

Graph Designs for Rapidly Assessing Budget Performance

Stephen Few

August 8, 2006

This article begins a five-part series that will present the winning solutions to our 2006 Data Visualization Competition. Each article will feature one of the five business scenarios that participants were asked to respond to by developing a visual display to communicate a particular set of data for a particular purpose.

Here's the first scenario as it was described to participants:

You have been asked by the Budget Manager of a large corporation to develop a visualization that will enable her to examine expense budget performance during the current year (it is now mid-November) across 12 departments. She needs to see more than aggregate measures of each department's performance for the year. She needs to see trends and specific performance problems that have occurred during the year clearly enough to recognize situations that demand more detailed analysis. She is primarily interested in budget variances to identify potential problems in the budget itself as well as problems in the management of the budget. Your solution need not consist only of a single graph, but should not exceed what you can fit on a single page.

In addition to this description, participants were given a specific set of data in an Excel file.

While judging the many solutions that were submitted for this scenario, in addition to clarity of communication and ease of use, I was looking for the following key characteristics:

- Ease in spotting problems in departmental performance, both year-to-date and during particular months.
- Ease in seeing the trends, the ups and downs of budget performance through the year.
- Ease in comparing the performance of individual departments.
- Direct communication of variance to budget, expressed as degree of variance (that is, variance percent), rather than dollar variance. This is because a variance of \$10,000 for a department with a \$1 million budget is equal in degree to a \$1,000 variance for a department with a \$100,000 budget, but this is only obvious when expressed as variance percent.
- Clear indication that no data has yet been recorded for December (the current date is November 15) and only partial data has been recorded so far for November.

- All of this information is displayed on a single page or screen, within eye span for optimal ability to make comparisons.

The Winning Solution

Dylan Cotter of Spotfire submitted the winning solution for this scenario, which he created using Spotfire DXP, an exceptional tool for visual analysis.

