

VISUAL PATTERN RECOGNITION



MEANINGFUL PATTERNS IN
QUANTITATIVE BUSINESS
INFORMATION

STEPHEN FEW
PERCEPTUAL EDGE

NOVEMBER 25, 2006

THE EXTRAORDINARY HUMAN ABILITY TO SEE PATTERNS

All creatures have the ability to sense the surrounding world, but in various ways and degrees. You might envy the bloodhound's exceptional nose, but humans possess visual prowess that (although it doesn't match the eagle's in distance) is unsurpassed in the ability to detect and make sense of patterns. Our eyes and brains work as a team to discover meaningful patterns that help us make sense of the world. When we look at the pointillist painting by Georges Seurat called "The Lighthouse at Honfleur" (Figure 1), rather than a random arrangement of dots, we recognize patterns that form a beautiful seashore.



Figure 1

What initially evolved to help us survive in fundamental ways (such as distinguishing nourishing versus toxic plants) now enables us to detect and make sense of patterns in abstract information as well. The power and subtlety of visual perception can be applied to the analysis of information, such as your company's sales, to tease out and make sense of the patterns that reveal what's happening and why. The ability to sift out and interpret meaningful patterns is a required skill for effective data analysis.

Visual perception is a process that consists of multiple stages. Pattern perception is one of these stages. Colin Ware of the University of New Hampshire explains the process as three sequential stages:

Stage 1: Parallel processing to extract low-level properties of the visual scene

Stage 2: Pattern perception

Stage 3: Sequential goal-directed processing

(*Information Visualization: Perception for Design*, Second Edition, Colin Ware, Morgan Kaufmann Publishers, San Francisco, 2004, pages 20-22).

In the first stage, billions of neurons in our eyes work simultaneously to extract features from our field of vision and then pass what they find to the primary visual cortex in the back of our brains. This stage of the process occurs rapidly and without conscious thought (that is, pre-attentively). The visual attributes that we detect during this stage of the process are called pre-attentive attributes of visual perception. They include features such as the location, size, shape, color, orientation, and texture of objects. Because perception of these attributes is powerful and immediate, certain aspects of what we see can be easily discerned and even jump out to grab our attention, such as dark objects in a field of mostly light objects or tilted objects in a field of mostly vertical objects. These pre-attentive attributes of visual perception assist us in grouping and arranging objects in ways that form patterns that can be discerned in the next stage of processing.

In the second stage, our initial perception of the visual field is divided into regions and simple patterns, such as groups of objects based on their proximity to one another or a series of objects that forms a linear slope based on their alignment (Figure 2).




Figure 2

In the third and final stage, we actively try to make sense of what we see—a goal-directed process—which involves attending to particular patterns. Representations of these patterns are stored temporarily in working memory (a.k.a. short-term memory), which makes them readily available to consciousness, so we can think about them. This process can go on for some time through a series of visual queries. We see something that catches our interest and provokes a question, which we pursue by searching through the patterns in our visual field (a visual query) to satisfy our interest and answer the question.

Our perception of patterns in the objects that we see is fundamental to the sense-making process. Pattern perception supports abstract thinking. We can visually encode abstract information, such as financial information, which forms patterns that allow us to explore and make sense of that information, leading to insights that might never occur if the data were examined in any other way.

The richness and speed of visual perception is a gift we cannot afford to ignore. The ability to encode data visually in ways that form perceptible and meaningful patterns is a critical skill that must be nourished if you hope to understand more than the most obvious facts that live in your data.

EFFECTIVE WAYS TO VISUALLY ENCODE PATTERNS

Visual representations of data are effective if they clearly, accurately, and efficiently communicate the meanings contained in the data. The efficacy of a graphically-encoded pattern can be tested by first understanding the meaning that it ought to encode and then checking to see if people grasp that meaning when they view it.

Encoding Relationships

Many characteristics of quantitative business data are potentially meaningful, depending on the nature of the information and the questions one is asking. Nevertheless, most conditions of interest can be encoded using a limited set of visual patterns.

When we examine quantitative business data, we are always looking for relationships between the values. Here's a list of the most common relationships:

Relationship	Example Question
Rank	In what order are the company's departments based on number of employees from high to low?
Part-to-whole	What portion of the market does our company command, and how does that compare to our competitors?
Time-series	Is the traffic on our website increasing or decreasing?
Deviation	To what degree do each department's expenses vary relative to the budget?
Distribution	What is the range of our employees' salaries, and how many employees fall into each subset of that range in increments of \$10,000?
Correlation	Is there a relationship between the amount of money we spend on marketing and the resulting sales revenues?

