

Telling Compelling Stories with Numbers




Stephen Few
of
Perceptual Edge

www.PerceptualEdge.com

Copyright 2010 Stephen Few

Upon this gifted age, in its dark hour,
rains from the sky a meteoric shower
of facts...they lie, unquestioned, uncombined.
Wisdom enough to leach us of our ill
is daily spun; but there exists no loom
to weave it into a fabric.



"Huntsman, What Quarry?", 1939, Edna St. Vincent Millay

Even though Millay wrote these words long before the information age as we're experiencing it today began, she poignantly described one of the fundamental challenges of our time: to make use of data that now exists in abundance.



We're buried in data today—overwhelmed—not because there's too much but because we haven't learned how to weave into something meaningful. For that reason, the promise of the information age still eludes us.



We must learn to tap into the steady stream of data in ways that allows us to make sense of it and then use what we learn to do good in the world. Data must be expressed ways that make the stories the dwell within visible to our eyes and meaningful to our brains. When words and numbers fail, we need pictures to coax the stories that live in our data from the shadows into the light.

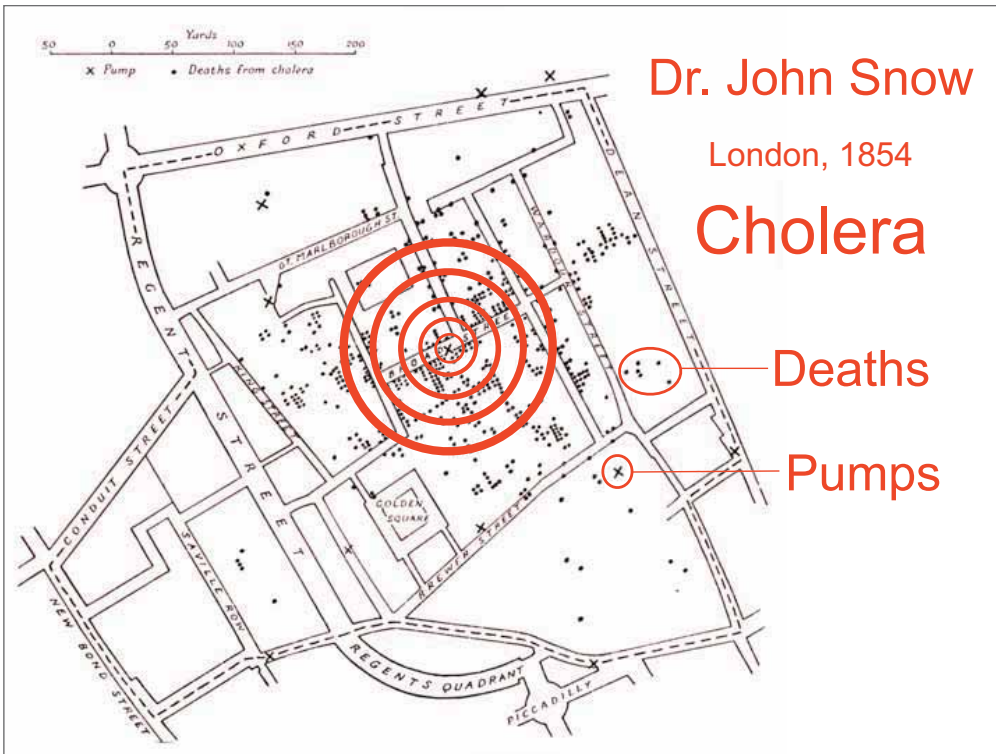


And not just any picture will do. We must craft pictures that clarify and enlighten.


A picture is worth a thousand words.



The right picture can sometimes tell a story in a way that no amount of words could ever match. To take advantage of visualization's great potential, we first must know when pictures rather than words are needed.



Here's a data visualization that was needed, for it told an enlightening story that saved hundreds of lives. This map was researched and drawn by Dr. John Snow in the mid-19th century during cholera epidemic in London. He was trying to determine how cholera was spread so he could figure out how to stop it. Contrary to the current miasma theory, that cholera was spread by noxious fumes in the air, Snow suspected that it was spread in the water supply. By marking the location of each death, shown here as the black dots, and the location of each well from which water was drawn, the connection between the disease and the city's water supply became visible as emanating from the Broad Street well. Snow used this evidence to persuade the city fathers to remove the handle from the Broad Street pump. Within a short time the epidemic ceased.

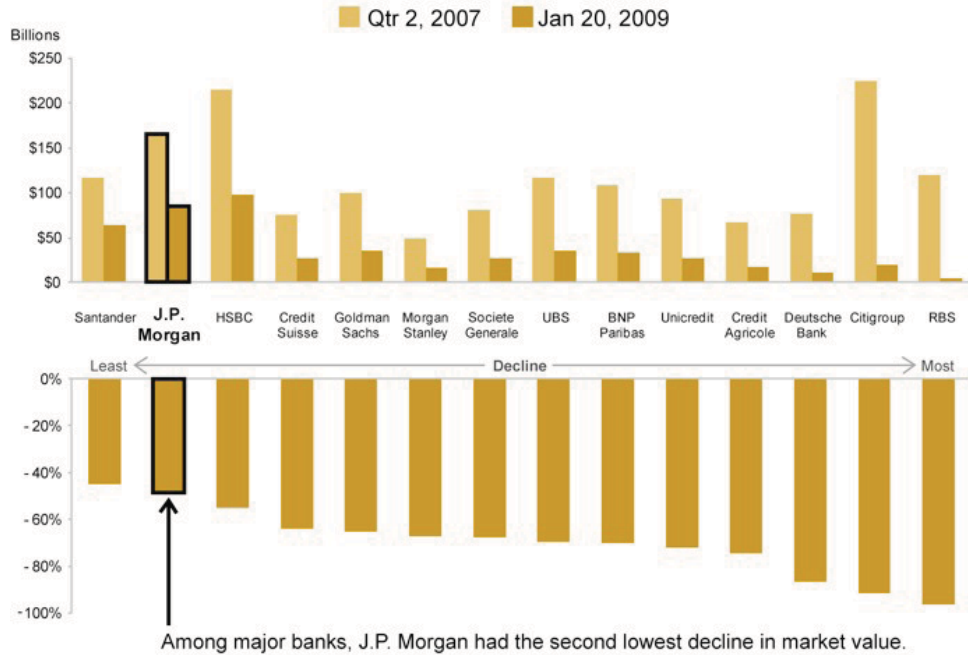


We must not only know when to use pictures of data, but also how to design them to tell stories clearly, accurately, and compellingly.

This recent visualization fails to do this. This series of circles within circles--blue for the market values of banks in quarter 2 of 2007, before the recent financial meltdown, and green for declined values as of January of 2009--was published by Bloomberg. You would never guess its purpose, however, which was to show that J. P. Morgan's decline in market value was less severe than all other major banks except one: Santander.


This picture of the data doesn't tell the story clearly, simply, or accurately. The comparative sizes of the circles are far from the comparative market values. Even if the sizes of the circles were accurate, we would still struggle with this chart because visual perception isn't well-tuned to handle size comparisons, but it is tuned to handle length comparisons,...

