perceptual edge

Use-Based Types of Quantitative Display

Stephen Few, Perceptual Edge
Visual Business Intelligence Newsletter
January/February/March 2012

We present quantitative information for various purposes. Each purpose requires that we design displays in particular ways to achieve relevant outcomes. What are these purposes? If we exclude quantitative displays that are used for mathematical purposes (e.g., a mathematical proof), we can identify five fairly distinct but at times interrelated purposes for which we display quantitative information:

1. Lookup
2. Narrative
3. Monitoring
4. Exploratory Data Analysis
5. Prediction

My purpose in this article is to describe these types and provide a few basic guidelines for their effective design.

Purpose #1: Lookup

Lookup displays, often called operational or production reports, are used to look up facts that are needed to do one's work. As such, the information is usually presented in tables (i.e., arranged into columns and rows). As such, they communicate primarily through verbal channels (words and numbers that must be read) rather than visually (graphics).

This is the type of display that business intelligence (BI) products support primarily. This is ironic, because lookup displays do not directly support the primary purpose of BI: improved decision making based on evidence derived from data. Looking up facts that are needed to do one's job is an operational task that rarely requires decisions, for the appropriate responses are usually prescribed.

The fundamental requirement of lookup displays is that they present information in a way that makes it easy and efficient to find the individual facts that you need. Tables that arrange information into logical groups and sort large lists of items alphabetically usually serve this purpose best. Graphics are rarely useful for lookup.

Purpose #2: Narrative

Narrative displays are used to inform, explain, or persuade. Narrative (story) is an explanation. It answers questions: "How did we get to this state?" "What caused this to happen?" Its primary goal is communication, resulting in understanding and informed action. The meanings that we construct and express to make sense of our world—whether in a grand sense to live our lives or in trivial matters that we must understand to achieve minor goals—are stories. To the degree that we all strive to make sense of things, reflect on the process, and pass what we find on to others, we are all storytellers. Being storytellers is in part what makes us uniquely human.

In the case of quantitative narrative, a particular story needs to be told that is based entirely or in part on quantitative information. Displays of this type can combine words, numbers written as text, and images. They

are sometimes presented live in meetings, but are more often presented in documents. Infographics involving numbers are a written form of quantitative narrative.

The effectiveness of a narrative display can be determined by measuring the degree to which the intended understanding and desired outcomes were produced. All of the information that people need should be presented in a manner that is clear, accurate, and understandable. The presentation must also be convincing as well, if persuasion is the goal.

Purpose #3: Monitoring

Monitoring displays support the following purposes: 1) Maintaining ongoing awareness of what's going on and how well things are doing, and 2) reporting situations that require prompt action, either to correct a problem or take advantage of an opportunity. As such, displays of this type may prompt and support decisions.

Monitoring displays can be further divided into two subtypes: 1) those that are updated on a near-real-time basis, and 2) those that are updated less often, usually daily. Displays that are updated less often than daily rarely serve the monitoring purpose in the true sense, but more often serve lookup or narrative purposes. Displays that are updated on a near-real-time basis must usually present information more simply than those that are updated daily, because it is more difficult for people to process a large amount of information when it is constantly changing. There are other ways in which near-real-time displays should be designed differently, including the ability to temporarily halt updates so the information can be studied more closely without distraction.

A monitoring display works best when it presents information in a way that is closely aligned with the mental model (a.k.a., conceptual model) of the person or group that uses it. For example, people who supervise a manufacturing process hold in their heads a model that includes each step in the process, each piece of machinery and worker who's involved in the process, and the relationships between them. They know how a particular condition in one part of the operation will affect others down the line. Increases in a supervisor's expertise relate directly to the sophistication of his or her mental model. Greater expertise coincides with more complex mental models. For a monitoring display that is designed for manufacturing supervisors to work, it must present information in a way that corresponds to their mental models. The greater the correspondence, the greater the ease with which a supervisor could use the display to update his or her awareness of what's going on in the manufacturing process.

Monitoring displays also work best when they present information graphically whenever possible. Graphical displays, when properly designed, can include a great deal more information on a screen than text-based displays. Graphical displays can also be perceived and understood more rapidly than text-based displays, because the information can be processed in parallel by the visual cortex, rather than serially by the verbal processing part of the brain.

Purpose #4: Exploratory Data Analysis

Exploratory data analysis (EDA) displays support the following two purposes: 1) data exploration to find facts of potential interest in a set of data, and 2) data sensemaking (a.k.a., data analysis or descriptive statistics) to determine what the facts mean. Understanding is the immediate goal of these activities. The ultimate goal is to improve decisions and the actions that follow.

EDA tools rely heavily on data visualization and must support the following features especially well:

1. *Data Interaction:* EDA requires rich, flexible, rapid, and fluid interaction with data.
2. *Data Comparison:* All uses of data visualization rely on comparisons to represent data in context, but the ability to make rich comparisons is especially central to EDA.
3. *Multi-perspective Views:* EDA requires that we view data from multiple perspectives simultaneously to spot relationships and connections—that is, to connect the dots.
4. *Statistical Functions:* Although statistics inform all uses of quantitative data visualization, ready access to a basic set of statistical functions is especially critical to EDA.

