

3

THIRTEEN COMMON MISTAKES IN DASHBOARD DESIGN

Preoccupation with superficial and functionally distracting visual characteristics of dashboards has led to a rash of visual design problems that undermine their usefulness. Thirteen visual design problems are frequently found in dashboards, including in the examples featured as exemplary by software vendors.

Exceeding the boundaries of a single screen
Supplying inadequate context for the data
Displaying excessive detail or precision
Choosing a deficient measure
Choosing inappropriate display media
Introducing meaningless variety
Using poorly designed display media
Encoding quantitative data inaccurately
Arranging the data poorly
Highlighting important data ineffectively or not at all
Cluttering the display with useless decoration
Misusing or overusing color
Designing an unattractive visual display

The fundamental challenge of dashboard design is the need to squeeze a great deal of information into a small amount of space, resulting in a display that is easily and immediately understandable. If this doesn't sound challenging, either you are an expert designer with extensive dashboard experience, or you are basking in the glow of naiveté. Attempt the task, and you will find that dashboards pose a unique data visualization challenge. And don't assume that you can look to your software vendor for help—if they have the necessary design talent, they're doing a great job of hiding it.

Sadly, it is easy to find many examples of the mistakes you should avoid by looking no further than the web sites of the software vendors themselves. Let's use some of these examples to examine design that doesn't work and learn why it doesn't.

In almost every case, I've chosen to use actual examples from vendor web sites to illustrate dashboard design mistakes. In doing so, I am not saying that the software that produced the example is bad—I'm not commenting on the quality of the software one way or another. What I am saying is that the design practice is bad. This results primarily from vendors' lack of expertise in or inattention to visual design. These vendors should know better, but they've chosen to focus their energies on other aspects of their products, often highlighting glitzy visual features that actually undermine effective communication. I hope that seeing their work used to illustrate poor dashboard design will serve as a wake-up call to start paying attention to the features that really matter.

Exceeding the Boundaries of a Single Screen

My insistence that a dashboard should confine its display to a single screen, with no need for scrolling or switching between multiple screens, might seem arbitrary and a bit finicky, but it is based on solid and practical rationale. After studying data visualization for a while, including visual perception, one discovers that something powerful happens when things are seen together, all within eye span. Likewise, something critical is lost when you lose sight of some data by scrolling or switching to another screen to see other data. Part of the problem is that we can hold only a few chunks of information at a time in short-term memory. Relying on the mind's eye to remember information that is no longer visible is a rocky venture.

One of the great benefits of a dashboard as a medium of communication is the simultaneity of vision that it offers: the ability to see everything that you need at once. This enables comparisons that lead to insights—those “Aha!” experiences that might not occur in any other way. Clearly, exceeding the boundaries of a single screen negates this benefit. Let's examine the two versions of this problem—fragmenting data into separate screens and requiring scrolling—independently.

Fragmenting Data into Separate Screens

Information that appears on dashboards is often fragmented in one of two ways:

- Separated into discrete screens to which one must navigate
- Separated into different instances of a single screen that are accessed through some form of interaction

Enabling users to navigate to discrete screens or different instances of a single screen to access additional information is not, in general, a bad practice. Allowing navigation to further detail or to a different set of information that achieves its purpose best by standing alone can be a powerful dashboard feature. However, when all the information should be seen at the same time to gain the desired insights, that fragmentation undermines the unique advantages of a dashboard. Fragmenting data that should be seen together is a mistake.

Let's look at an example. The dashboard in Figure 3-1 fragments the data that executives need into 10 separate dashboards. This would be fine if the executives wouldn't benefit from seeing these various measures together, but that is hardly the case.




Figure 3-1. This dashboard fragments the data in a way that undermines the viewer's ability to see meaningful relationships.

In this example, a banking executive is forced to examine the performance of the following aspects of the business separately:

- Highlights
- Deposits
- Past due loans
- Profitability
- Growth
- Loans
- Risk
- Deposit mix
- Channels
- Market share

Each of these screens presents a separate, high-level snapshot of a single set of measures that ought to be integrated into a single screen. Despite what you might assume about the available screen labeled "Highlights," it does not provide a consolidated visual overview of the data but consists primarily of a text table that contains several of the measures. A banking executive needs to see these measures together in a way that enables comparisons to understand how they relate to and influence one another.

Splitting the big picture into a series of separate small pictures is a mistake whenever seeing the big picture is worthwhile.

