

VISUAL COMMUNICATION



**CORE DESIGN PRINCIPLES
FOR DISPLAYING
QUANTITATIVE
INFORMATION**

**STEPHEN FEW
PERCEPTUAL EDGE**

SEPTEMBER 13, 2006



WHY VISUAL COMMUNICATION?

We experience the world predominantly through our eyes. Recognition of vision's unique power has led to the development of many new forms of visual communication. Our eyes are now seen as valuable targets for visual content—messages written in light that reach out to tickle our retinas and thus our minds—hoping to make an impression. Visual communication comes in many forms, many of which are designed to entertain us through moving images, such as film, television, and video games. Visual technologies such as these have become quite sophisticated, but some forms of visual communication remain primitive by comparison, crudely attempting to deliver information that is far too important to be displayed so poorly. Graphs—the visual representation of quantitative information—are often sad examples of a crude visual medium. This is particularly sad, because the skills and technology needed to effectively present quantitative information in graphs are not complicated, but they remain rare nonetheless.

Graphs were invented to bring meanings in quantitative data to light, which could not be discerned from a table of numbers. Whether you display data in a table or a graph should not be an arbitrary decision. They serve very different purposes. Tables work marvelously when you wish to look up particular values or you need precise values. Graphs, however, make meaningful relationships between values visible by giving them size, shape, and color. There is no substitute for a well-designed graph when you wish to see or communicate meaningful trends, patterns, and exceptions in quantitative data.

Take a moment to study the table in Figure 1. If you want to know the amount of domestic sales for the month of July, it's extremely easy to find. If you want to make a simple comparison of two values, such as domestic and international sales in April, the table supports this operation as well. If you wish to discern trends or patterns of any sort, however, tables aren't very helpful.

2003 Sales (U.S. dollars in thousands)

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Domestic	1,983	2,343	2,593	2,283	2,574	2,838	2,382	2,634	2,938	2,739	2,983	3,493
International	574	636	673	593	644	679	593	139	599	583	602	690
	\$2,557	\$2,979	\$3,266	\$2,876	\$3,218	\$3,517	\$2,975	\$2,773	\$3,537	\$3,322	\$3,585	\$4,183

Figure 1: A typical table of values.

Take the same set of values and display them in a graph, however, and trends, patterns, and exceptions become immediately recognizable. Look at Figure 2 and see for yourself. What are some of the stories this set of data reveals that were not obvious when you examined the table?




Figure 2: A typical graph, which brings trends, patterns, and exceptions to light.

Here are some of the stories that you probably found quite easily while examining this graph:

- Domestic sales are trending upwards across the year as a whole.
- International sales are relatively flat across the year as a whole.
- An exception to the norm of international sales occurred in the month of August, when they took a dive.
- There is a cyclical pattern in domestic sales; one that repeats itself every quarter. Sales are lowest during the first month, they grow during the second month, and they reach their quarterly peak during the last month, then they go down again in the first month of the next quarter.

These characteristics are obvious in the graph because they are expressed in a medium that is particularly good at revealing such characteristics. Graphs are the only means of communication that reveals relationships between values in this clear and direct manner. They do so by taking advantage of visual perception—our most powerful sense.

PREREQUISITES TO EFFECTIVE VISUAL COMMUNICATION

As with all forms of communication, the first step is to determine what you want to say. To communicate your message effectively, you must know what your message is. I suspect that most people who create graphs don't take the time to think about what they want to say before producing them. If you examine a graph and its message is not clear, there's a good chance that the person who created it didn't have a particular message in mind.

Once you know what you want to say, effective visual communication is achieved by displaying information in a way that enables people to clearly see an accurate representation of your message and understand what they see. To do this, you must understand a few things about how people see (visual perception) and how people think (cognition). You must present your message in a way that takes advantage of the strengths of visual perception while avoiding its weaknesses, and in a way that matches the human thought process, augmenting it when necessary to work around limitations.