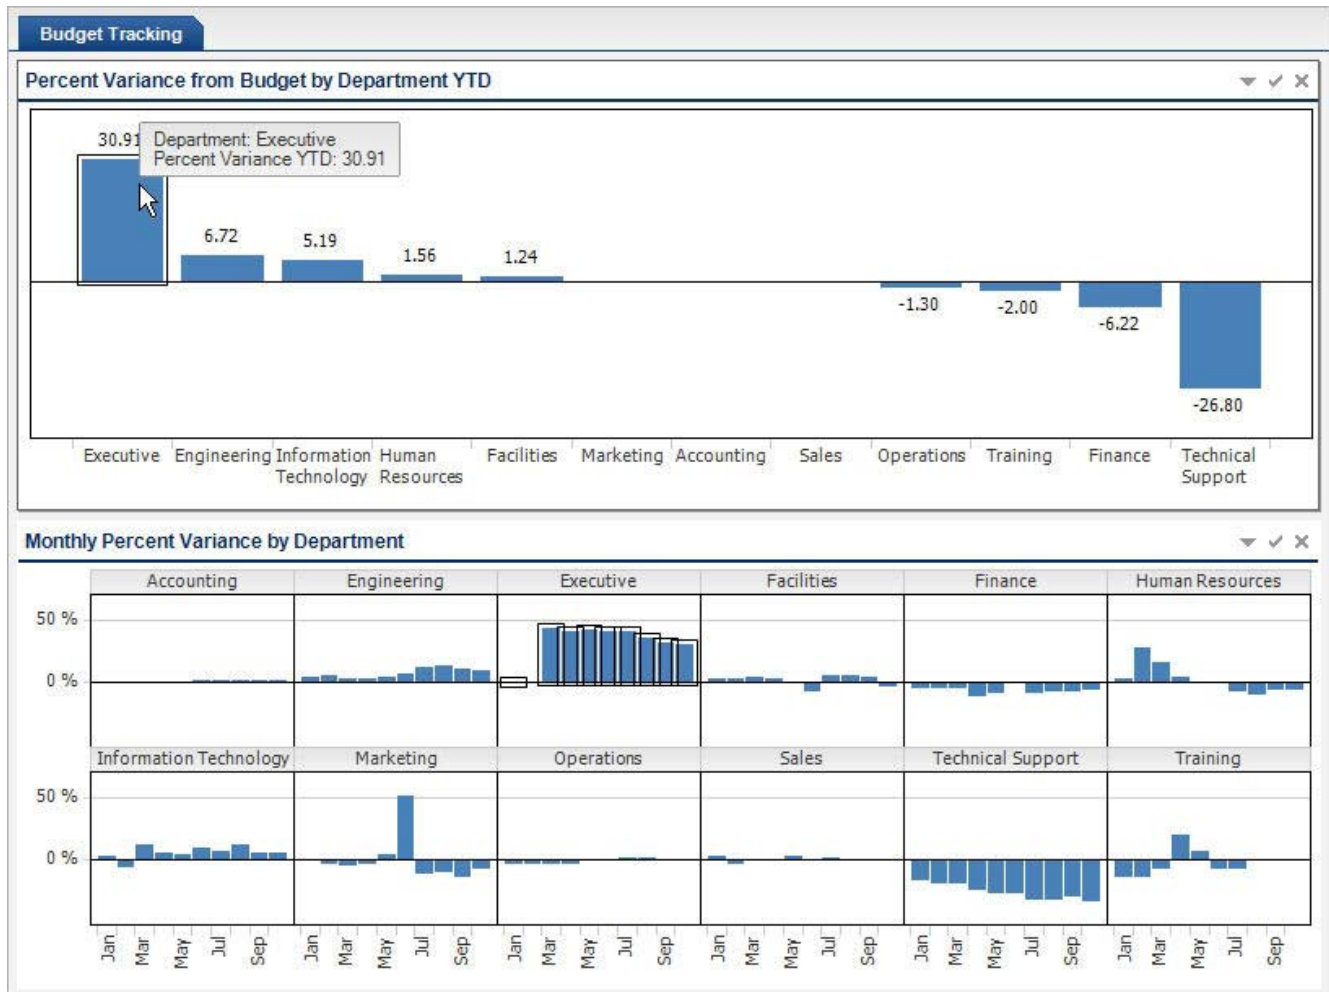


Figure 1: The Winning Solution for Scenario 1, created by Dylan Cotter of Spotfire.

Here are some of the highlights that made Dylan's solution stand out:

- Budget performance is directly expressed as a variance percent.
- It is easy to compare the overall performance of departments to one another using the ranked bar graph at the top.
- It is easy to examine the performance for a specific department month by month throughout the year.
- It is easy to visually link a selected department's overall performance in the upper graph to its monthly performance in the corresponding graph at the bottom simply

by pointing to a bar in the upper graph, which automatically highlights all associated bars in the corresponding small graph below.

- The fact that December data isn't included is clear from its omission in the small month-level graphs. The partial month of November has also been excluded. It could have been included but displayed differently than the other months (for example, by using a dimmer color of blue for that bar), but Dylan correctly reasoned that it would not make much sense and might create confusion to express the percentage variance to budget for only a half-month's worth of data, which would look like good performance even if there were a large under-budget variance.
- The small graphs at the bottom—one per department—can be compared to one another because they use a consistent quantitative scale.

In addition to learning from Dylan's wise design decisions, let's put his solution under the microscope to see if it could be improved. Given the purpose of this display, a few design changes could improve its ability to communicate.