Declines in Bank Market Values Since the Financial Crisis Began



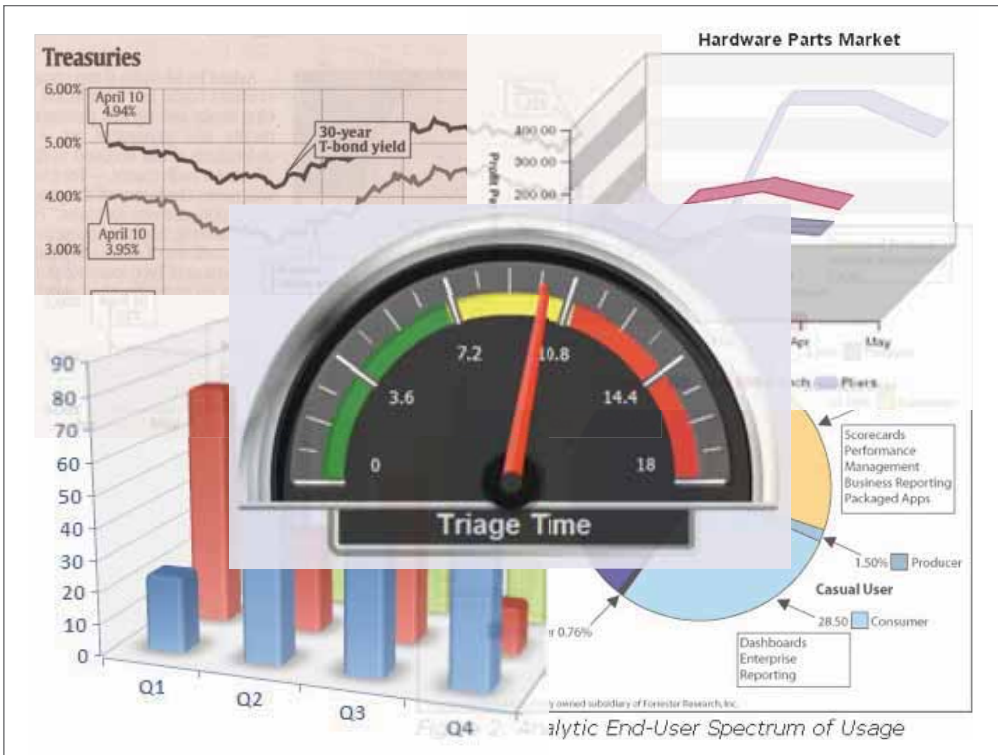
...such as the lengths of these bars in my redesign of the chart. We can now easily see that J. P. Morgan lost roughly half of its market value during this period, but the fact that its losses were less severe than all by one bank—Santander—still isn't obvious. The right addition to the picture, however, such as this one that displays the losses directly, can make this part of the story clear as well.

We've been telling them with graphs for quite awhile.




William Playfair – 1786

In 1786, a roguish Scot – William Playfair – published a small atlas that introduced or greatly improved most of the quantitative graphs that we use today. Prior to this, graphs of quantitative data were little known.



Today, 220 years later, partly due to the arrival of the PC, graphs are commonplace, fully integrated into the fabric of modern communication. Surprisingly, however, Playfair's innovative efforts – sprung from meager precedent – are still superior to most of the graphs produced today.



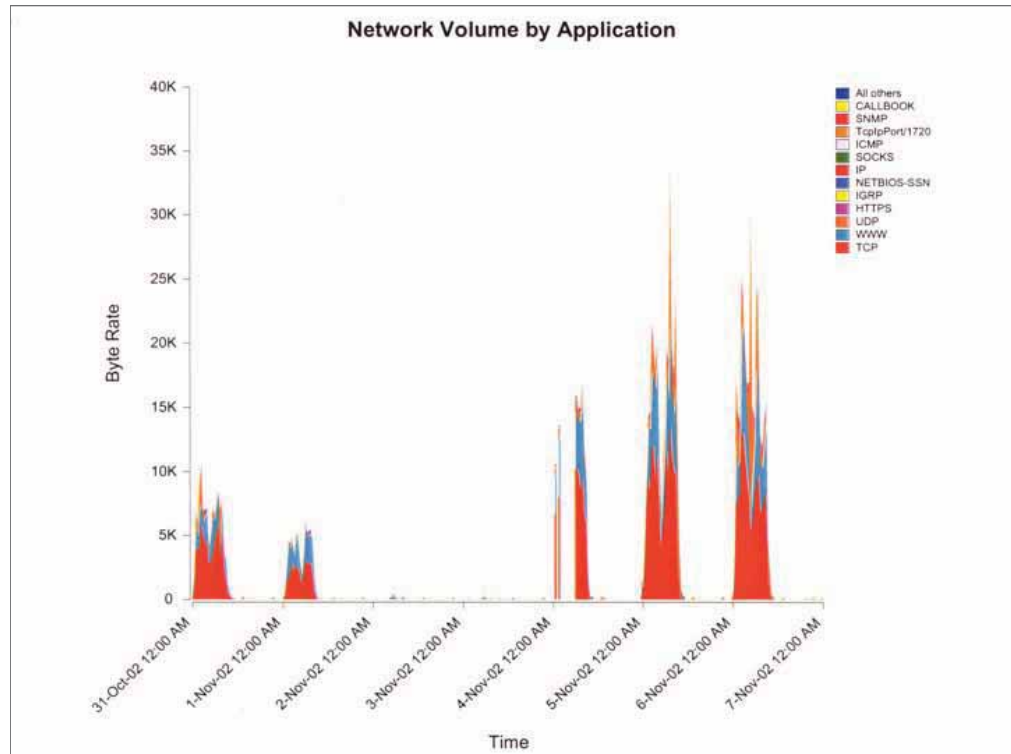
Problems like this silly pie chart on Fox News, which adds up to 193%, are far too common.

Finally...


Effective network monitoring has arrived!

- Near real-time
- Phenomenally user-friendly
- Instant insight → effective response

Imagine that you've been invited to another of those many meetings that you're required to attend. You're one of several managers in the IT department. Like most meetings, this one begins with the light of a projector suddenly illuminating a screen. Bursting with excitement, a young fellow at the front of the room announces that you will now receive a daily report that will inform you how the network is being utilized, and then the graph on the next slide appears.




You stare at this graph intently, trying your best to keep any hint of confusion from showing on your face. From your peripheral vision you can see that the CIO (Chief Information Officer) is smiling broadly and nodding with obvious understanding. You and everyone else in the room begin to nod enthusiastically as well. You feel dumb, because you have no idea what this graph is trying to say. What you don't realize is that you are not alone.



I wrote the book *Show Me the Numbers: Designing Tables and Graphs to Enlighten* in 2004 to help people like you respond to the challenges that you face every day when presenting quantitative information.

Department	Expenses			
	Jan	Feb	Variance	Change %
Sales	9,933	9,293	-640	-6%
Marketing	5,385	5,832	+447	+8%
Operations	8,375	7,937	-438	-5%
Total	\$23,693	\$23,062	-\$1,327	-3%



Quantitative information is primarily communicated through tables and graphs, but few communicate effectively.

Why? Few people are trained.


Why? Few people recognize the need.

Why? Few examples of good design exist to expose the problem.


“Poor documents are so commonplace that deciphering bad writing and bad visual design have become part of the coping skills needed to navigate in the so-called information age.” Karen A. Schriver, *Dynamics in Document Design*, John Wiley & Sons, Inc., 1997.

“The public is more familiar with bad design than good design. It is, in effect, conditioned to prefer bad design, because that is what it lives with. The new becomes threatening, the old reassuring.” (Kevin Mullet and Darrel Sano, *Designing Visual Interfaces*, Sun Microsystems, Inc., 1995 – quoting Paul Rand, *Design, Form, and Chaos*)

Effective communication is not always intuitive – it must be learned.



Intentional deceit is no longer our biggest problem. In 1954, Darrell Huff wrote his best-selling book about how people were often intentionally using statistics, including graphs, to spread misinformation, especially in favor of their own products or causes. Today, vastly more misinformation is disseminated unintentionally because people don't know how to use charts to communicate what they intend.



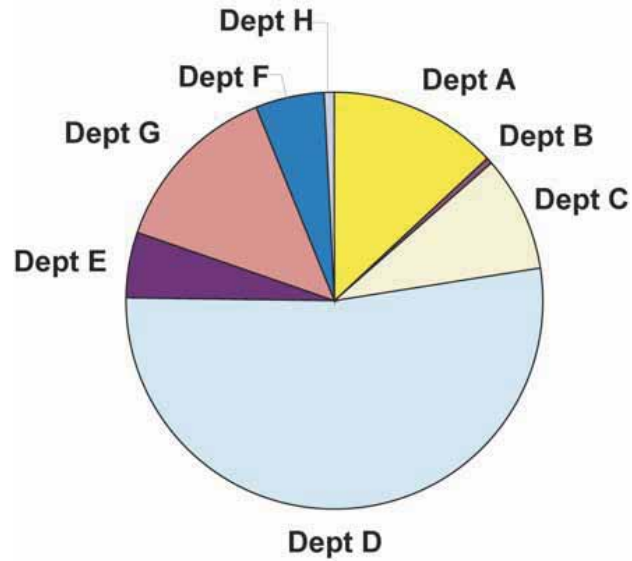
Numbers are commonly obscured, then dressed up to look *sexy*. Why?

When the PC was introduced, software soon made the arduous task of table and graph creation as easy as 1-2-3 (literally “Lotus 1-2-3”, the software that was the first to legitimize the PC as a viable tool for business). Unfortunately, this improvement in ease and efficiency was not accompanied by instruction in visual design for communication. People today think that if they know how to click with the mouse to create a table or graph, they know how to present data effectively.