5. *Rich Access:* Most uses of data visualization rely on data stores that are readily available and contain the information that's needed, but with EDA, the full set of information that's needed is often discovered during the process and involves data that cannot be accessed from a single, well-structured source. Consequently, EDA tools must be able to access data from many sources and to combine data on the fly in efficient ways.

Purpose #5: Prediction

Displays of this type are used for predictive analysis (a.k.a., predictive statistics or "what if" analysis) to anticipate what might happen in the future given specific conditions, based on an understanding of what has happened in the past. Effective predictive analysis requires an understanding of probability theory and the ability to build reliable probability models.

Statisticians often develop predictive models that involve no use of visualization. Users enter input variables, the computer runs the mathematical model in a manner that is hidden from view, and results are produced. Models of this type do not engage the user in thinking about the process. Users must trust a model that was developed by someone else because they have no view into its inner workings. This may be appropriate at times, but when it is useful for the user understand the model, this can be accomplished by visualizing the effects that input variables have on all of the intervening variables, not just on the final result. Visual predictive models can be built such that users can manipulate variables at will and immediately see the effects on all of the parts that work together to produce particular outcomes. Visual models assist us in thinking, which builds our skills and our ability to judge the merits of the results.

An Integrated View

EDA and predictive analysis both belong to what we currently call "analytics." Whether we're exploring and making sense of what is or using what is to make sense of what might be, we are engaged in analytics.

EDA precedes all of the other uses of quantitative displays. We engage in EDA to find the stories that dwell in numbers and then determine what they mean. Before we can tell the stories, we must find and understand them. Before we can monitor performance, we must identify the data that measures performance most effectively. Before we can predict what will likely happen given specific conditions, we must understand historical behavior.

A single quantitative display can be used for multiple purposes. A common combination involves one that is used primarily for monitoring performance, but when conditions appear that require a response, can then be used to lookup relevant facts or to perform analysis to determine how to respond. Sometimes this requires a single display to be designed to support multiple purposes as well as possible, but is usually better served by designing separate displays that are linked for easy navigation from one to another and back again.

Factors other than the purpose for which quantitative displays are used can also influence their design, including the display platform (e.g., a large screen on a desktop vs. a smaller mobile device such as a smartphone) and the skills of the user. A display's purpose, however, is the primary factor that informs its design.

The State of Technology

No single software product serves all of these purposes equally well. The best data visualization tools today tend to focus primarily on one or perhaps two of these purposes. Tableau, Spotfire, and SAS JMP all focus primarily on EDA. The charting capabilities of Excel, sometimes combined with PowerPoint, are best reserved for narrative. Several products work primarily for building monitoring displays and simple, prebuilt analytical applications, including QlikView and Panopticon. It is difficult for software vendors to build a single product that serves all of these purposes well. This difficulty is due, in part, to the fact that these different purposes tend to serve distinct groups of people with quite different interests and skills. If you need to look up facts to do your work and you never analyze data, tell stories that dwell in numbers, monitor performance, or try to predict future outcomes, you probably wouldn't want a tool that serves these other purposes because they would be distractions, constantly looming in the periphery of your vision. A simple tool for building tabular

reports would probably work best. On the other hand, if you spend most of your time exploring and analyzing data, but also must explain what you've found to others, you would benefit from an EDA tool that could also be used to produce presentation quality output that you could easily incorporate into slides or written documents. Some features can be added to extend the reach of a data visualization tool into adjacent territory without compromising the tool's ability to serve its primary purpose. Smart vendors are looking for ways to expand the purposes for which their products can be used while also respecting the natural boundaries that could only be exceeded by compromising the effectiveness of their products.

In future articles, I will build on this content by listing and describing the full set of features that data visualization products require to support each of these purposes effectively. These articles will provide guidelines for evaluating the merits of products when you're shopping for one to use for specific data visualization purposes.

Discuss this Article

Share your thoughts about this article by visiting the Use-Based Types of Quantitative Display thread in our discussion forum.

About the Author

Stephen Few has worked for over 25 years as an IT innovator, consultant, and teacher. Today, as Principal of the consultancy Perceptual Edge, Stephen focuses on data visualization for analyzing and communicating quantitative business information. He provides training and consulting services, writes the quarterly *Visual Business Intelligence Newsletter*, and speaks frequently at conferences. He is the author of three books: *Show Me the Numbers: Designing Tables and Graphs to Enlighten*, *Information Dashboard Design: The Effective Visual Communication of Data*, and *Now You See It: Simple Visualization Techniques for Quantitative Analysis*. You can learn more about Stephen's work and access an entire library of articles at www.perceptualedge.com. Between articles, you can read Stephen's thoughts on the industry in his blog.