A similar example, from the same software vendor, is shown in Figure 3-2. This time the picture of daily sales has been split into a separate dashboard for each of 20 products. If the intention is to serve the needs of product managers who are each exclusively interested in a single product and never want to compare sales of that product to others, this design doesn't fragment the data in a harmful way. If, however, any benefit can be gained by viewing the sales of multiple products together, which is almost surely the case, this design fails.




Figure 3-2. This dashboard requires viewers to click on a desired product and view information for only one product at a time.

Requiring Scrolling

The dashboard in Figure 3-3 illustrates the problem that's created when scrolling is required to see all the data. Not only are we left wondering what lies below the bottom of the screen in the dashboard as a whole, but we're also given immediate visual access only to the first of many metrics that appear in the scrollable box at the top right, beginning with "No. Transactions." We'd be better off reading a printed report extending across multiple pages, because at least then we could lay out all of the pages at once for simultaneous viewing. People commonly assume that anything that lies beyond their immediate field of vision and requires scrolling to see

is of less importance than what's immediately visible. Many viewers won't bother to look at what lies off the screen, and those who take the time will likely resent the effort.





Figure 3-3. This dashboard demonstrates the effectiveness that is sacrificed when scrolling is required to see all the information.

Supplying Inadequate Context for the Data

Measures of what's currently going on in the business rarely do well as a solo act; they need a good supporting cast to succeed. For example, to state that quarter-to-date sales total \$736,502 without any context means little. Compared to what? Is this good or bad? How good or bad? Are we on track? Are we doing better than we have in the past, or worse than we've forecasted? Supplying the right context for key measures makes the difference between numbers that just sit there on the screen and those that enlighten and inspire action.

The gauges in Figure 3-4 could easily have incorporated useful context, but they fall short of their potential. For instance, the center gauge tells us only that 7,822 units have sold this year to date, and that this number is good (indicated by the green arrow). A quantitative scale on a graph, such

as the radial scales of tick marks on these gauges, is meant to provide an approximation of the measure, but it can only do so if the scale is labeled with numbers, which these gauges lack. If the numbers had been present, the positions of the arrows might have been meaningful, but here the presence of the tick marks along a radial axis suggests useful information that hasn't actually been included.



These gauges use up a great deal of space to tell us nothing whatsoever. The same information could have been communicated simply as text in much less space, without any loss of meaning:

YTD Units	7,822
October Units	869
Returns Rate	0.26%

Another failure of these gauges is that they tease us by coloring the arrows to indicate good or bad performance, without telling us *how* good or bad it is. They could easily have done this by labeling the quantitative scales and visually encoding sections along the scales as good or bad, rather than just encoding the arrows in this manner. Had this been done, we would be able to see at a glance how good or bad a measure is by how far the arrow points into the good or bad ranges.

The gauge that appears in Figure 3-5 does a better job of incorporating context in the form of meaningful comparisons. Here, the potential of the graphical display is more fully realized. The gauge measures the average duration of phone calls and is part of a larger dashboard of call-center data.

Supplying context for measures need not always involve a choice of the single best comparison—rather, several contexts may be given. For instance, quarter-to-date sales of \$736,502 might benefit from comparisons to the budget target of \$1,000,000; sales on this day last year of \$856,923; and a time-series of sales figures for the last six quarters. Such a display would provide much richer insight than a simple display of the

current sales figure, with or without an indication of whether it's "good" or "bad." You must be careful, however, when incorporating rich context such as this to do so in a way that doesn't force the viewer to get bogged down in reading the details to get the basic message. It is useful to provide a visually prominent display of the primary information and to subdue the supporting context somewhat, so that it doesn't get in the way when the dashboard is being quickly scanned for key points.



Figure 3-5. This dashboard gauge (found in a paper entitled "Making Dashboards Actionable," written by Laurie M. Orlov and published in December 2003 by Forrester Research, Inc.) does a better job than those in Figure 3-4 of using a gauge effectively.

I believe that the circular shape used by gauges like this one wastes valuable space on a dashboard, as I'll explain in Chapter 6, *Effective Dashboard Display Media*. Nevertheless, I commend this gauge for displaying richer information than most.

The amount of context that ought to be incorporated to enrich the measures on a dashboard depends on its purpose and the needs of its viewers. More is not always better, but when more provides real value, it ought to be included in a way that supports both a quick overview without distraction as well as contextual information for richer understanding.