“In the two centuries since [the invention of the first graphs], ...charts have become commonplace. With the advent of modern computer tools, creating graphs from data involves trivial effort. In fact, it has probably become too easy. Graphs are often produced without thought for their main purpose: to enlighten and inform the reader.” (Jonathan G. Koomey, *Turning Numbers into Knowledge*, Analytics Press, 2001)

I can talk about this all day, but the best way to make my point convincingly is to show you.


Departmental Expenses



The purpose of this graph is to display how *Department G* is doing regarding expenses compared to the other departments. Is the message clear?


Often, when someone creates a graph that appears inadequate somehow, they try to fix it with sizzle, as in the next slide.

Departmental Expenses



Does the addition of 3D improve this pie chart? Definitely not. In fact, it actually makes it harder to read.

Departmental Expenses



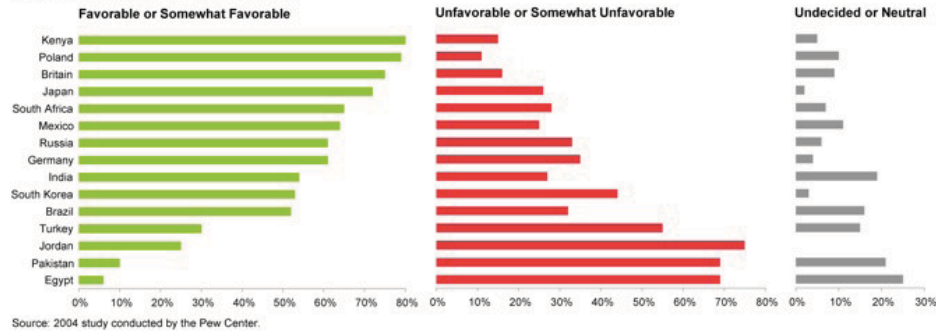
Though it lacks flash and dazzle, this simple bar graph tells the story elegantly.

Favorable or Unfavorable View of the U.S.

Brazil: % with somewhat or very favorable opinion of the U.S.:	52%
Brazil: % with somewhat or very unfavorable opinion of the U.S.:	32%
Mexico: % with somewhat or very favorable opinion of the U.S.:	64%
Mexico: % with somewhat or very unfavorable opinion of the U.S.:	25%
Britain: % with somewhat or very favorable opinion of the U.S.:	75%
Britain: % with somewhat or very unfavorable opinion of the U.S.:	16%
Germany: % with somewhat or very favorable opinion of the U.S.:	61%
Germany: % with somewhat or very unfavorable opinion of the U.S.:	35%
Russia: % with somewhat or very favorable opinion of the U.S.:	61%
Russia: % with somewhat or very unfavorable opinion of the U.S.:	33%
Poland: % with somewhat or very favorable opinion of the U.S.:	79%
Poland: % with somewhat or very unfavorable opinion of the U.S.:	11%
South Africa: % with somewhat or very favorable opinion of the U.S.:	65%
South Africa: % with somewhat or very unfavorable opinion of the U.S.:	28%
Kenya: % with somewhat or very favorable opinion of the U.S.:	80%
Kenya: % with somewhat or very unfavorable opinion of the U.S.:	15%
India: % with somewhat or very favorable opinion of the U.S.:	54%
India: % with somewhat or very unfavorable opinion of the U.S.:	27%
Japan: % with somewhat or very favorable opinion of the U.S.:	72%
Japan: % with somewhat or very unfavorable opinion of the U.S.:	26%
South Korea: % with somewhat or very favorable opinion of the U.S.:	53%
South Korea: % with somewhat or very unfavorable opinion of the U.S.:	44%
Egypt: % with somewhat or very favorable opinion of the U.S.:	6%
Egypt: % with somewhat or very unfavorable opinion of the U.S.:	69%
Pakistan: % with somewhat or very favorable opinion of the U.S.:	10%
Pakistan: % with somewhat or very unfavorable opinion of the U.S.:	69%
Turkey: % with somewhat or very favorable opinion of the U.S.:	30%
Turkey: % with somewhat or very unfavorable opinion of the U.S.:	55%
Jordan: % with somewhat or very favorable opinion of the U.S.:	25%
Jordan: % with somewhat or very unfavorable opinion of the U.S.:	75%

I found this table on the Web site for Bill Moyers' public television show "Now". I felt that it provided important information that deserved a better form of presentation. In this case the story could be told much better in visual form.


Current World Opinions About the U.S.A



Source: 2004 study conducted by the Pew Center.


This series of related graphs tells the story in vivid terms and brings facts to light that might not ever be noticed in the table.

Figure 2: Maine AIDS Cases and Deaths Among Persons with AIDS, 1984-2005



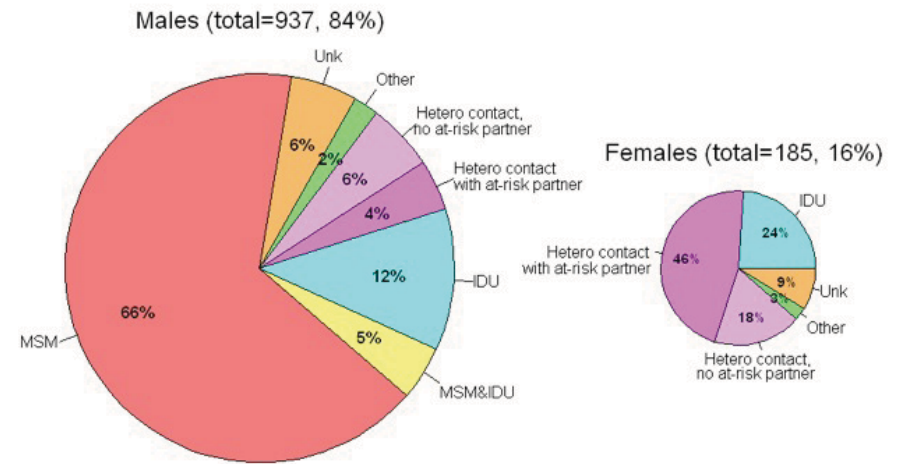
Note: Data concerning 2005 deaths will be available July, 2006

Here's a public health example from the state of Maine. This graph contains important patterns that are difficult to discern due to clutter. It's hard to independently discern the patterns of change through time of diagnosed AIDS cases vs. deaths or to compare these patterns to one another.



But in this graph there is no clutter and the patterns are crystal clear and easy to compare.

Figure 5: People Living with Diagnosed HIV Infection¹, Mode of Transmission by Sex



If you were asked to tell the story contained in this display, it would take you some time to put it together before you could even begin to explain it to others.

How are HIV infections transmitted? males vs. females




In this display of the same information, however, the story is clear and aspects of the story that weren't apparent in the pie charts jump right out.

Fundamental challenges of data presentation


1. Determining the medium that tells the story best
2. Designing the visual components to tell the story clearly

Either

Product	Units Sold	Revenue
Food	34,837	746,383
Beverage	42,374	845,984
Total	77,211	1,592,367



and which kind?



1. You begin by determining the best medium for your data and the message you wish to emphasize. Does it require a table or a graph? Which kind of table or graph?
2. Once you've decided, you must then design the individual components of that display to present the data and your message as clearly and efficiently as possible.

The solutions to both of these challenges are rooted in an understanding of visual perception.

Tables work best when...

