



Data Analysis At the Speed of Thought

Stephen Few

April 1, 2005

As a data visualization specialist, I'm a harsh critic of software vendors and rarely mention them except when citing examples of what *not* to do. However, I have encouraging news to share about a new, largely unknown vendor named Tableau Software. Its 1.0 release introduces an exciting new way to explore data and mine its meaning, shaping insight at nearly the speed of thought. An earlier version is already used at AOL, Google, Wells Fargo, Dow Chemical, Safeway and the U.S. Navy.

In the early 1980s, the electronic spreadsheet legitimized the PC as a viable business tool. Tableau could similarly help business intelligence to achieve its bright but seldom realized promise to deliver important insights. Tableau extends the familiar model of the tabular data display to a new visual form that lets our eyes detect what our minds could never discover in rows and columns of text and numbers. Imagine for a moment the power of a crosstab (a.k.a. pivot table) and then replace the rows and columns of numbers with rows and columns of graphs, each representing a distinct cross-section of data dimensions.

You're probably thinking: "But I can already query and massage my data and then turn the results into graphs." True, but something quite different and exciting is happening here. First, what you can produce today probably looks nothing like Tableau's multidimensional matrix of related graphs: the simplest and richest example that I've seen in commercial software of what data visualization guru Edward Tufte calls "small multiples"—a powerful vision that has largely been ignored by BI vendors. Second, you could experience a leap in productivity and insight because you'd no longer need to hone your queries; filter, group and sort your data; and then create a graph to visualize the results. Instead, you can watch each step of this process unfold as you go. The visualization isn't the result of the data access and manipulation process, it's part of the process itself. Data manipulation, vision and thought become one in a way that eases and expands the thinking process, delivering insight that would never result from a conventional approach.

Tableau is the best commercial incarnation I've ever seen of information visualization research involving quantitative data. It incorporates best practices derived from years of Stanford University research into the visualization and process flow. These approaches discourage the poor visualization practices that most software not only allows but often encourages.

The screen capture below shows a sample interface and the type of data visualizations that can be constructed quickly. The data I examined with Tableau represents one business's monthly sales of coffee and tea for 2003. It's displayed by regions across the columns and product types down the rows, sorted in descending order by sales. Two obvious facts pop out from the visualization: Best overall sales for a single region and product type occurred in the west for espresso products, and the worst occurred in the south, where no tea was sold. In the region with the highest single product sales—the east, leading with Columbian coffee—the other two coffee products sold extremely poorly by comparison. Further investigation might reveal that marketing efforts for the two poorly performing coffee products were unusually low in the east.