Displaying Excessive Detail or Precision

Dashboards almost always require fairly high-level information to support the viewer's need for a quick overview. Too much detail, or measures that are expressed too precisely (for example, \$3,848,305.93 rather than \$3,848,305, or perhaps even \$3.8M), just slow viewers down without providing them any benefit. In a way, this problem is the opposite extreme of the one we examined in the previous section—too much information rather than too little.

The dashboard in Figure 3-6 illustrates this type of excess. Examine the two sections that I've enclosed in red rectangles. The lower-right section

displays from 4 to 10 decimal digits for each measure, which might be useful in some contexts, but doubtfully in a dashboard. The highlighted section above displays time down to the level of seconds, which also seems like overkill in this context. With a dashboard, every unnecessary piece of information results in time wasted trying to filter out what's important, which is intolerable when time is of the essence.




Figure 3-6. This dashboard shows unnecessary detail, such as times expressed to the second and measures expressed to 10 decimal places.

Choosing a Deficient Measure

For a measure to be meaningful, we must know what is being measured and the units in which the measure is being expressed. A measure is deficient if it isn't the one that most clearly and efficiently communicates the meaning that the dashboard viewer should discern. It can be accurate, yet

not the best choice for the intended message. For example, if the dashboard viewer only needs to know to what degree actual revenue differs from budgeted revenue, it would be more direct to simply express the variance as -9% (and perhaps display the variance of $-\$8,066$ as well) rather than displaying the actual revenue amount of $\$76,934$ and the budgeted revenue amount of $\$85,000$ and leaving it to the viewer to calculate the difference. In this case, a percentage clearly focuses attention on the variance in a manner that is directly intelligible.

Figure 3-7 illustrates this point. While this graph displays actual and budgeted revenues separately, its purpose is to communicate the variance of actual revenues from the budget.




Figure 3-7. This graph illustrates the use of measures that fail to directly express the intended message.

The variance, however, could have been displayed more vividly by encoding budgeted revenue as a reference line of 0% and the variance as a line that meanders above and below budget (expressed in units of positive and negative percentages, as shown on the next page in Figure 3-8). The point here is to always think carefully about the message that most directly supports the viewer's needs, and then select the measure that most directly supports that message.




Figure 3-8. This graph is designed to emphasize deviation from a target, which it accomplishes in part by expressing the difference between budgeted and actual revenues using percentages.

Choosing Inappropriate Display Media

Choosing inappropriate display media is one of the most common design mistakes made, not just in dashboards, but in all forms of quantitative data presentation. For instance, using a graph when a table of numbers would work better, and vice versa, is a frequent mistake. Allow me to illustrate using several examples beginning with the pie chart in Figure 3-9.




Figure 3-9. This chart illustrates a common problem with pie charts.

This pie chart is part of a dashboard that displays breast cancer statistics. Look at it for a moment and see if anything seems odd.

Pie charts are designed specifically to present parts of a whole, and the whole should always add up to 100%. Here, the slice labeled “Breast 13.30%” looks like it represents around 40% of the pie—a far cry from 13.3%. Despite the meaning that a pie chart suggests, these slices are not parts of a whole; they represent the probability that a woman will develop a particular form of cancer (breast, lung, colon, and six types that aren’t labeled). This misuse of a pie chart invites confusion.

The truth is, I never recommend the use of pie charts. The only thing they have going for them is the fact that everybody immediately knows when they see a pie chart that they are seeing parts of a whole (or ought to be). Beyond that, pie charts don’t display quantitative data very effectively. As you’ll see in Chapter 4, *Tapping into the Power of Visual Perception*, humans can’t compare two-dimensional areas or angles very accurately—and these are the two means that pie charts use to encode quantitative data. Bar graphs are a much better way to display this information.

The pie chart in Figure 3-10 shows that even when correctly used to present parts of a whole, these graphs don’t work very well. Without the value labels, you would only be able to discern that opportunities rated as “Fair” represent the largest group, those rated as “Field Sales: 2-Very High” represent a miniscule group, and the other ratings groups are roughly equal in size.

Refer to my book *Show Me the Numbers: Designing Tables and Graphs to Enlighten* (Oakland, CA: Analytics Press, 2004) for a thorough treatment of the types of graphs that work best for the most common quantitative messages communicated in business.