You must then develop a simple set of skills (graphicacy) based on this knowledge. This is mostly science, not art, based on clear-cut principles about what works and what doesn't, rooted in an understanding of why. It also helps if you have graph-producing software that was likewise built on this knowledge, making it easy to display information meaningfully and difficult to do the opposite. Unfortunately, few software products today were built by people who understand visual perception and cognition, but my, my, my can they ever make those pie charts sparkle and spin. If only this were useful.

THE POWER OF VISUAL PERCEPTION AND THE MIND

Of the total sense receptors in the human body, 70% reside in our eyes. Visual perception delivers the world to our brains at high speeds and with exquisite subtlety, but it does so in particular ways that are not necessarily intuitive. If we want to present information to people's eyes, we must understand a little about how the eyes work, including some very real limitations. One of the leading lights in visualization research today is Dr. Colin Ware of the University of New Hampshire. He expertly explains the importance of visualization and how it works.

Why should we be interested in visualization? Because the human visual system is a pattern seeker of enormous power and subtlety. The eye and the visual cortex of the brain form a massively parallel processor that provides the highest-bandwidth channel into human cognitive centers. At higher levels of processing, perception and cognition are closely interrelated, which is the reason why the words 'understanding' and 'seeing' are synonymous. However, the visual system has its own rules. We can easily see patterns presented in certain ways, but if they are presented

in other ways, they become invisible...The more general point is that when data is presented in certain ways, the patterns can be readily perceived. If we can understand how perception works, our knowledge can be translated into rules for displaying information. Following perception-based rules, we can present our data in such a way that the important and informative patterns stand out. If we disobey the rules, our data will be incomprehensible or misleading. (Information Visualization, Second Edition, Colin Ware, Morgan Kaufmann Publishers, 2004, page xxi)

We must design our displays according to the rules of perception and cognition. These rules have emerged from an extensive body of research conducted over many years, beginning as early as the work of the Gestalt School of Psychology, founded in 1912. Gestalt is simply the German word for “pattern.” These early 20th century scientists were interested in how we see patterns and the laws that guide what we see. We can apply these and other hard-earned insights to graph design in the form of simple principles. Within the confines of this brief paper, I cannot present all of the principles that should be followed to communicate quantitative information in every circumstance, but I can list and explain seven core principles to set you on the path and whet your appetite for more.

Research Finding: Communication is most effective when you say neither more nor less than what is relevant to your message.

Principle #1: Display neither more nor less than what is relevant to your message.

When you wish to get your message across—any message—whether in conversation, in writing, or in a graph, irrelevant content is distracting. Don’t make people wade through meaningless visual content in your display to find what really matters. It has become common today, even in business graphs, to include all sorts of nonsense, such as cute pictures in the background or the addition of a third dimension to bars, lines, and pies. Despite good intentions (if you consider attempts to entertain or impress good), visual content of this sort is something that people’s eyes must scan and brains must process, without any payback, for it is meaningless. Extraneous content not only wastes people’s time, it makes it harder for them to get at the message.

The reverse is true as well. Don't design a display that doesn't contain everything people need to make sense of it. Include every piece of information that is part of your message—even notes to explain what might not be clear—otherwise you're communicating poorly.

This principle is broken in many graphs today by adding a 3-D effect to bars, lines, data points, and pies. The left-hand graph in Figure 3 illustrates this popular practice. What does the third dimension of depth that has been added to the bars mean? Does it encode any new meaning that resides in the data? No, it is pure visual fluff. Actually, it is worse than wasted visual content, for the gratuitous decoration and added dimension of depth make the data much harder to read than the corresponding graph on the right, which sticks to a 2-D display.