- Even though the bars encode the monthly values in the bottom set of graphs, trends and change from one month to the next would be easier to see if the data were encoded as a line (that is, as line graphs).
- The link between the departments in the top graph and the corresponding small graph below would have been stronger if the small graphs were sequenced in the same order as the departments in the graph above.
- The arrangement of the small departmental graphs into two rows of six graphs each is not ideal for comparing the performance of various departments in a specific month. This is only a secondary need, but a rearrangement of the graphs would support it better without sacrificing other functionality. Arranging the graphs vertically, one above the other, would align the months from one graph to the next, making comparisons between departments for a particular month much easier. For instance, notice how much easier it is to compare the month of March for the Human Resources and Training departments, because these departments are positioned one on top of the other, than for departments that are not arranged in this manner. To arrange the small graphs vertically would require a change to the top graph to free up space, which could be accomplished by making it a horizontal bar graph positioned along the left edge of the display.
- It is always useful to state the "as of date" of the data on a report. Looking at this display, nothing currently tells us that we are looking at data that is complete through the end of October. People might be able to assume the information is current when they view a live version of this report on their computer screens, but what if they print and distribute it to others or file it for use at a later date? If you place the "as of date" of the data on the report, no one would ever need to guess.

Solutions that Fell Short

Several of the solutions exhibited particular problems that are worth noting so you'll know to avoid them. The first example (Figure 2) is actually quite good in most respects, but one aspect of its design in particular undermines its effectiveness. Time-series data ought to be displayed horizontally, from left to right. Notice how the series of small graphs in the bottom left section displays months from bottom to top, forming a pattern that is much harder to read when looking for trends. Placing the months on the horizontal X-axis and connecting the data points in each graph with a line would solve this problem. While we're improving this series of small graphs, I should also point out that, although it isn't obvious, these graphs are in order of departmental performance, from best (Technical Support) to worst (Executive), but the sequence runs counter-intuitively from right to left. Arranging them from left to right would work better, unless your audience reads from right to left, which is not the case here.