This list of relationships is certainly not comprehensive, but relatively few of the questions that we ask about quantitative business data fall outside of it. Most business questions lead us to examine one or more of these relationships.

Patterns that Reveal Relationships

Even though the list of specific conditions in business data that are worth examining is limitless, the patterns that we must detect and make sense of to analyze these conditions are limited. Graphical representations only need to handle a limited set of relationships, and the number of meaningful patterns that these relationships can form is limited. I'll illustrate this by taking one of the relationships mentioned previously—time-series (change through time)—and listing the most

common patterns of interest to business analysts:

Pattern	Example Question(s)
Patterns that reveal the relative magnitude of separate sets of values	Which set of values is the higher or lower?
Patterns that reveal the general nature of change	Are the values going up, going down, or remaining relatively flat?
Patterns that reveal the rate of change of the comparative rate of change	Are sales increasing at a progressively greater rate over time? Which of these two sets of values is decreasing at a greater rate?
Patterns that reveal seasonal or cyclical events	Are there particular times of the year when sales peak?
Patterns that reveal a possible temporal co-variation between two sets of values	Are the costs of radio ads a reliable leading indicator of sales?
Patterns that reveal when one set of values surpasses or falls below another	At what point during the year did the sale of widgets surpass the sale of gadgets?
The stability or volatility of a set of values	Has our Web traffic remained stable over the year, or has it fluctuated significantly?

Using time-series analysis as an example, you can see that the number of meaningful patterns for most analytical tasks is hardly overwhelming. Consequently, it should be possible to develop a comprehensive graphical vocabulary that can be used to represent these patterns.

Visual Building Blocks for Encoding Quantitative Data

Certain pre-attentive attributes of visual perception—the basic building blocks of visual objects and patterns—work exceptionally well for encoding quantitative values because they map intuitively to our concept of greater than or less than (quantitative comparisons) and we are able to make these comparisons fairly accurately. The most powerful attribute of this type is 2-D location. The position of an object in our field of vision arranged in 2-D space (up/down, right/left) is ideal for encoding quantity.

In Figure 3, the fact that the dot on the right is greater (because it is higher) than the one on the left requires little training to comprehend. Furthermore, we can also see that the dot on the right is twice as great (twice as high) as the one on the left.




Figure 3

The Efficacy of Simple XY Graphs

It is not hard to understand why Rene Descartes' invention of the graphical coordinate system for encoding values on a 2-D plane relative to two perpendicular axes (one horizontal, labeled X, and one vertical, labeled Y) caught on and has continued to be the prominent means of visually representing quantitative values. It works exceptionally well.

In addition to 2-D location, the other highly effective attribute for encoding quantitative values, which also works in XY axes graphs, is the length or height of a simple object. This is how the bars in a bar graph work. With bars, we can easily compare values to one another by simply comparing the heights of the bars.

Although there are a few exceptions, most of the effective ways that have been developed to graphically encode quantitative values use the 2-D location or length of an object in an XY axes graph. You might be wondering why I haven't mentioned 3-D graphs, despite the fact that 3-D location (depth perception, which is possible because of our stereoscopic vision) is another pre-attentive visual attribute. Unfortunately, 3-D graphs are rarely able to display business information in a way that communicates effectively because of technological limitations. Computer screens and paper are both flat, 2-D surfaces. Consequently, they cannot be used for true 3-D displays. What they can do is use visual cues (for example, perspective drawing) to fool our eyes into thinking we are seeing 3-D space. This illusion, however, creates a visual experience of depth that is far from perfect, especially on a computer screen, often producing perceptually confusing representations of depth when applied to graphs.

The Best Objects for Encoding Quantitative Data

Graphs encode values as objects that appear in the plot area. Most graphs use one or more of only three objects to encode quantitative data: points, lines, and bars. These simple objects work well, primarily because they are so simple, but also because they form images that make sense to our brains. Each is particularly good at representing particular types of data for particular analytical purposes.




Figure 4

Meaningful Patterns

Meaningful patterns in quantitative data fall into three general categories:

- **Large-scale patterns** (a.k.a. trends). These are patterns that reveal what's going on in general (that is, as a whole). For example, sales have trended downwards over the course of the year as a whole.
- **Small-scale patterns**. These are patterns that reveal what's going on in specific subsets of data. For example, sales and customers' ages correlate exclusively among the elderly.
- **Exceptions**. These are values that appear outside of what's normal or acceptable. For example, out of 100 orders that were shipped, three were not shipped until 10 or more days from the time the order was received, whereas all other orders shipped within five days.