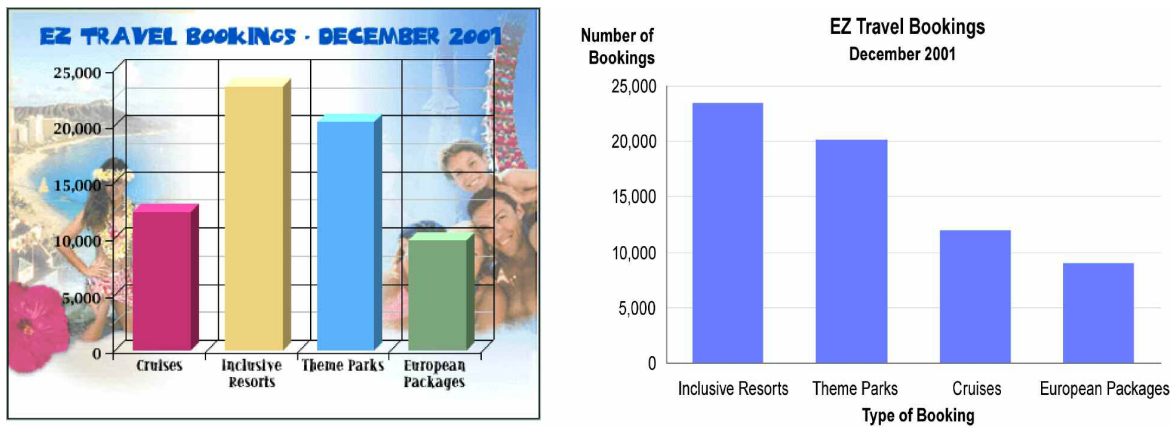


Figure 3: The meaningless 3-D effect in the graph on the left and the background images distract from the simple message that comes through clearly in the graph on the right.

Research Finding: People perceive visual differences in an information display as differences in meaning.

Principle #2: Do not include visual differences in a graph that do not correspond to actual differences in the data.

Graphs encode quantitative data in the form of objects, such as bars, lines, and data points, and visual properties of those objects, such as color to group objects. For example, Figure 4 distinguishes two sets of bars using distinct hues.




Figure 4: A visual property, hue in this case, was used to distinguish bookings and billings.

Because differences in visual properties, such as color, are used to communicate actual differences in the information itself, visual differences should never be used arbitrarily. When people notice visual differences, they try to discern the meaning of those differences. Don't confuse people and waste their time by including visual differences that are meaningless. Figure 5 shows a common example of how this rule is broken. What is the meaning of the different colors that appear on the bars? The answer is "nothing." We already know what the bars represent, because they are labeled as years along the X-axis. Meaningless visual differences such as this gratuitous use of color not only cause people to search for meanings that don't exist, but in this case they clutter the graph with an eye-assaulting abundance of color.

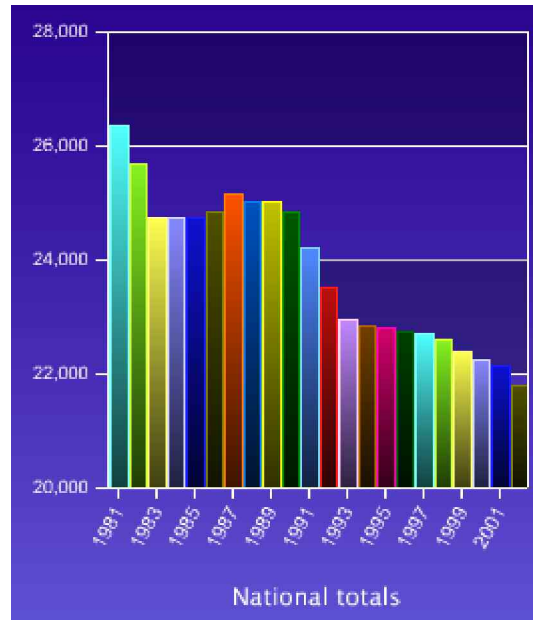


Figure 5: The differences in the color of the bars in this graph are meaningless.

Research Finding: The visual properties that work best for representing quantitative values are the length or 2-D location of objects.

Principle #3: Use the lengths or 2-D locations of objects to encode quantitative values in graphs unless they have already been used for other variables.

There is an important set of visual properties that are called “preattentive attributes” of visual perception. They are preattentive in that the process of perceiving them does not involve conscious thought; it is automatic and immediate. This includes such properties as an object’s length (for example, the length of bar in a bar graph), its 2-D location (for example, the position of a data point in a scatterplot), its size, its shape, its orientation, its hue, and so on. If objects in a graph vary from one another along one of these properties to a great enough degree to appear different, we see those differences immediately, without conscious effort. For example, if a single data point is orange in a scatterplot that contains 100 data points, 99 of which are black, the orange dot will stand out as different. We can use this knowledge to intentionally make particular items in a graph stand out as different or important.