Figure 3-10. This example shows that even when they are used correctly to present parts of a whole, pie charts are difficult to interpret accurately.

Figure 3-11 displays the same data as Figure 3-10, this time using a horizontal bar graph that can be interpreted much more efficiently and accurately.




Figure 3-11. This horizontal bar graph does a much better job of displaying part-to-whole data than the preceding pie charts.

Other types of graphs can be equally ineffective. For example, the graph in Figure 3-12 shows little regard for the viewer's time and no understanding of visual perception. This graph compares revenue to operating costs across five months, using the size of overlapping circles (sometimes called bubbles) to encode the quantities. Just as with the slices of a pie, using circles to encode quantity relies on the viewer's ability to compare two-dimensional areas, which we simply cannot accurately do. Take the values for the month of February as an example. Assuming that operating costs equal \$10,000, what is the revenue value?






Figure 3-12. This graph uses the two-dimensional area of circles to encode their values, which needlessly obscures the data.

Our natural tendency is to compare the sizes of the two circles using a single dimension—length or width—equal to the diameter of each, which suggests that revenue is about three times that of operating costs, or about \$30,000. This conclusion is wrong, however, to a huge degree. The two-dimensional area of the revenue circle is actually about nine times bigger than that of the operating costs circle, resulting in a value of \$90,000. Oops! Not even close.

Now compare operating costs for the months of February and May. It appears that costs in May are greater than those in February, right? In fact, the interior circles are the same size—measure them and see. The revenue bubble in May is smaller than the one in February, which makes the enclosed operating costs bubble in May seem bigger, but this is an optical illusion. As you can see, the use of a bubble chart for this financial data was a poor choice. A simple bar graph like the one in Figure 3-13 works much better.



Actual versus budgeted revenue is also the subject of Figure 3-14, but this time it's subdivided into geographical regions rather than time slices and displayed as a radar graph. The quantitative scale on a radar graph is laid along each of the axis lines that extend from the center to the perimeter, like radius lines of a circle. The smallest values are those with the shortest distance between the center point and the perimeter.



The lack of labeled axes in this graph limits its meaning, but the choice of a radar graph to display this information in the first place is an even more fundamental error. Once again, a simple bar graph like the one in Figure 3-15 would communicate this data much more effectively. Radar graphs are rarely appropriate media for displaying business data. Their circular shape obscures data that would be quite clear in a linear display such as a bar graph.




Figure 3-15. This bar graph effectively compares actual to budgeted revenue data.

The last example that I'll use to illustrate my point about choosing inappropriate means of display appears in Figure 3-16.




Figure 3-16. This display uselessly encodes quantitative values on a map of the United States.

There are times when it is very useful to arrange data spatially, such as in the form of a map or the floor plan of a building, but this isn't one of them. We don't derive any insight by laying out revenue information—in this case, whether revenues are good (green light), mediocre (yellow light), or poor (red light), in the geographical regions South (brown), Central (orange), West (tan), and East (blue)—on a map.

If the graphical display were presenting meaningful geographical relationships—say, for shipments of wine from California, to indicate where special taxes must be paid whenever deliveries cross state lines—perhaps a geographical display would provide some insight. With this simple set of four regions with no particular factors attached to geographical location, however, the use of a map simply takes up a lot of space to say no more than we find in the table that appears on this same dashboard, which is shown in Figure 3-17.

	Actual	Budget
East	\$87,398	\$78,950
West	\$132,931	\$119,850
South	\$50,846	\$49,100
Central	\$129,680	\$125,180

Figure 3-17. This table, from the same dashboard, provides a more appropriate display of the regional revenue data that appears in Figure 3-16.

Introducing Meaningless Variety

The mistake of introducing meaningless variety into a dashboard design is closely tied to the one we just examined. I've found that people often hesitate to use the same type of display medium multiple times on a dashboard, out of what I assume is a sense that viewers will be bored by the sameness. Variety might be the spice of life, but if it is introduced on a dashboard for its own sake, the display suffers. You should always select the means of display that works best, even if that results in a dashboard that is filled with nothing but multiple instances of the same type of graph. If you are giving viewers the information that they desperately need to do their jobs, the data won't bore them just because it's all displayed in the same way. They will definitely get aggravated, however, if forced to work harder than necessary to get the information they need due to arbitrary variety in the display media. In fact, wherever appropriate, consistency in the means of display allows viewers to use the same perceptual strategy for interpreting the data, which saves time and energy.