Points

Small points in the simple shape of circles (dots), squares, triangles, etc, are ideal for pinpointing the precise location of individual values in a graph. If you wish, you can make them very small and display many of them in a single graph, which is sometimes precisely what's needed. Points show their greatest strength in the form of a scatterplot, which is primarily designed to show the correlation (or lack of one) between two sets of quantitative values. For example, you might want to know if there is a correlation between the seniority of sales people (how long they've been working in sales) and profits that result from their sales. To test this, you could represent each salesperson with a single point on a scatterplot, positioning it along the X axis to encode that person's years of sales experience, and along the Y axis to encode profits resulting from that person's sales activity. In this way, a single point is able to encode two very different quantitative values for a single entity.

Let's look at one possible result (Figure 5).




Figure 5

This picture tells us that the longer a person has worked in sales, the greater the profits that have been earned, but not forever. It appears that sales people reach a point when the profitability of their work starts to decline. (By the way, this data is completely made up to illustrate the kind of story that can be told using points to encode quantitative data, so don't worry if you work in sales and are approaching your 10th anniversary.)

Here's a partial list of meaningful patterns that can be found when points are used to encode values:

Category	Pattern	Examples
Linear trends	Points are arranged in a way that forms a pattern that looks like a straight line	<ul style="list-style-type: none"> • Points are going up from left to right, indicating a positive correlation between employee's heights in inches and their salaries in dollars • Points are going down from left to right, indicating a negative correlation between per capita wealth and infant mortality
Non-linear trends	Points are arranged in a way that forms a pattern that looks like a curved line	<ul style="list-style-type: none"> • Points are going up from left to right at an ever-increasing rate, indicating not just growing sales, but also a growing rate of sales • Points are going down from left to right, indicating fewer device failures as temperature rises until the temperature reaches 100° and failures begin to increase
Concentrations	Dense sets of points appear in particular areas of the graph	<ul style="list-style-type: none"> • A large concentration of points appears in the upper left corner of the graph, indicating a large number of small orders with large profits • A large collection of points appears in the center of the graph, indicating that most orders are moderate in both size and profits
Clusters	Sets of points appear to be set apart from other points	<ul style="list-style-type: none"> • A small group of isolated points appears in the lower right corner of the graph, indicating a distinct group of large orders with low profits
Gaps	Area in the graph where no points appear in the midst of surrounding points	<ul style="list-style-type: none"> • In a scatterplot that correlates home sales amounts and the number of days on the market, there is a glaring gap between 45 and 55 days where no homes appear
Randomness	Points are arranged randomly in the graph, without a discernible pattern	<ul style="list-style-type: none"> • A random scattering of points throughout the graph indicates an unexpected lack of correlation between marketing expenses and sales
Exceptions	Points that stand out as different from the norm	<ul style="list-style-type: none"> • A single point appears alone in the lower right corner of the graph indicating a customer who places many orders but receives a low discount

Lines

By connecting points with a line, we can transform a display of individual values into a story of transition from one value to the next. Using a line to connect the points allows us to readily perceive an actual connection that exists between the values, making it easy for our eyes to trace the transition from one value to the next. This is especially useful when the values measure change through time, which can be seen in the up or down, dramatic or subtle slopes that encode temporal transitions. Only two of quantitative relationships that I listed previously involve values that are intimately connected to one another: time-series and distributions. I talked a bit about time-series relationships earlier, but let's look at them again, this time focusing on the visual patterns that are meaningful, which lines can display when encoding values across time.

Pattern	Examples
Unidirectional lines	<ul style="list-style-type: none"> • A line that encodes monthly sales is heading down during the course of the year • A line that encodes a company's total headcount has remained flat for the last five years, despite ups and downs in particular departments
Curved lines	<ul style="list-style-type: none"> • Sales of a product increased rapidly soon after its introduction but began to increase at a much slower rate after a few months and eventually began to decrease after a few years
Repeating line patterns	<ul style="list-style-type: none"> • Sales are at their lowest in the first month of each quarter, rise somewhat in the second month, and hit their peak in the final month, then drop again at the beginning of the next quarter to begin the cycle again
Co-variation among multiple lines	<ul style="list-style-type: none"> • Revenues and profits went up and down together throughout the year • Sales always increased in relation to the number of marketing mailings, but didn't register until five days after the mailings
Intersecting lines	<ul style="list-style-type: none"> • In June of the year, the decline in coffee sales and the increase in soft drink sales intersected and changed positions in a ranking of sales by product
Smooth vs. jagged lines	<ul style="list-style-type: none"> • Expenses in the Human Resources department tended to change slowly and smoothly compared to the Information Systems department, which exhibited large and rapid increases in expenses from time to time
Exceptions	<ul style="list-style-type: none"> • Sales in the month of August decreased far below the norm