Of the full set of preattentive attributes, a few are perceived quantitatively. By this I mean that we perceive differences between varying expressions of a visual property (for example, length, exhibited as long bars, short bars, medium-length bars, etc.) as greater than or less than one another. Apart from these preattentive attributes, those that are not perceived quantitatively are simply seen as different, such as the different hues of black, green, blue, orange, purple, and so on. Two of the preattentive attributes that are perceived quantitatively are also perceived with a fair amount of quantitative precision: length and 2-D position. It is easy to determine, when looking at the two sets of objects in Figure 6, that the taller bar is about twice as long as the shorter bar, and that the higher dot is about three times as high as the lower dot.




Figure 6: Length and 2-D location are the two preattentive visual attributes that work best for encoding quantitative values.

It isn't accidental that the primary means of encoding quantitative values in graphs involves the use of length, as in the length of bars, and 2-D location, as in the position of data points. When lines are used to encode values in graphs in the form of a line graph, the 2-D locations where the data points are positioned along the line are what encode their values. Other attributes that we perceive quantitatively, but unfortunately with a low degree of accuracy (see Figure 7), include size (the 2-D area of an object, such as slices of a pie chart or of circles in a bubble chart) and color intensity (ranging from light to dark or pale to fully saturated).

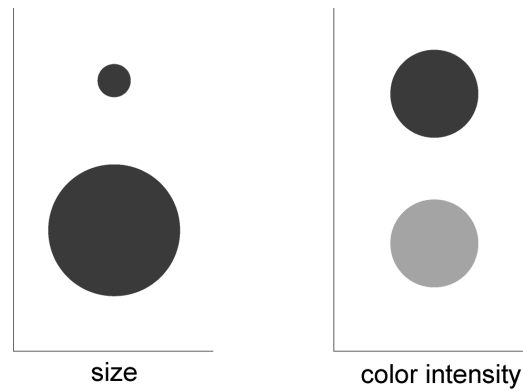


Figure 7: These visual attributes—size and color intensity—are perceived quantitatively, but imprecisely.

Because quantitative differences in these particular properties are not perceived accurately, at best they can provide an approximation of quantity, and should only be used when this is sufficient for your message. The bubble chart in Figure 8 uses the 2-D location of the circles to encode marketing expenses and associated sales revenues for individual states, but the size of the circles has also been used to encode profits—the larger the circle the greater the amount of profit. Size or color intensity may be used to encode quantities in a pinch, but only if approximate values will do the job.




Figure 8: This graph uses the 2-D location of circles (bubbles) to encode marketing expenses and sales revenues, plus the size of the circles as well to encode profits.

Research Finding: People perceive differences in the lengths or 2-D locations of objects fairly accurately and interpret them as differences in the actual values that they represent.

Principle #4: Differences in the visual properties that represent values (that is, differences in their lengths or 2-D locations) should accurately correspond to the actual differences in the values they represent.

Graphs are sometimes intentionally designed to deceive, to misrepresent the truth by visually encoding values in a way that does not correspond to the actual values themselves and the differences between them. Even more often, however, people unintentionally misrepresent data in this manner, simply because they don't understand this principle and how to follow it. The most common way that this occurs involves bar graphs with quantitative scales that don't begin at zero. Because the lengths of bars encode the values they represent, the full length of the bar must be displayed, beginning from zero, for the values to be encoded properly. Take a look at Figure 9. Notice that actual sales in the East region appear to be twice as great as planned sales, but in fact, this is far from the truth. Actual sales are only 5% greater than the plan. When you use a graph to communicate, people should be able to look at the graphical representation alone to compare differences in values. If the graph doesn't support this operation, what's the point of using a graph?