Figure 3-18 illustrates variety gone amok. This visual jumble requires a shift in perceptual strategy for each display item on the dashboard, which means extra time and effort on the user's part.




Figure 3-18. This dashboard exhibits an unnecessary variety of display media.

Using Poorly Designed Display Media

It isn't enough to choose the right medium to display the data and its message—you also must design the components of that medium to communicate clearly and efficiently, without distraction. Most graphs used in business today are poorly designed. The reason is simple: almost no one has been trained in the fundamental principles and practices of effective graph design. This content is thoroughly covered in my book *Show Me the Numbers: Designing Tables and Graphs to Enlighten*, so I won't repeat myself here. Instead, I'll simply illustrate the problem with a few examples.

In addition to the fact that a bar graph would have been a better choice to display this data (the division of revenue between six sales), Figure 3-19 exhibits several design problems. Look at it for a moment and see if you can identify aspects of its design that inhibit quick and easy interpretation.




Figure 3-19. This pie chart illustrates several design problems.

Here are the primary problems that I see:

A legend was used to label and assign values to the slices of the pie. This forces our eyes to bounce back and forth between the graph and the legend to glean meaning, which is a waste of time and effort when the slices could have been labeled directly.

The order of the slices and the corresponding labels appears random. Ordering them by size would have provided useful information that could have been assimilated instantly.

The bright colors of the pie slices produce sensory overkill. Bright colors ought to be reserved for specific data that should stand out from the rest.

The pie chart in Figure 3-20 also illustrates a problem with color choice.




Figure 3-20. This pie chart uses colors for the slices that are too much alike to be clearly distinguished.

In this case, the 11 colors that were chosen are too similar. It is difficult to determine which of the hues along the yellow through orange to red spectrum in the legend corresponds to each slice of the pie. This kind of eye-straining exercise is deadly, especially on a dashboard.

Another example of an ineffective display medium is shown in Figure 3-21. These meters are an attempt to be true to the metaphor of a car dashboard. Notice that the numbers look just like they would on an odometer: they lack the commas normally used to delineate every set of three digits to help us distinguish thousands from millions, and so on. In a misguided effort to make these meters look realistic, their developers made the numbers harder to read—engineers designed these meters, not business people. Notice also that numbers along the quantitative scale are positioned inside rather than outside the axis, which will cause them to be obscured by the needle when it points directly to them, and that the positioning of the text at the bottom of each meter (for example, “4382934 Amount Sold” on the “Internet Revenue” meter) obstructs the needle for measures near the bottom or top of the scale.




Figure 3-21. These dashboard meters have definitely taken the dashboard metaphor too far.

In the last section, I spoke of bar graphs as a preferable alternative to certain other display media. However, while bar graphs can do an excellent job of displaying quantitative data, they can be misused as well. Examine the graph in Figure 3-22, and take a moment to list any problems with its design that you see. Write down your observations below before reading on, if you'd like.




Figure 3-22. This bar graph, found on a dashboard, exhibits several design problems.

You might have noticed that the grid lines on the graph (not to mention the background pattern of colored rectangles) do nothing but distract from the data. Grid lines such as these, especially when visually prominent, make it more difficult to see the shape of the data. Perhaps you also noticed that the 3-D effect of the graph not only added no value, but also made the values encoded by the bars harder to interpret. Anything else? Well, this graph illustrates a common problem with color. Why is each of the bars a different color? The colors aren't being used to identify the bars, as each one has a label to its left. Differences in the color of data-encoding objects should always be meaningful; otherwise, they needlessly grab our attention and cause us to search for meaning that isn't there.

The distinct colors of the bars in Figure 3-23 do, thankfully, carry meaning, but here the colors are distractingly bright and the 3-D effect makes them hard to read.




Figure 3-23. This bar graph, found on a dashboard, was poorly designed in a number of ways.

However, this isn't the problem that I most want you to notice. The purpose of the graph is to compare actual to budgeted revenues for each of the four regions, but something about its design makes this difficult. Can you see the problem? Given its purpose, the bars for actual and budgeted revenues for each region should have been placed next to one another. As they are, it is unnecessarily difficult to compare them. Simple design mistakes like this can significantly undermine the success of a dashboard.