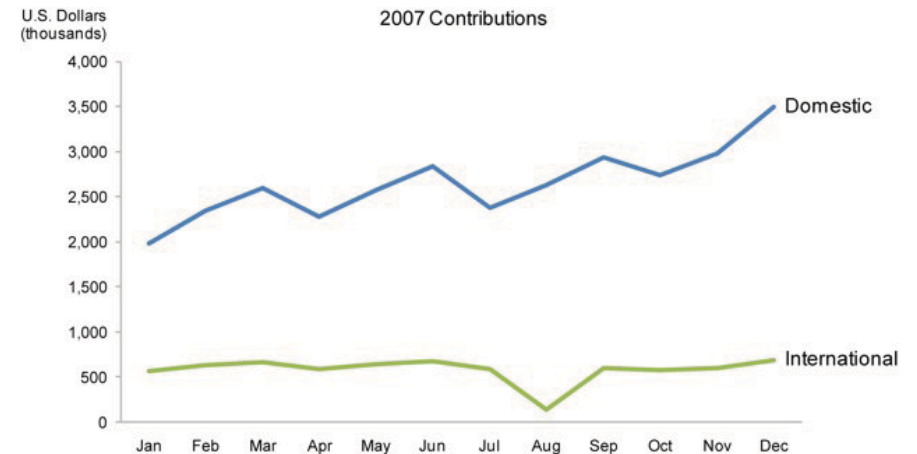
Year	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Annual
1990	127.4	128.0	128.7	128.9	129.2	129.9	130.4	131.6	132.7	133.5	133.8	133.8	130.7
1991	134.6	134.8	135.0	135.2	135.6	136.0	136.2	136.6	137.2	137.4	137.8	137.9	136.2
1992	138.1	138.6	139.3	139.5	139.7	140.2	140.5	140.9	141.3	141.8	142.0	141.9	140.3
1993	142.6	143.1	143.6	144.0	144.2	144.4	144.4	144.8	145.1	145.7	145.8	145.8	144.5
1994	146.2	146.7	147.2	147.4	147.5	148.0	148.4	149.0	149.4	149.5	149.7	149.7	148.2
1995	150.3	150.9	151.4	151.9	152.2	152.5	152.5	152.9	153.2	153.7	153.6	153.5	152.4
1996	154.4	154.9	155.7	156.3	156.6	156.7	157.0	157.3	157.8	158.3	158.6	158.6	156.9
1997	159.1	159.6	160.0	160.2	160.1	160.3	160.5	160.8	161.2	161.6	161.5	161.3	160.5
1998	161.6	161.9	162.2	162.5	162.8	163.0	163.2	163.4	163.6	164.0	164.0	163.9	163.0
1999	164.3	164.5	165.0	166.2	166.2	166.2	166.7	167.1	167.9	168.2	168.3	168.3	166.6
2000	168.8	169.8	171.2	171.3	171.5	172.4	172.8	172.8	173.7	174.0	174.1	174.0	172.2
2001	175.1	175.8	176.2	176.9	177.7	178.0	177.5	177.5	178.3	177.7	177.4	176.7	177.1
2002	177.1	177.8	178.8	179.8	179.8	179.9	180.1	180.7	181.0	181.3	181.3	180.9	179.9

- Used to look up individual values
- Data must be precise

What do graphs do well?



Feature patterns, trends, and exceptions.



The saying, “A picture is worth a thousand words,” applies quite literally to quantitative graphs. By displaying quantitative information in visual form, graphs efficiently reveal information that would otherwise require a thousand words or more to adequately describe.

In this example of purely manufactured data to illustrate my point...

Joseph Berkson once stated what happens quite powerfully: When we visualize the data effectively and suddenly, we experience “interocular traumatic impact”—a conclusion that hits us between the eyes.


Can you see trends, patterns and exceptions in this table of numbers?

Data table for figure 1. Total population and older population: United States, 1950-2050

Year	All ages	65 years and over	75 years and over
	Number		
1950	150,216,110	12,256,850	3,852,395
1960	179,325,657	16,207,237	5,359,338
1970	203,211,926	20,065,502	7,630,046
1980	226,545,805	25,549,427	9,968,822
1990	248,709,873	31,078,895	13,033,400
2000	281,421,906	34,991,753	16,600,767
2004	293,622,764	36,221,951	17,777,261
2010	308,935,581	40,243,713	18,974,204
2020	335,804,546	54,631,891	22,852,732
2030	363,584,435	71,453,471	33,505,538
2040	391,945,658	80,049,634	44,579,726
2050	419,853,587	86,705,637	48,763,200

Tables work great for looking up individual values, but they don't reveal trends, patterns, and exceptions very well.

Graphs make trends, patterns and exceptions visible.



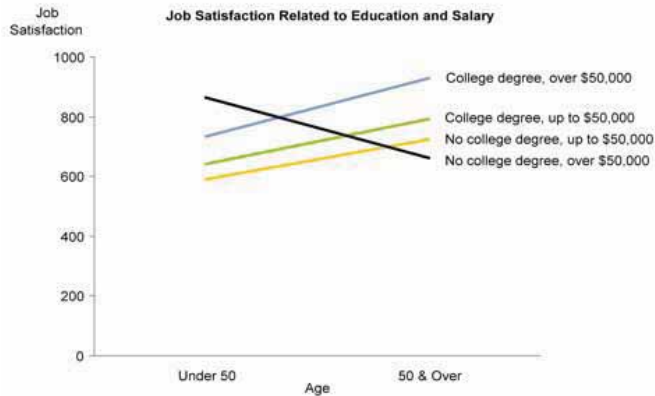
Now, however, by expressing this same information visually, giving shape to the data, the trends come alive.

Even small data sets can benefit from graphs.

What's the message?

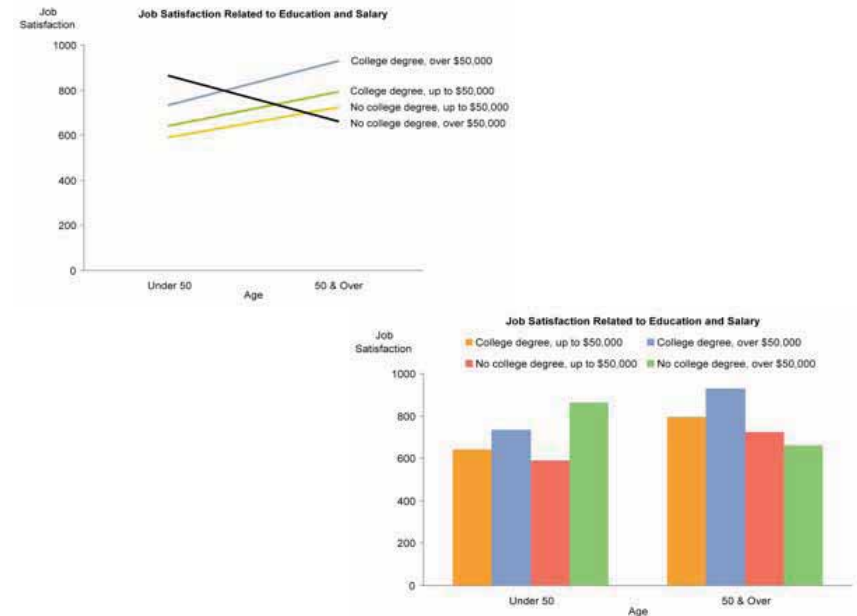
Job Satisfaction By Income, Education & Age

Income	College Degrees		No College Degrees	
	Under 50	50 & over	Under 50	50 & over
Up to \$50,000	643	793	590	724
Over \$50,000	735	928	863	662



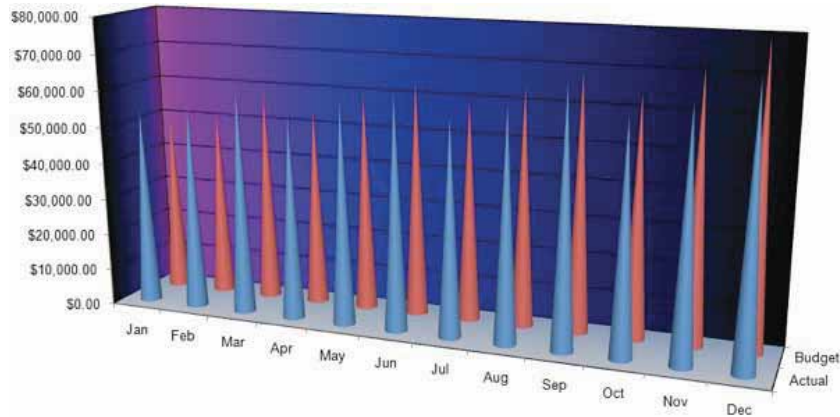
The fact that job satisfaction for employees without a college degree decreases significantly in their later years doesn't jump out at you when you examine the table, but it is immediately obvious when you examine the graph.

...but only if you design them correctly.



The type of graph that is selected and the way it's designed also have great impact on the message that is communicated. By simply switching from a line graph to a bar graph, the decrease in job satisfaction among those without college degrees in their later years is no longer as obvious.