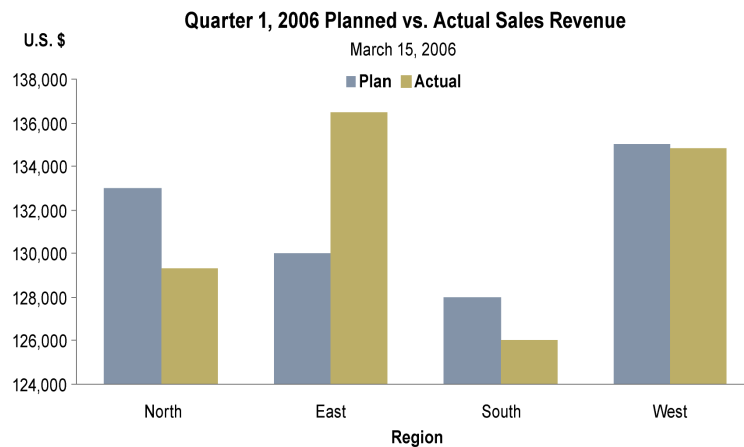


Figure 9: The differences in the lengths of these bars do not accurately reflect the differences in the values they represent. Bars must begin at a baseline of zero.

Research Finding: People perceive things that appear connected as wholes and things that appear disconnected as discrete.

Principle #5: Do not visually connect values that are discrete, thereby suggesting a relationship that does not exist in the data.

Values that we display in graphs are sometimes intimately related to one another and sometimes they are discrete. The way we visually display these values should make it easy to see, without effort, this distinction. Figure 10 illustrates one of the most common ways that this principle is broken. In this graph, lines connect values, suggesting a relationship between them that doesn't exist. The regions North, East, South, and West are discrete, so values that measure something going on in these regions should be displayed as discrete. Connecting them with a line is misleading. Doing so forms a pattern of upwards and downwards slopes that are utterly meaningless.




Figure 10: This graph illustrates a misuse of lines, because the values that the lines connect are discrete, not intimately related.

In contrast to what you've just seen, Figure 11 illustrates an appropriate use of lines to encode values in a graph. Measurements of something at equal intervals across time are intimately related to one another. A line does a wonderful job of displaying this connectedness in the values, using the slopes formed by the connection of one value to the next as a meaningful representation of change.




Figure 11: This graph illustrates an appropriate use of lines, for connection from one value to the next accurately represents the nature of the data.

Bars may be used to encode values across time as well, but should only be used when the graph is used to help people attend to and compare individual (discrete) values, rather than to see the overall shape of change through time.

In Figure 12, notice how the use of bars makes it easy to focus on individual quarters independently and to compare budgeted and actual expenses in a given quarter—something lines would not support quite as well.

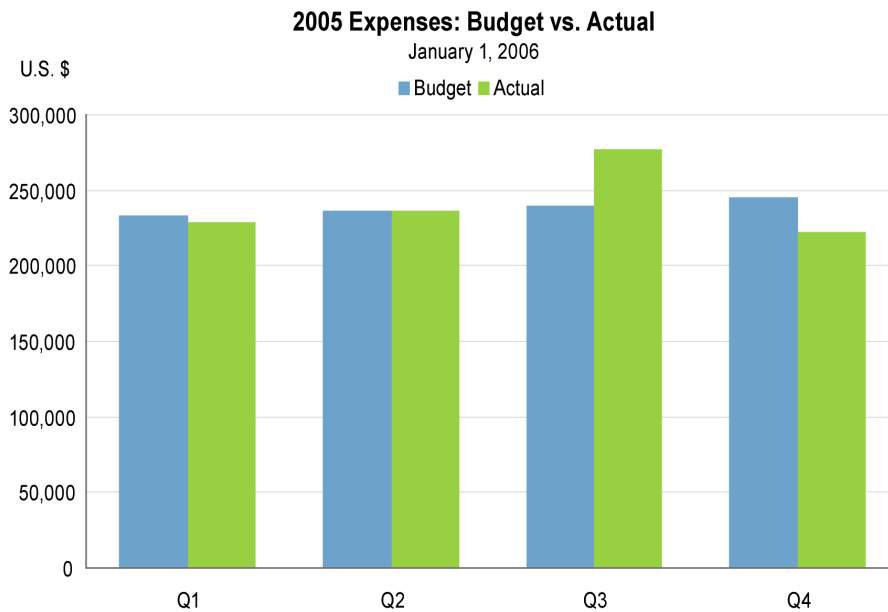


Figure 12: Bars may be used to encode time-series values, but only when people should focus on and compare individual values, such as budgeted and actual expenses.