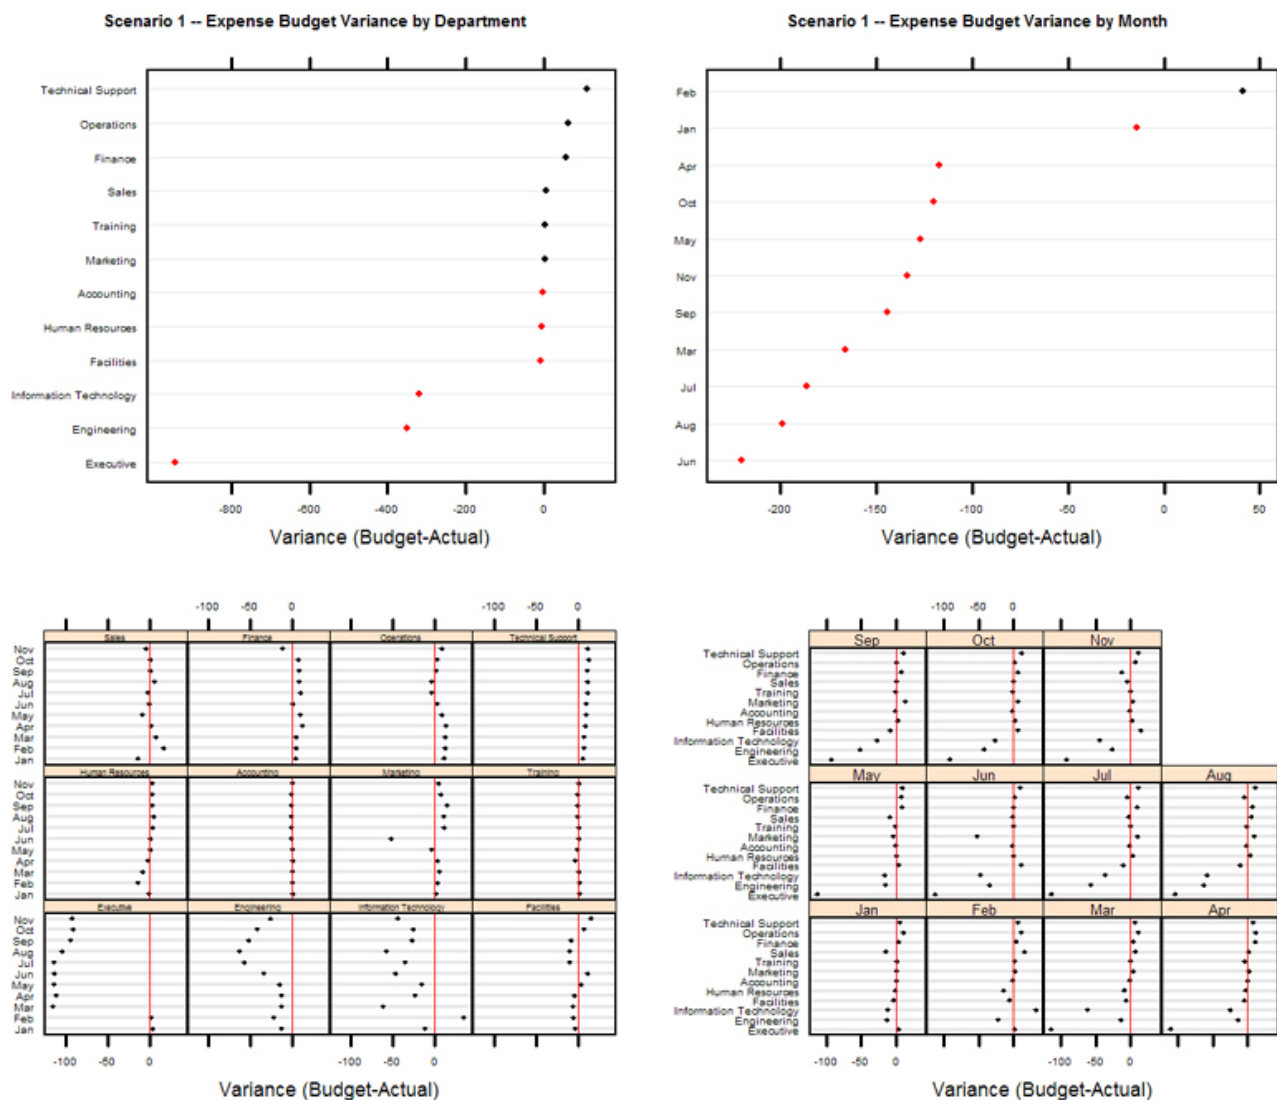


Figure 2

The next example (Figure 3) also has several good qualities, but a couple of problems hinder its effectiveness. I'm mostly concerned with the graph on the right. Three-dimensional graphs

are difficult to read. This is the only graph that displays monthly performance, but it is difficult to see trends and to interpret and compare the monthly values. In fact, something that ought to be obvious—the difference between positive and negative variances—is difficult to see. Which bars are extending upward from zero and which are heading downward into negative values? My other concern primarily involves the middle graph (“Adj Budget Variances per Month”). Nothing distinguishes the bar for November from the others, even though it represents only a half-month’s worth of data. This problem is minor compared to the month of December, however. The long bar extending downward for December suggests a large negative variance (that is, an extremely good value, well under budget), but there actually isn’t any data for the month of December. No bar whatsoever should appear for December.