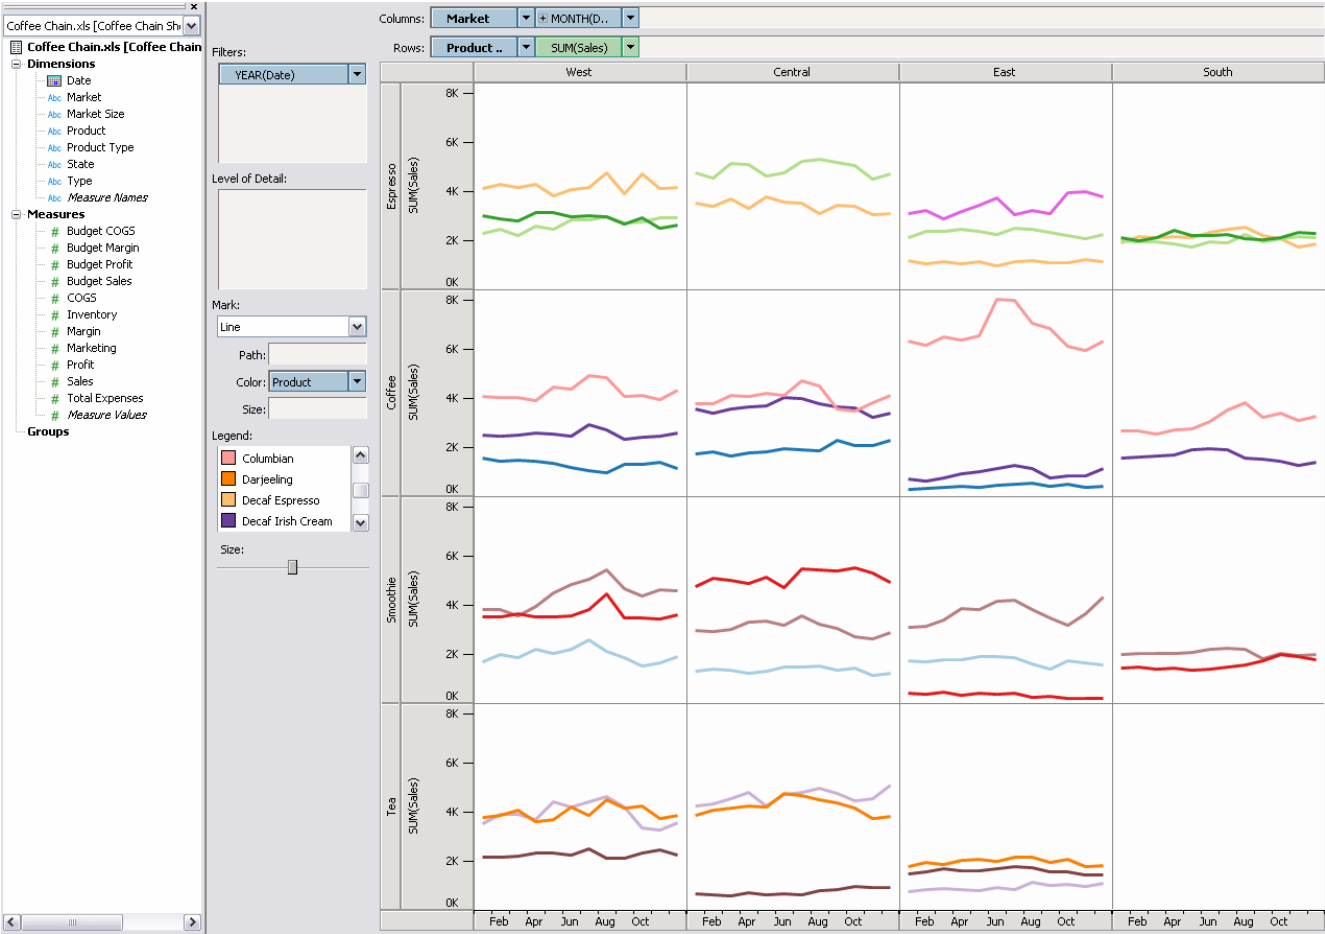


Figure 1: The grid of charts lets you quickly identify what's important in an otherwise overwhelming set of multidimensional data.

The next screen capture focuses on tea profits by region, which shows a loss for green tea in the west. With a simple drag-and-drop operation I was able to use the same bars that displayed profits to show sales as well, encoded as varying intensities of the color green. This display confirmed immediately that the loss wasn't because of especially low sales.