Research Finding: People pay most attention to and consider most important those parts of a visual display that are most salient.

Principle #6: Make the information that is most important to your message more visually salient in a graph than information that is less important.

Not all information is created equal. It is often the case that some information is more important to your message than other information. You can communicate this fact in a graph by making those items that are most important more visually dominant (salient). It is your job, if you wish to communicate effectively, to direct people's eyes to the most important parts of the display, so they are sure to adequately focus on them.

The title of the graph in Figure 13 clearly states at least part of its purpose: to highlight what happened in March. This purpose was visually reinforced by making the bars in March more salient than the others, in this case by placing borders around them.




Figure 13: Data in the month of March has been highlighted through the use of borders around the bars.

Many visual properties can be used to make something stand out as important, including a brighter or darker color, or even a color that is simply different from the norm. Just as you might raise your voice and speak more slowly to emphasize something that you are saying, you should use visual means to emphasize particular parts of your message as well.

Research Finding: Short-term memory is limited to about four chunks of information at a time.

Principle #7: Augment people's short-term memory by combining multiple facts into a single visual pattern that can be stored as a chunk of memory and by presenting all the information they need to compare within eye span.


Short-term memory, also known as working memory, is similar to RAM (random access memory) in a computer. It is the place where information is stored while it is being processed. Unfortunately, unlike the large amount of data that can be stored in RAM for ready access, short-term memory has an extremely limited capacity. Only about four chunks of memory can be stored in short-term memory at any one time. Information that is stored temporarily in short-term memory comes either from the outside world through our five senses or from long-term memory, where it is stored until it is needed.

Computers have an advantage over the human brain in terms of fast-access working memory capacity, but our brains excel by comparison in their ability to combine multiple pieces of information into a single chunk of memory. Remember the graph in Figure 2? Each of the two lines in this line graph combines 12 different sales figures, one per month, into a single pattern of upward and downward sloping line segments. When encoded in a visual pattern such as this, these 12 numbers can be stored together as a single chunk of information in short-term memory. By presenting quantitative information visually as patterns, more information can be simultaneously stored in short-term memory, thereby augmenting it in a way that extends people's ability to think about it.

When a person examines information on a computer screen or the page of a printed report, a limited amount of the information can be held in short-term memory if comparisons must be made to information on another computer screen or page. Information that is never attended to never gets stored in short-term memory, and even if several chunks of information are attended to, only around four will be remembered when the person moves from one display to the next. This makes comparisons difficult. You can augment people's short-term memory, however, by placing everything that needs to be compared within eye span, so it is readily available for rapid swapping in and out of short-term memory as it is being processed. A popular example is a dashboard that has been properly designed to display everything within eye span on a single screen that people must monitor for rapid processing and comprehension.

A FINAL WORD

What you've read in these pages is merely an introduction to some of the core principles that must be followed to effectively communicate quantitatively information in graphs. Don't be disheartened, however, for the skills that are required are all quite easy to learn. Continue to read the white papers in this series and you will build these skills one concept at a time.



Stephen Few has worked for 24 years as an IT innovator, consultant, and educator. Today, as Principal of the consultancy Perceptual Edge, Stephen focuses on data visualization for analyzing and communicating quantitative business information. He provides consulting and training services, writes the monthly Business Visualization newsletter, speaks frequently at conferences like *TDWI* and *DAMA*, and teaches in the MBA program at the University of California in Berkeley. He is the author of two books: *Show Me the Numbers: Designing Tables and Graphs to Enlighten* and a new book entitled *Information Dashboard Design: The Effective Visual Communication of Data*. You can learn more about Stephen's work at www.perceptualedge.com.