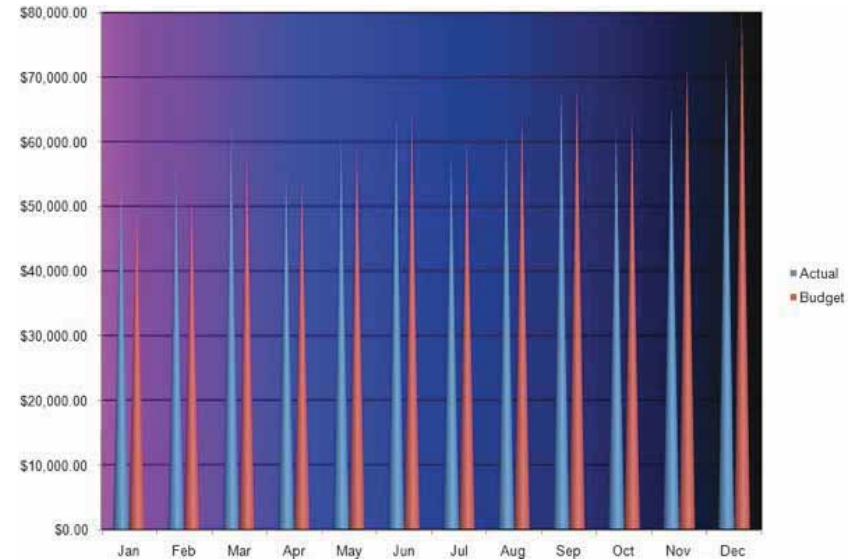
Most graphing software misleads us.




This is the kind of graph that software products, including Excel, encourage us to create. They give us an infinite selection of poorly-designed graphs from which to choose. What we really need, however, is a small selection of graphs that really work.

Using this graph, try to see the pattern of change across the months in actual expenses. Try to determine one of the actual values. Try to compare actual expenses to the budget across time.


Let's transform this graph into one that communicates.




We have now removed the useless 3-D effects and angle, which makes the data easier to read.




We have now removed the background fill color.




We have now replaced the silly cones with regular bars.




We have now removed the tick marks, which aren't necessary. Tick marks are not needed to separate the months along the X-axis and because horizontal grid lines are being displayed, there is no need for tick marks on the Y-axis either.




We have now enlarged the text, making it easier to read.




We have now removed the unnecessary decimal places in the dollar amounts along the Y-axis.




We have now removed the redundant dollar signs and labeled the unit of measure (U.S. \$) clearly.




We have now reoriented the Y-axis label to the horizontal and placed it above the axis to make it easier to read.




We have now reoriented and repositioned the legend to make it easier to associate it with the data bars.




We have now changed the color of the Budget bar to be more visually pleasing in relation to the blue Actual bars. Changing from the color red also removed the possibility people interpreting the data as something bad or a warning, which red is often used to represent.




We have now reduced the visual salience of the Budget values, because they are less important than the Actual values, and have done so in a way that reduced clutter.




We have now made it much easier to see the pattern of change through time by using lines rather than bars to represent the data.




We have now made it much easier to examine the differences between actual expenses and the budget by spreading them across more space.




We have now labeled the lines directly, removing the need for a legend.



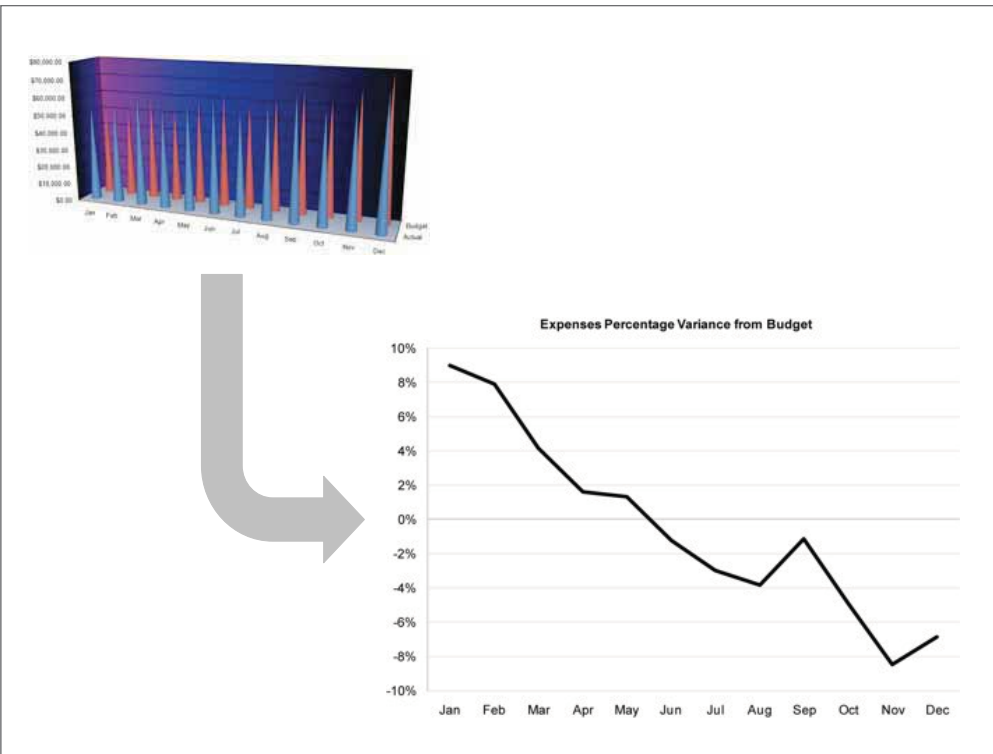
We have now changed the lines to two shades of gray to guaranty that even if the graph is printed on a black-and-white printer or photocopier, they will still look distinctly different from one another.



We have now represented the variance of actual expenses from the budget directly, as a single line.

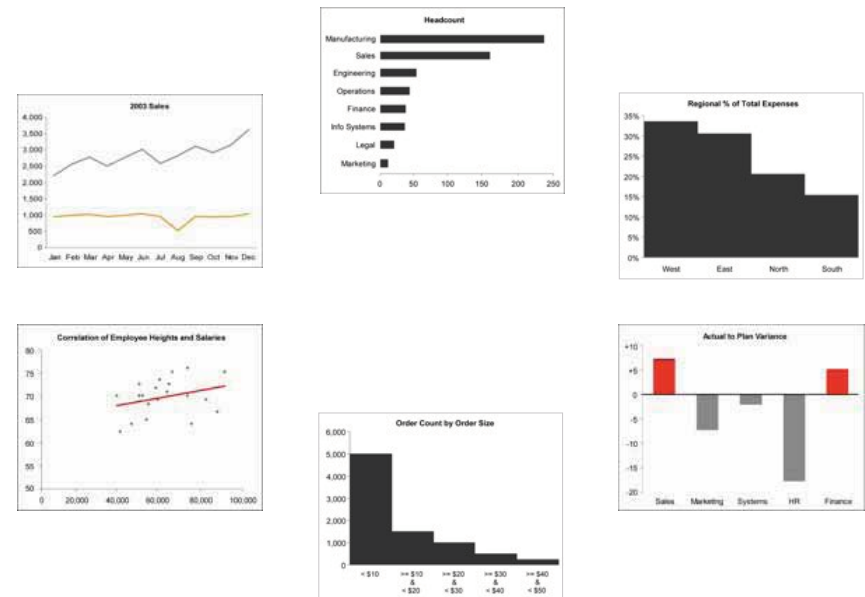


As our final step, we have expressed variance as a percentage, to provide a better measure of performance.




Our final solution, which we produced in sixteen steps, could have easily been our original solution. It usually takes no longer to design effective graphs than those that communicate poorly, if at all.

Six common quantitative relationships




Quantitative messages always feature relationships. Each of these graphs illustrates a different type of quantitative relationship. Just as in life in general, the interesting and important content of a graph always involves relationships.

Time Series




A time-series graph has a categorical scale that represents time, subdivided into a particular unit of time, such as years, quarters, months, days, or even hours. These graphs provide a powerful means to see patterns in the values as they march through time.

Could this pattern of change be displayed more clearly?



Here's an attempt to display a time-series relationship regarding HIV diagnoses, which works fairly well, but the trend and patterns could be much more clearly displayed.


Bars and lines tell time differently.



Here's the same exact data presented in two ways: the top graph uses bars and the bottom graph uses a line. Which displays the shape of change through time more clearly?


Bars work well for comparing individual values to one another, but lines show the shape of change through time much more clearly.

Ranking




Ranking graphs show the sequence of a series of categorical subdivisions, based on the measures associated with them.

Unsorted




Sorted



A simple sort adds meaning and clarity. In the display of trauma registry injuries by county on the left, notice how difficult it is to compare the values and to get a sense of rank when they aren't sequenced by size.


The same information is displayed on the right, this time with the counties arranged by the number of injuries. If the purpose of the display is to look up individual values, which is the only thing that alphabetical order supports, a table would work much better. The ranking display on the right, however, tells a useful story.