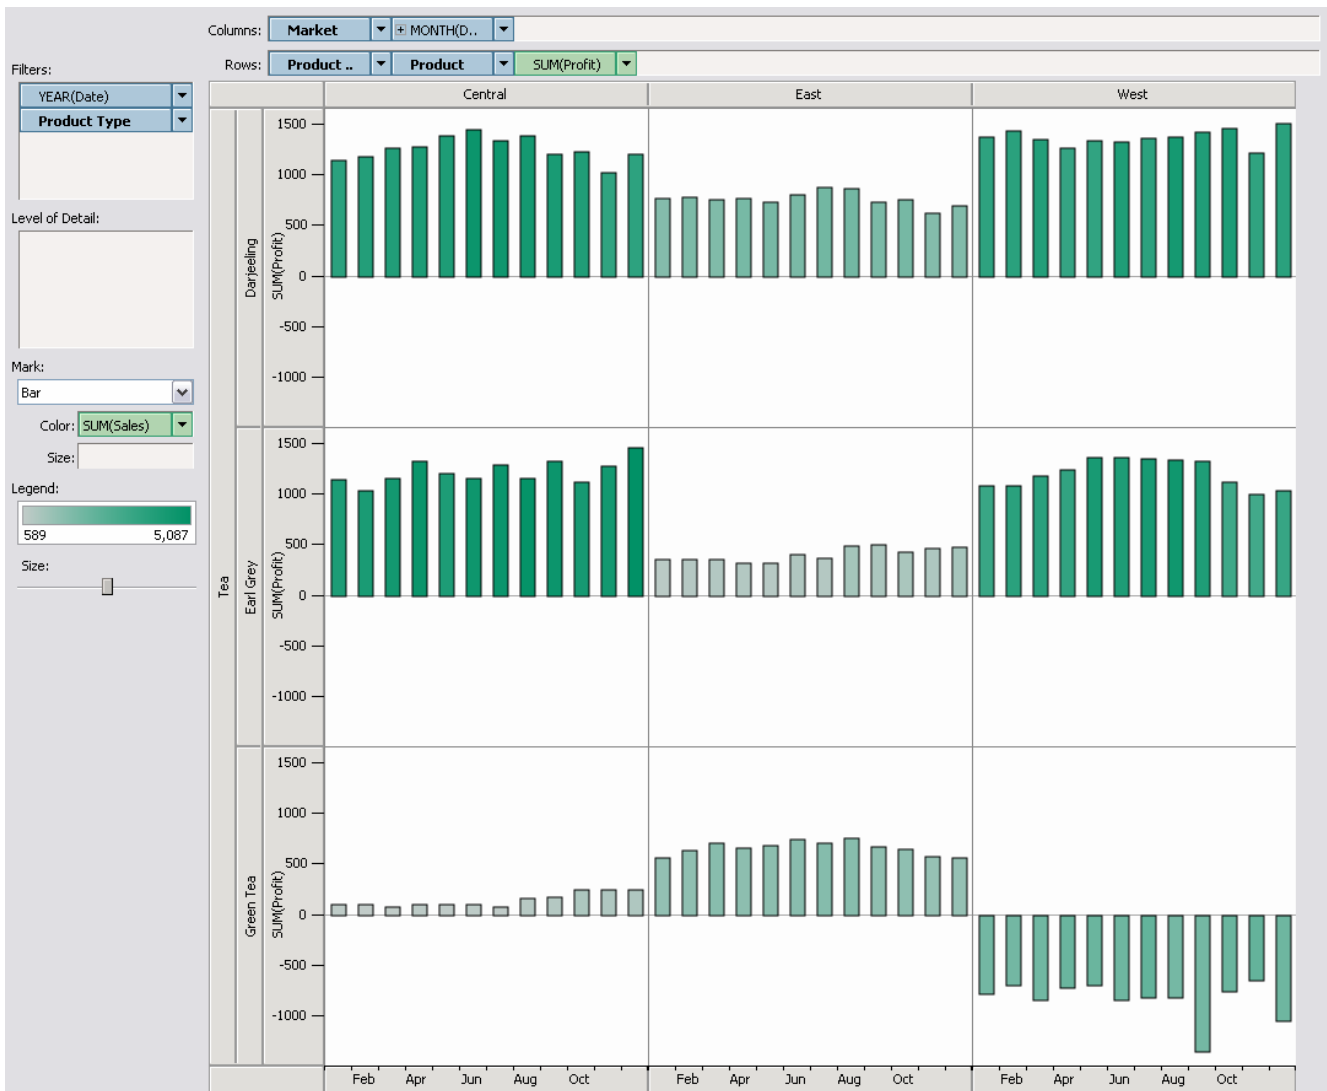


Figure 2: With one drag-and-drop operation, bars that showed profits changed to display sales as well.

Despite my enthusiasm for Tableau's promising beginning, this initial release certainly isn't complete or perfect. While its greatest strength is the unique way that it enables the analysis process through powerful visualizations and easy analytic navigation, its primary weakness is limited data access. The only enterprise-level relational database that it currently supports is Microsoft SQL Server, and even then the data must first be consolidated into a single denormalized table or view. If your data exists in Excel, a flat file or in one of two OLAP databases—Hyperion's Essbase or Microsoft's SQL Server Analysis Services—you can proceed immediately to analyze data. But if it resides in Oracle or DB2, you'll have to wait until at least version 1.5. Tableau says it plans to add support for direct access to star schemas in a future release.

Even in this initial release, Tableau has already incorporated many satisfying features that would be worth description if space allowed, such as:

- Unlimited undo and redo
- The ability to break hierarchies when useful for analysis
- The ability to define groups of data values on the fly, right in the visual display
- The ability to transform a measure into a dimension by automatically subdividing the full range of values into consecutive bins of equal size
- The automatic creation of hierarchy levels for date fields (year, quarter, month and so on)

If you're looking for software that allows the analytic process to flow, rather than shuffling along in its usual fits and starts, Tableau Software deserves your full attention.

Tableau Software 1.0 requires Microsoft Windows 2000 or later release. License ranges from \$999 to \$1,799, depending on edition. For more product information, go to www.tableausoftware.com.

(This article was originally published in *Intelligent Enterprise*.)

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](#).