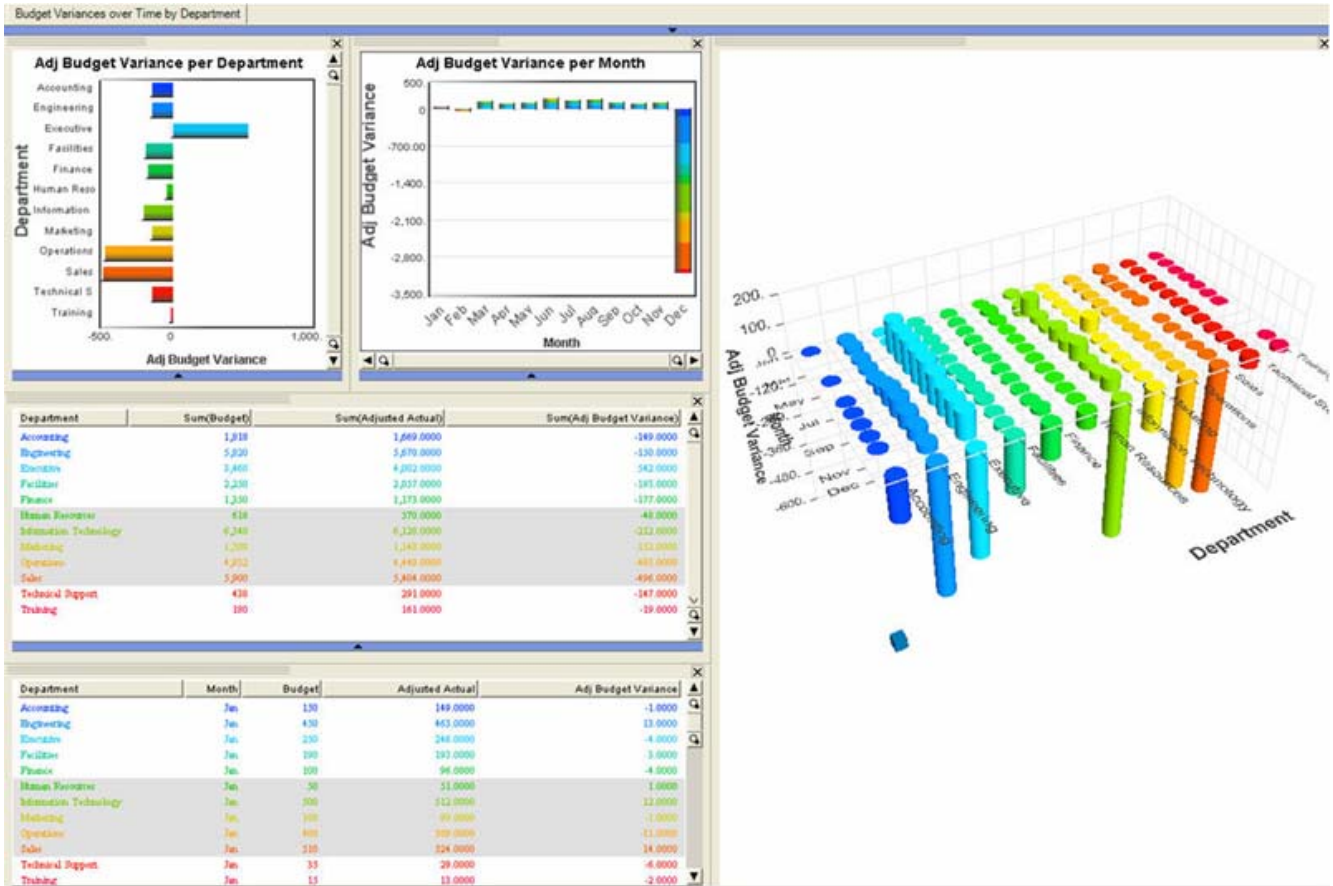


Figure 3

The next example (Figure 4) falls far short in effectiveness compared to those we’ve already seen. This is a simple table that has been enhanced slightly by color-coding individual cells to indicate “Within Budget” (green) “Up to 4% Over” (yellow), or “> 5% Over Budget” (red). The primary problem with this display is that it communicates too little. We can’t see trends through time, nor can we see the relative values of each month or change from month to month, except in terms of these imprecise performance bins. And, in which bin do the values greater than 4% through 5% fit? There is a gap between two of the bins. Finally, by using the colors red and green to reveal important information, 10% of men and 1% of women—those who are color blind—are kept in the dark.

Even if a table were an adequate solution, this design forces us to work harder than necessary by centering the numbers and varying their levels of precision (the number of

decimal digits), rather than making them easy to compare by displaying the same level of precision for every number (such as one decimal digit) and right aligning them, as numbers should always be aligned in tables.