Part-to-Whole




A part-to-whole graph shows how the measures associated with the individual categorical subdivisions of a full set relate to the whole and to one another.

Pie charts are difficult to read.




Part-to-whole relationships are typically displayed as pie charts, but they don't communicate very effectively. If you want to see the order of items and to compare the size of one to another, with this display you would struggle,...

A bar graph tells this story clearly.



...but with this bar graph the story simply—nothing fancy, but clear.

Save the pies for dessert!




Coda Hale once expressed his opinion of pie charts quite colorfully:

Pie charts are the information visualization equivalent of a roofing hammer to the frontal lobe...[Piecharts] have no place in the world of grownups, and occupy the same semiotic space as short pants, a runny nose, and chocolate smeared on one's face. They are as professional as a pair of assless chaps. Anyone who suggests their use should be instinctively slapped.

For more information about pie charts, read my article "Save the Pies for Dessert."

Deviation



A deviation graph shows how one or more sets of values differ from a reference set of values.

Don't force people to calculate differences in their heads.

Median Annual Household Income (Current Dollars), Utah and U.S., 1984-2004



When people primarily need to see the differences between things, show them the difference directly, rather than showing them the two sets of values and forcing them to construct a new picture in their heads of how they differ.


The difference between the median annual household income in Utah and in the U.S. as a whole isn't shown directly in this graph,...

Display deviations directly.



...but this graph directly expresses how household income in Utah has differed over time from the U.S. as a whole in positive and negative dollars.

Distribution

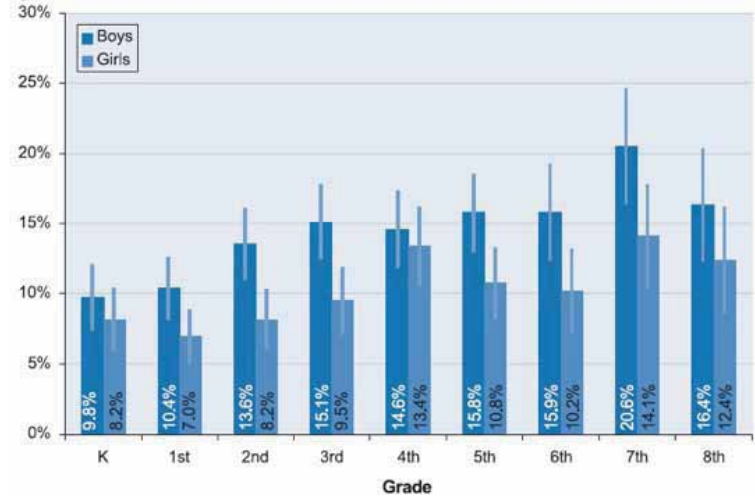


This type of distribution graph, called a frequency distribution, shows the number of times something occurs across consecutive intervals of a larger quantitative range. In a frequency distribution, a quantitative scale (in this case the range of dollar values of orders) is converted to a categorical scale by subdividing the range and giving each of the subdivisions a categorical label (“< \$10”, and so on).

Histograms with multiple data sets can look cluttered.


Overweight* by Grade and Gender

Figure 1. Percentage of Utah children who were overweight by grade and gender, Utah, 2002




Here's a graph that attempts to show the distribution of overweight children by grade separately for boys and girls, but doing it in this way results in clutter that makes the patterns difficult to segregate and compare.

Clutter can be easily eliminated.



or




This pair of histograms—one for boys and one for girls—however, are arranged in a way that makes the patterns of each easy to see, yet still easy to compare.

Even better, by using lines rather than bars, the separate patterns can be shown in the same graph in a way that features the shape of the patterns and how they differ.


Correlation

Avg Salary
(U.S. dollars)

Relationship between Employee Heights and Salaries



A correlation graph shows whether two paired sets of measures vary in relation to one another, and if so, in which direction (positive or negative) and to what degree (strong or weak). If the trend line moves upwards, the correlation is positive; if it moves downwards, it is negative. A positive correlation indicates that as the values in one data set increase, so do the values in the other data set. A negative correlation indicates that as the values in one data set increase, the values in the other data set decrease. In a scatter plot like this, the more tightly the data points are grouped around the trend line, the stronger the correlation.



I don't run across many correlation displays that are used for business analysis. This is a shame, because correlations can reveal important information about the causes of things. Unless we understand causation, we can't change what's happening.

This example, based on WHO data, explores the correlation between adult literacy and fertility rate by country. A correlation clearly exists: higher literacy corresponds to lower rates of fertility. It is also clear from this display that the highest rates of fertility all occur in Africa (the blue circles), which the one exception of Yemen (the one green circle at the high end of fertility).

Six common relationships in graphs

- Time-series
- Ranking
- Part-to-whole
- Deviation
- Distribution
- Correlation

Almost every quantitative story involves one or more of these fundamental relationships, these comparisons between numbers. The skills required to tell these stories clearly are easy to learn.

Dashboards are everywhere!



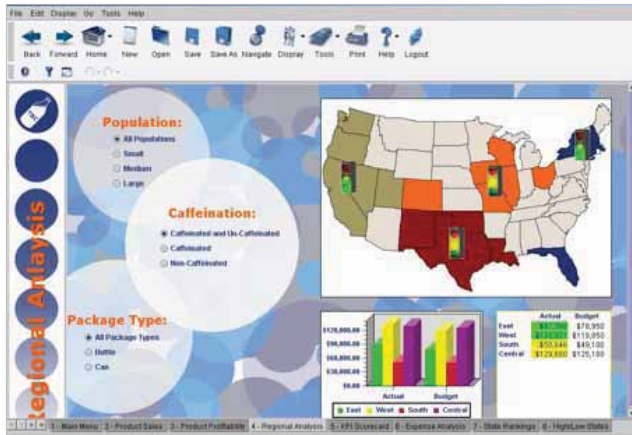
Dashboards are extremely popular today. Everyone assumes that dashboards are useful. Like all aspects of business intelligence, however, they can be done well or they can be done poorly.

But most dashboards fail.



Dashboards like this are typical. You can find them on the websites of most business intelligence vendors. In their attempt to dazzle us visually, they fail entirely to present information in a way that can be understood at a glance.


The problem?



- They say too little
- What they do say, they say poorly

This failure is not rooted in technology; it is a failure of communication, rooted in poor design. The main problem that plagues dashboards today is that they don't say enough and what they do say they don't say very well. In other words, they are not very informative. When designed properly, a dashboard provides an overview of what's going on, clearly and rapidly.

Dashboards are not for show.



They are for communication.

Dashboards are not for show. No amount of cuteness and technical wizardry can substitute for clear communication. My purpose today is to present the design practices you cannot ignore if you wish to build dashboards that set the stage for informed response.

Dashboard



Everyone defines dashboards differently, if at all, including the vendors who sell them. Before proceeding, we need a definition.

The examples that follow in the next few slides are all screen prints of displays that their creators call “dashboards” (or in one case a “cockpit”). See if you can come up with a definition of “dashboard” based on these examples.


Dashboard

A visual display
of
the most important information needed
to achieve one or more objectives
that has been
consolidated on a single computer screen
so it can be
monitored and understood at a glance

A dashboard is a visual display of the most important information needed to achieve one or more objectives; consolidated and arranged on a single screen so the information can be monitored at a glance.

This definition first appeared in the March 20, 2004 issue of *Intelligent Enterprise* magazine in an article written by Stephen Few entitled “Dashboard Confusion.”

The fundamental design challenge



The fundamental challenge of dashboard design is to display all the required information on a single screen:


- clearly and without distraction
- in a manner that can be quickly examined and understood

Think about the cockpit of a commercial jet. Years of effort went into its design to enable the pilot to see what's going on at a glance, even though there is much information to monitor. Every time I board a plane, I'm grateful that knowledgeable designers worked hard to present this information effectively. Similar care is needed for the design of our dashboards. This is a science that few of those responsible for creating dashboards have studied.

Like airplane cockpits, dashboards require thoughtful design.



Think about the cockpit of a commercial jet. Years of effort went into its design to enable the pilot to see what's going on at a glance, even though there is much information to monitor. Every time I board a plane, I'm grateful that knowledgeable designers worked hard to present this information effectively. Similar care is needed for the design of our dashboards. This is a science that few of those responsible for creating dashboards have studied.




It is because of these unique design requirements that I wrote the book *Information Dashboard Design*. The principles are not difficult to learn, but they aren't obvious until someone points them out.