Several of the examples that we've examined have been rendered in 3D, even though the third dimension of depth doesn't encode any meaning. Even when the third dimension is used to encode a separate variable, however, it still poses a problem. The graph in Figure 3-24 uses the third dimension of depth to represent time (the four quarters of the year 2001). The problem in this case isn't that the third dimension is meaningless, but rather that you can't read everything on the chart. This is caused by *occlusion*. Adding the dimension of depth causes some of the bars to be hidden behind—or occluded by—others. For instance, what were fax revenues for Quarter 3? You can't tell because the bar is completely hidden. Whether the third dimension is used to encode data or not, you should almost always avoid 3-D graphs. Exceptions to this rule are rare when displaying typical business data.





Figure 3-24. This 3-D bar graph illustrates the problem of occlusion.

Encoding Quantitative Data Inaccurately

Sometimes graphical representations of quantitative data are mistakenly designed in ways that display inaccurate values. In Figure 3-25, for instance, the quantitative scale along the vertical axis was improperly set

for a graph that encodes data in the form of bars. The length of a bar represents its quantitative value. The bars in this graph that represent revenue and costs for the month of January suggest that revenue was about four times costs. An examination of the scale, however, reveals the error of this natural assumption: the revenue is actually less than double the costs. The problem is that the values begin at \$500,000 rather than \$0, as they always should in a bar graph.



Arranging the Data Poorly

Dashboards often need to present a large amount of information in a limited amount of space. If the information isn't organized well, with appropriate placement of information based on importance and desired viewing sequence, along with a visual design that segregates data into meaningful groups without fragmenting it into a confusing labyrinth, the result is a cluttered mess. Most examples of dashboards found on the Web are composed of a small amount of data to avoid the need for skilled visual design, but they still often manage to look cluttered and thrown together. The goal is not simply to make the dashboard look good, but to arrange the data in a manner that fits the way it's used. The most important data ought to be prominent. Data that require immediate attention ought to stand out. Data that should be compared ought to be arranged and visually designed to encourage comparisons.

The dashboard in Figure 3-26 illustrates some of the problems often associated with poor arrangement of data. Notice first of all that the most prominent position on this dashboard—the top left—is used to display the vendor's logo and navigational controls. What a waste of prime real estate! As you scan down the screen, the next information that you see is a gauge that

presents the average order size. It's possible that average order size might be someone's primary interest, but it's unlikely that, of all the information that appears on this dashboard, this is the most important. As I'll discuss in Chapter 5, *Eloquence Through Simplicity*, the least prominent real estate on the screen is the lower-right corner. However, in this example the large amount of space taken up by the graphs that present "Computers Returns Across Models," as well as the larger font sizes used in this section, tends to draw attention to data that seems tangential to the rest. This dashboard lacks an appropriate visual sequence and balance based on the nature and importance of the data. Notice also that the bright red bands of color above each section of the display, where the titles appear in white, are far more eye-catching than is necessary to declare the meanings of the individual displays. This visually segments the space to an unnecessary degree. Lastly, note that the similarity of the line graphs that display order size and profit trends invites our eyes to compare them. This is probably a useful comparison, but the positional separation and side-by-side rather than over-under arrangement of the two graphs makes close comparison difficult. As this example illustrates, you can't just throw information onto the screen wherever you can make it fit and expect the dashboard to do its job effectively.




Figure 3-26. This dashboard exemplifies poorly arranged data.

Highlighting Important Data Ineffectively or Not at All

When you look at a dashboard, your eyes should immediately be drawn to the information that is most important, even when it does not reside in the most visually prominent areas of the screen. In Chapter 5, *Eloquence Through Simplicity*, we'll examine several visual techniques that can be used to achieve this end. For now, we'll look at what happens when this isn't done at all, or isn't done well.

The problem with the dashboard in Figure 3-27 is that *everything* is visually prominent, and consequently nothing stands out. The logo and navigation controls (the buttons on the left) are prominent both as a result of their placement on the screen and the use of strong borders, but these aren't data and therefore shouldn't be emphasized. Then there are the graphs where the data reside: all the data are equally bold and colorful, leaving us with a wash of sameness and no clue where to focus. Everything that deserves space on a dashboard is important, but not equally so—the viewer's eyes should always be directed to the most crucial information first.




Figure 3-27. This dashboard fails to differentiate data by its importance, giving relatively equal prominence to everything on the screen.