Different from time series, when we examine distributions, we focus primarily on three characteristics of the data: 1) the full range across which the full set of values extends (spread), 2) the central tendency of the values, which serves as a single-value summary of the full set (for example, the median), and 3) the shape of the distribution, which informs us where items fall at various locations across the distribution.

Pattern	Examples
Symmetrical	Most employees are between 40 and 49 years of age, followed by those in the 30 to 39 and 50 to 59 age groups, with only a small number in both the 20 to 29 and the 60 to 69 age groups.
Skewed	The median salary for employees is only \$32,000, even though salaries begin at \$20,000 and extend to \$100,000.
Uniform	The number of visitors to the website is evenly distributed across all age groups.
Multi-modal (multiple peaks)	Percentage contributions to the 401K plan were especially high among employees in two distinct salary ranges: \$20,000-\$30,000 and \$70,000-\$80,000.
Gaps	The distribution of orders by dollar amount exhibited a strange gap in the \$50-\$60 range
Exceptions	All but one salary fell within the range of \$20,000 to \$100,000, which was the CEO's salary of \$1,000,000.

Bars

Bars are less versatile in their ability to form meaningful patterns, but they support one analytical task exceptionally well: comparisons of magnitudes. Because bars are so salient, standing firm in the graph like a series of monuments, they draw our eyes to individual values and make it as easy to compare the magnitudes as it is to compare the heights or lengths of two bars to one another. Consequently, in a bar graph, it is easy to rapidly find the lowest and the highest values in the bunch.

This is especially useful when you want to see how a collection of values are ranked from highest to lowest or vice versa. Nothing reveals a ranking relationship better than a series of sorted bars. Bars work well for part-to-whole relationships also.

People generally use pie charts to display part-to-whole relationships, such as each product's percentage contribution to total sales, but it is much harder to compare the 2-D areas of pie slices (a comparison that visual perception can only approximate) than it is to compare the lengths or heights of bars (illustrated in Figure 6).

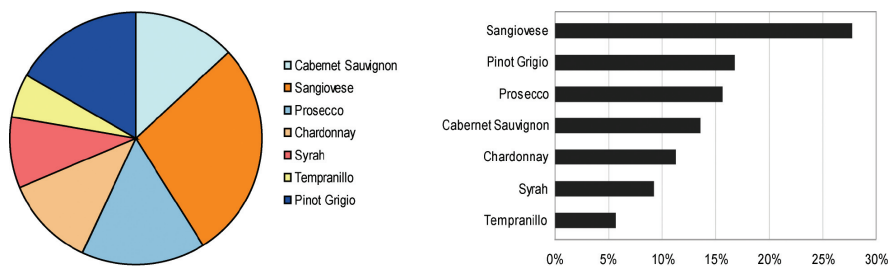


Figure 6

Bars may always be used in place of a line when you wish to emphasize the individuality of the values and to compare one to another, instead of examining the overall shape of the values. Consider a distribution relationship involving the number of days that it took to ship the orders that were received in a given month. Figure 7 shows two ways to encode the same distribution: one that uses a line and one that uses bars. Although they display the same information, the picture that they paint is somewhat different.





Figure 7

The line emphasizes the overall shape of the distribution. The bars show the shape as well, but not as directly and clearly as the line. What the bars do better than the line, however, is feature the individual values, in this case the number of orders that shipped in a particular number of days, making it easier to compare individual values, such as the number of orders that shipped in three days versus four days. Whether lines or bars works better depends on what you're doing with the data—what you're looking for.

FINAL WORD

Our visual sense is highly tuned to spot patterns, and this natural ability can be trained to achieve even greater levels of acuity. Improvements result primarily from learning which of the many possible patterns in data are meaningful to your work and then also learning—through practice—to spot these patterns quickly, even in a dense forest of visual noise. As with all skills, practice in pattern recognition is required to move from average ability to true expertise.



Stephen Few has worked for over 20 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 Visual Business Intelligence Newsletter, speaks frequently at conferences, and teaches in the MBA program at the University of California at 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.