Visual monitoring


1. Scan the big picture



3. Link to supporting details



2. Zoom in on important specifics



The process of visual monitoring involves a series of sequential steps that the dashboard should be designed to support. The user should begin by getting an overview of what's going on and quickly identifying what needs attention. Next, the user should look more closely at each of those areas that need attention to be able to understand them well enough to determine if something should be done about them. Lastly, if additional details are needed to complete the user's understanding before deciding how to respond, the dashboard should serve as a seamless launch pad to that information, and perhaps even provide the means to initiate automated responses, such as sending emails to those who should take action.

Monitoring is most efficiently done with our eyes.



70%



30%


Monitoring is a cognitive activity that receives input primarily through the visual channel because this is our most powerful sense, working at high speeds of parallel input, able to detect subtle distinctions and complex patterns. Seventy percent of the sense receptors in the human body reside in our eyes.

It's sometimes appropriate to use auditory signals to get people's attention when they aren't looking at the screen, but we must make use of their eyes to perceive the rest.

The advantage of graphical communication

Metric	Variance to Target
Revenue	-0.68%
Profit	-17.90%
Avg Order Size	-5.90%
On Time Delivery	2.35%
New Customers	58.38%
Cust Satisfaction	-12.22%
Market Share	-13.33%

Text must be read, processed serially



Graphics can be perceived at a glance, processed in parallel

The visual orientation of dashboards is important due to the speed of perception that's usually required to monitor information. The faster you must assess what's going on, the more you should rely on graphical means to display the information.

Text must be read, which involves a relatively slow, serial process.

Certain visual properties, however, can be perceived at a glance, without conscious thought. With the graphical display on the right, it's quick and easy to see which bars exceed target, marked by the short vertical line, and which fall short.

13 common mistakes in dashboard design

6



1. Exceeding the boundaries of a single screen

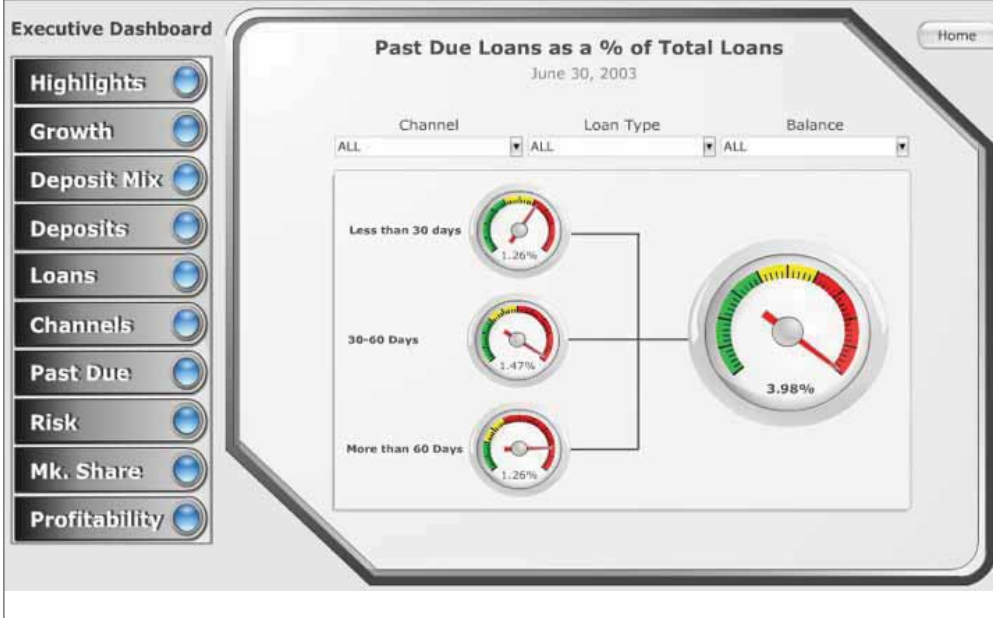


My insistence that a dashboard confine its display to a single screen with no need for scrolling might seem arbitrary, but it is based on the findings of perceptual and cognitive research. Something powerful happens when you see things together, all within eye span. You are able to make comparisons, spot relationships, and see the big picture. This ability is lost when you must lose sight of some data in order to scroll down or over to see other data. Part of the problem is that we can only hold a few chunks of information at a time in short-term memory. Relying on the mind's eye to retain a visualization that is no longer visible is a limited venture. One of the great benefits of a dashboard is the simultaneity of vision, the ability to see everything that you need at once. This enables comparisons that lead to insights that might not occur in any other way.

The dashboard shown above not only leaves us wondering what lies below the bottom of the screen, we're also only given immediate visual access to the first of many metrics that appear at the top right, under the heading "No. of transactions". Don't force your viewers to scroll around to see what they need. I'd prefer a printed report that extends across multiple pages, because at least then I could lay out all of the pages at once for simultaneous viewing.

(Source: Website of Business Objects.)

1. Exceeding the boundaries of a single screen



Data is often fragmented into separate screens in one of two ways:


- It is separated into discrete screens to which one must navigate.
- It is separated into different versions of a single screen that are accessed individually.

When the information should all be seen at the same time to gain the desired insights and to make the needed connections, this fragmentation undermines the unique advantages of a dashboard.

The dashboard above fragments the data that the executives need into 10 separate dashboards. This would be fine if the executives would not benefit from seeing these various measures together, but that is hardly the case. Splitting the big picture into a series of separate, small pictures is a mistake when seeing the big picture is worthwhile.

(Source: Website of Informmersion Incorporated.)

2. Supplying inadequate context for the data




Measures of what's going on in the business rarely do well as solo acts; they need a good supporting cast to get their message across. 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? Is this better than before? The right context for the key measures makes the difference between numbers that just sit there on the screen and those that enlighten and inspire action.

The gauges above could have easily incorporated useful context, but they fall short of their potential. The center gauge tells us only that 7,822 units have sold year to date and that this number is good, indicated by the green arrow.

Quantitative scales on a graphic, such as those suggested by the tick marks around these gauges, are meant to help us interpret the measures, but they can only do so when scales are labeled with numbers, which these gauges lack. A great deal of the space that is used by these gauges tells us nothing whatsoever.

(Source: Website of Informmersion Incorporated.)

3. Choosing inappropriate display media

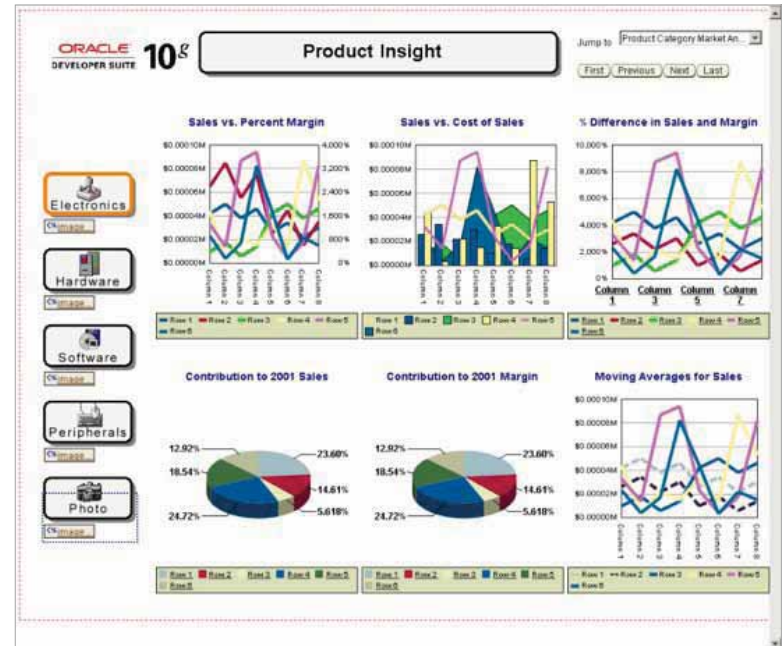


This is one of the most common design mistakes made, not just on dashboards, but in all forms of data presentation. Using a graph when a table of numbers would work better and vice versa is a frequent mistake, but the one that stands out as the most common and egregious is using the wrong type of graphic.

Without the value labels on the pie chart above, you would conclude that all of the slices are roughly the equal in size. The bar graph below it, however, tells the story clearly and quickly, because it is a better medium of display for this information.

(Source: Website of Corda Technologies Incorporated.)