Cluttering the Display with Useless Decoration

Another common problem on the dashboards that I find on vendor web sites is the abundance of useless decoration. They either hope that we will be drawn in by the artistry or assume that the decorative flourishes are necessary to entertain us. I assure you, however, that even people who enjoy the decoration upon first sight will grow weary of it in a few days.

The makers of the dashboard in Figure 3-28 did an exceptional job of making it look like an electronic control panel. If the purpose were to train people in the use of some real equipment by means of a simulation, this would be great, but that isn't the purpose of a dashboard. The graphics dedicated to this end are pure decoration, visual content that the viewer must process to get to the data.




Figure 3-28. This dashboard is trying to look like something that it is not, resulting in useless and distracting decoration.

I suspect that the dashboard in Figure 3-29 looked too plain to its designer, so she decided to make it look like a page in a spiral-bound book—cute, but a distracting waste of space.



Figure 3-29. This dashboard is another example of useless decoration—the designer tried to make the dashboard look like a page in a spiral-bound notebook.

Likewise, I'd guess that the designer of the dashboard in Figure 3-30—after creating a map, a bar graph, and a table that all display the same data—decided that he had to fill up the remaining space, so he went wild with an explosion of blue and gray circles. Blank space is better than meaningless decoration. Can you imagine yourself looking at this every day?




Figure 3-30. This dashboard is a vivid example of distracting ornamentation.

The last example, Figure 3-31, includes several elements of decoration that ought to be eliminated. To begin with, a visually ornate logo and title use

up the most valuable real estate across the entire top of the dashboard. If a logo must be included for branding purposes, make it small and visually subtle, and place it somewhere out of the way. The background colors of gold and blue certainly draw our eyes to the data, but they do so in an unnecessarily heavy-handed manner. Also, the color gradients from dark to light provide visual interest that supports no real purpose and is therefore distracting. Lastly, the maps in the background of the three upper graphs, though visually muted, still distract from the data itself.




Figure 3-31. This dashboard exhibits several examples of dysfunctional decoration.

As data visualization expert Edward Tufte observes:

Edward R. Tufte, *The Visual Display of Quantitative Information* (Cheshire, CT: Graphics Press, 1983), 80.

Inept graphics also flourish because many graphic artists believe that statistics are boring and tedious. It then follows that decorated graphics must pep up, animate, and all too often exaggerate what evidence there is in the data... If the statistics are boring, then you've got the wrong numbers.

Misusing or Overusing Color

We've already seen several examples of misused or overused color. The remaining point that I want to emphasize here is that color should not be used haphazardly.

Color choices should be made thoughtfully, with an understanding of how we perceive color and the significance of color differences. Some colors are hot and demand our attention, while others are cooler and less visible. When any color appears as a contrast relative to the norm, our eyes pay attention and our brains attempt to assign meaning to that contrast. When colors in two different sections of a dashboard are the same, we are tempted to relate them to one another. We merrily assume that we can use colors such as red, yellow, and green to assign important meanings to data, but in doing so we exclude the 10% of males and 1% of females who are color-blind. In Chapter 4, *Tapping into the Power of Visual Perception*, we'll learn a bit about color and how it can be used meaningfully and powerfully.

Designing an Unattractive Visual Display

Not being one to mince words for the sake of propriety, I'll state quite directly that some dashboards are just plain ugly. When we see them, we're inclined to avert our eyes—hardly the desired reaction to a screen that's supposed to be supplying us with important information. You might have assumed from my earlier warning against unnecessary decoration that I have no concern for dashboard aesthetics, but that's not the case. When a dashboard is unattractive—unpleasant to look at—the viewer is put in a frame of mind that is not conducive to its use. I'm not advocating that we add touches to make dashboards pretty, but rather that we attractively display the data itself, without adding anything that distracts from or obscures it. (We'll examine the aesthetics of dashboard design a bit in Chapter 7, *Designing Dashboards for Usability*.)

Figure 3-32 on the next page is a stellar example of unattractive dashboard design. It appears that the person who created this dashboard attempted to make it look nice, but he just didn't have the visual design skills needed to succeed. For instance, in an effort to fill up the space, some sections (such as the graph at the bottom right) were simply stretched. Also, although shades of gray can be used effectively as the background color of graphs, this particular shade is too dark. The image that appears under the title "Manufacturing" is clearly an attempt to redeem this dreary dashboard with a splash of decoration, but it only serves to distract from the data and isn't even particularly nice to look at. The guiding design principle of simplicity alone would have saved this dashboard from its current agony.