places where that same pattern occurs in an effort to discern a common cause. Consequently, TimeSearcher allows you to select a pattern directly in the visualization and ask it to automatically highlight all other instances of that pattern wherever they occur in the data. This is engaging visual functionality designed by people who have taken the time to understand our needs and address them in ways that match how we perceive and reason.

My purpose here is not to call attention to TimeSearcher in particular, but to illustrate the practical and powerful ways that data visualization can be applied to business data when the tools are well designed and you know what to look for in the data. It doesn't take an exceptional imagination to recognize the value of data visualizations like this. Having recognized this potential, the next logical step is to demand it.

(This article was originally published in *DM Review*.)

About the Author

Stephen Few has worked for over 20 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 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 two books: *Show Me the Numbers: Designing Tables and Graphs to Enlighten* and *Information Dashboard Design: The Effective Visual Communication of Data*. 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.