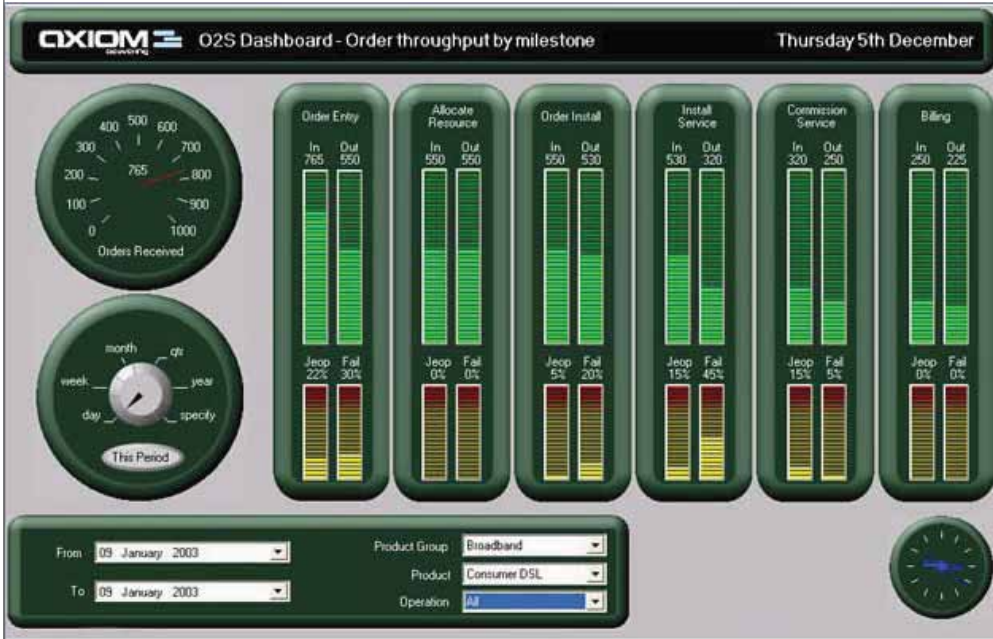
4. Ineffectively highlighting what's important



You should be able to look at a dashboard and have your eyes immediately drawn to the information that is most important. The problem with the dashboard above is that everything is visually prominent, which results in nothing standing 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 should therefore be subdued. Then there are the graphs, where the data resides, but all the data is 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.

(Source: Website of Oracle Corporation.)

5. Cluttering it with useless decoration




One of the most common problems on dashboards that are found on vendor websites 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 keep us entertained. I assure you, however, that even people who enjoy the decoration upon first sight will grow weary of it in a short time.

The makers of the dashboard above did an exceptional job of making it look like an electronic control panel. If the purpose were to train people in the use of equipment that actually looks like this by simulating it, then 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.

(Source: Website of Axiom Systems.)

6. Misusing or overusing color




Color can be used to highlight data, encode data, and create a relationship between individual items on the screen. Color choices should be made thoughtfully, based on 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 demanding. When any color appears as a contrast to the norm, our eyes pay attention and our brains attempt to assign meaning to that difference. When colors in two different displays are the same we are tempted to relate them to one another. We merrily assume that we can use colors like 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 colorblind.

Using too many colors is a common problem, especially bright colors. Because dashboards are often densely packed with information, the visual content must be kept as simple as possible. Using of too many colors can be visually assaulting.

The graph above, taken from a dashboard, misuses color in several ways, but one problem stands out as most egregious. What is the meaning of the separate color for each bar? The correct answer is that the colors mean nothing. There is no reason to assign different colors to the bars for they are already labeled along the Y axis. Nevertheless, time is wasted as our brains—whether consciously or unconsciously—search for the meaning of these differences which isn't there. It is best to keep colors subdued and neutral, except when you are using color to highlight something as especially important.

(Source: Website of Corda Technologies.)



Understand the uses and limitations of color.



We merrily assume that we can use colors like 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 colorblind.. Despite this fact, many vendors actually promote the use of red, yellow, and green stoplight colors as “business intelligence” colors, the most unintelligent choice.

Here on the right you can see how the red, yellow, and green stoplight colors appear to someone who is colorblind. The red and the green look the same. If you can't tell the colors apart, they are useless, they fail to communicate.

Too much color undermines its power.

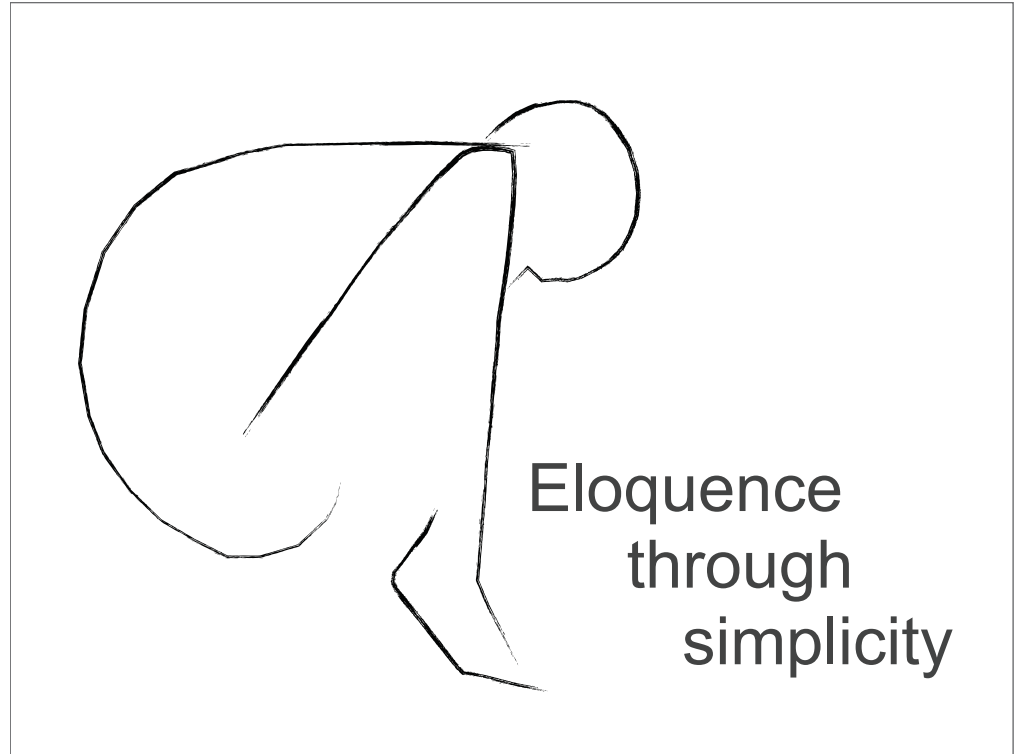


People tend to overuse color on dashboards, assuming that the brighter and more colorful they are, the better they work, but the opposite is true. People who understand color and how it can be used to display information use it sparingly.

Too much color is visually overwhelming; it tires our eyes. Also, if you use color gratuitously, you undermine its ability to be used to make things stand out. Notice how the red alerts clearly stand out in the bottom display in contrast to the neutral grays and blacks that been used elsewhere, rather than being lost in the meaninglessly colorful display above.



Which would you rather look at? What if you are Jessica's dermatologist? Dressing things up is appropriate for advertising, because the illusion pleases and sells. When you're responsible for discovering the truth and understanding it, makeup only gets in the way.



To clearly present everything on a single screen, even the slightest lack of organization will result in a confusing mess. You must condense the information, you must include only what you absolutely need, and you must use display media that can be easily read and understood even when they are small, which is often necessary.

Elegance in communication can be achieved through simplicity of design. Too often we smear a thick layer of gaudy makeup on top of the data in an effort to impress, rather than to communicate the truth in the clearest possible way.

“Simplify, simplify, simplify.”

Henry David Thoreau



We finish up now with an example of a well-designed dashboard. Notice the way that all of the principles that we've covered in this workshop have been combined to create a data-rich, yet simple and accessible sales dashboard.

Notice the following characteristics:

- Color has been used sparingly.
- The prime real estate on the screen has been used for the most important data.
- Small, concise display media have been used to support the display of a dense set of data in a small amount of space.
- Some measures have been presented both graphically and as text when precise values are frequently needed.
- Subtle means have been used to delineate and group data.
- The dashboard has not been cluttered with instructions and descriptions that will seldom be needed.



You have a choice.

The good news is, although the skills required to present data effectively are not all intuitive, they are easy to learn. The resources are available, but it won't happen unless you recognize the seriousness of the problem and commit yourself to solving it. It is up to you.