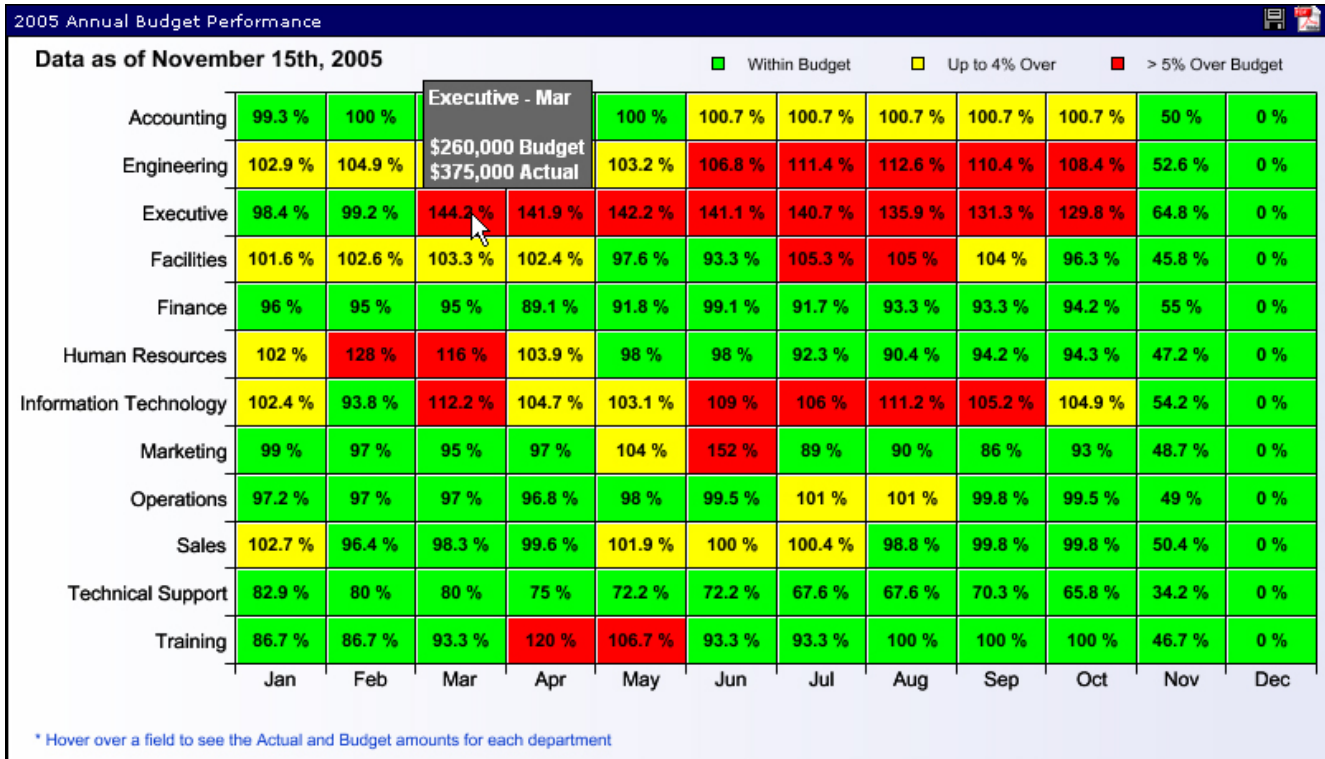


Figure 4

I've saved the example that is least effective for last (Figure 5). This approach is certainly creative, but not in a good way. Can you figure out how to read it? Certainly not without instructions, and even with instruction, this method of encoding the data would remain too difficult to decipher. Conventional solutions, such as simple line graphs in this case, work quite effectively and require no training. Innovation isn't necessary unless convention doesn't do the job. When innovation is required, the test of effectiveness is that, once trained, can people easily see what they need to see and easily understand what it means.




Figure 5

We can learn from one another's data visualization successes and failures. We all produce some of both. Next month, we'll examine solutions to a scenario that asks for a display of checking account transactions and the resulting balance interspersed through the course of a single month.

Thanks again to everyone who participated in the Business Intelligence Network's 2006 Data Visualization Competition. Through this unique form of collaboration, we can surely advance the state of effective data visualization!

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](#).