



Advisor Solutions Invites You to a Double Wedding

Stephen Few

January 1, 2006

Advisor Solutions, one of the few business intelligence (BI) vendors that understand data visualization, has released a new version of its analysis software: Advisor Analyst 5.0. Advisor Solutions presides over two important marriages in its software: the union of effective visual analysis techniques developed by the academic community with commercial BI software; and the joining of data mining algorithms with data visualization. It's a double wedding that's worthy of notice.

Most of the really interesting products that incorporate data visualization have been developed by fairly small vendors. For years, a growing community of academic researchers has studied how interactive computer visualizations can be used to tap into human visual perception to make sense of data. Some of the findings are slowly being adopted in commercial software, and Advisor Solutions is among the vendors taking the lead.

Explore the Multiple Views

Two of the most powerful research-derived techniques featured in Advisor Analyst are the use of multiple concurrent views of the same or related data, which reveals relationships and patterns that might otherwise remain hidden, and a technique called *brushing*, which takes advantage of tightly coupled views of data by letting you highlight or filter data in one view and immediately see the effects of that action in all other views as well. The screen capture in Figure 1 illustrates a typical arrangement of complementary displays of a single data set.




Figure 1

This simple example shows the advantage of examining one set of data simultaneously from multiple perspectives. Patterns and relationships become visible that would be difficult to find in any other way. The bar graph on the upper left displays sales revenue by region. The stacked bar graph displays sales revenue by industry, segmented by region using the same colors per region that were used in the upper bar graph. The scatter plot to the right correlates sales revenue and profit margin with a data point for each customer, and it's also segmented by region using the same colors used in the other two graphs. Finally, a table at the bottom provides supporting details as text.

Now let's see how brushing can be used to radically expand the analytical landscape. To find out more about the customers with the lowest profit margins, I drew a rectangle around the low-profit-margin data points in the scatter plot. The results are shown in the second screen capture in which everything that wasn't selected has turned to gray.

Looking at the scatter plot alone, low-profit-margin customers seem to belong primarily to the West region (depicted by red dots), but this is deceiving because the largest percentage of revenue and probably customers also comes from the West. Looking at the bar graph on the upper left of the brushed screen, you can see that the revenue associated with low profit margins represents about the same percentage of each region's customers. But look now at the stacked bar graph below: It clearly shows that the health-care industry (fourth bar from the top) is responsible for the greatest number of low-profit-margin customers. The transportation industry (third bar from the bottom) has the highest percentage of its revenues belonging to the low-profit-margin group. When multiple perspectives are displayed concurrently and inquiry is enabled through the brushing technique, you're encouraged to explore the data and ask questions that you might have never previously imagined.

Understand the Influences

Advisor Analyst's promising new predictive analytics feature combines the potential for discovery offered by a good data-mining algorithm with the ability to see meaningful patterns in data using visualizations. You begin by selecting a field of data to serve as your target of interest—usually a quantitative measure, such as profit margin. The idea is to better understand this target by searching for the variables that seem to affect it the most. Once you've selected your target and specify the variables you want to test, you can let the data-mining algorithm do its work. The result is a list of the variables ranked in order of their influence on the target, including the percentage of each variable's contribution to the total effect. A bar graph further describes the influence of each variable by segmenting them into individual values or groups of values based on the degree to which they influence the target. For example, if you were examining the effect of marketing expenditures (one of many potential variables) on profit margin (your target), the full range of marketing campaign costs would be segmented into groups of similar effect (such as \$0-5,000, \$5,000-12,000 and so on) with a bar to represent each group. The direction of the bar (extending either right or left of zero) indicates the nature of the correlation to the target, either positive or negative, and the bar's length shows the strength of the correlation.




Figure 2

The interesting potential of this data-mining functionality resides in its marriage to data visualization. You can gain a great deal of insight by selecting particular results and seeing how they relate to various visual representations of the data. This data-mining functionality confirmed what I had already observed in the visualizations—that low profit margins seemed to be related primarily to particular industries—by selecting profit-margin percentage as a target and letting the algorithm search for patterns of influence on profits among a host of variables. The results revealed that the industries to which customers belong are responsible for a whopping 78 percent of the total influence on profit margin, and that health care contributed most to low profit margins. By selecting health care in the data-mining results, all health-care-related data was highlighted in the graphs and table in the multiple concurrent views.

You'll need some training to make good use of Advizor Analyst 5.0's data-mining algorithms. Unfortunately, Advizor Solutions' documentation reads a bit too much like a statistics textbook.

The Perfect Solution?

Does Advizor Solutions offer the perfect solution for data analysis? Not today, but it's definitely headed in the right direction and is already doing a number of important things right.

Advisor Analyst 5.0's entry-level pricing is \$6,000 for a single seat. For more product information or to contact Advizor Solutions, go to www.advizorsolutions.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](http://www.perceptualedge.com) of articles at www.perceptualedge.com. Between articles, you can read Stephen's thoughts on the industry in his [blog](#).