

Today's dashboard environments provide developers with a rich set of indicators and charts, ranging from bullet graphs to sparklines and bar charts to bubble graphs. Additionally, numerous formatting options such as splines and 3D displays are available, adding even more complexity. To the casual observer, chart types and options may appear to be largely interchangeable and are simply left to the discretion of the developer. But as the art and science of effective dashboard design has evolved, data visualization gurus have established some guidelines regarding which charts and indicators to use for various types of data and how to present those charts most effectively. This issue is receiving increasing media attention as dashboards migrate to the limited display "real estate" found on mobile devices, placing a premium on efficient dashboard based information delivery.

In this series of Dashboard Insight articles we'll look at some popular chart types and describe how to use them most effectively in your dashboards, as well as pitfalls to avoid. In Part 1, we'll discuss effective dashboard indicators. Part 2 will deal with basic charts and in Part 3 we'll move on to charts that display more than one data element. In our final article, we'll present some special purpose charts. With this collection of indicators and charts, your dashboards will stand head and shoulders above the competition, and leave your users with an unambiguous and actionable view of their data.

Part 1 - Using Dashboard Indicators

Before deciding which indicator to use for displaying a particular data entity, it helps to classify indicators based on the information content they are capable of displaying. Binary indicators can only communicate two information states such as acceptable/unacceptable, below/above goal, problematic/normal. Categorical indicators take this a step further allowing the dashboard designer to encode an ordered set of categories such as poor, satisfactory, excellent. At the top of the information pyramid we find compound indicators that can communicate not only directionality and order, but actual quantities.

Since compound indicators typically display the most information per square inch of dashboard space, it might seem logical to use them as often as possible. But understanding your target audience plays an even more important role in choosing the appropriate combination of indicators. A simple binary indicator can motivate an assembly line machine operator to take immediate action far more effectively than an indicator formatted with numerical values. On the other hand, an operations manager may need to see actual quantities and historical context when making staffing decisions.

Dashboard indicators such as **Gauges**, **Dials**, **Bullets**, **Plus/Minus** indicators and **Sparklines** can be used to illustrate:

- actual performance versus goal
- current status compared to performance thresholds
- deviation from expected trend
- direction and magnitude from a benchmark value

Indicators also add context to tabular data views, providing a top-level summary for users who don't require or have the time to examine additional detail data. Indicators can be designed to display both qualitative metrics, typically via the use of color or color intensity, as well as quantitative metrics which are encoded as numeric labels.

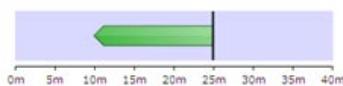
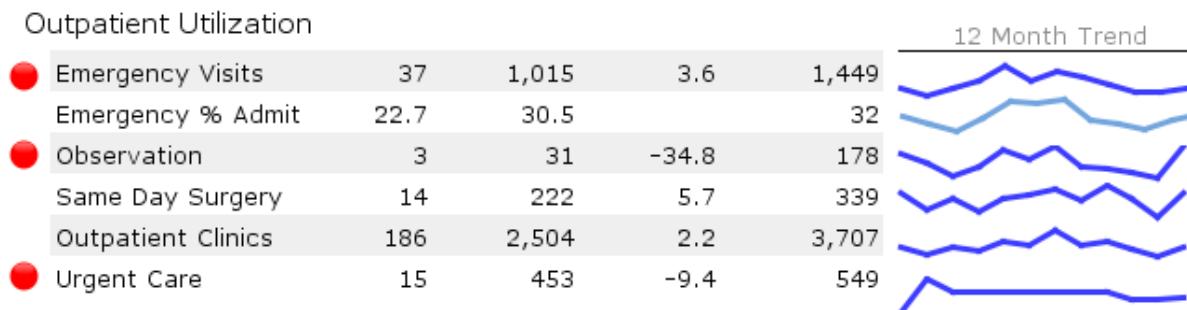


Alerts are a single shape whose color and/or direction changes based on whether the value is above or below a specified threshold.



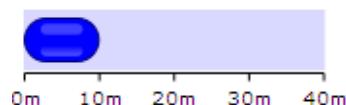
Sparklines are compact, sparsely formatted line charts. Their purpose is to show the overall trend of the displayed data and to show the most recent (rightmost) data point in the context of its history. When stacked in a dashboard, **Sparklines** provide immediate and unambiguous trend comparisons among related data entities.

The following example combines **Alerts** with **Sparklines** to enhance the information content of the Outpatient Utilization table. The red **Alerts** are positioned immediately to the left of the text labels so that the entities requiring attention are clear to the user. Quantitative information can be gained by examining the numerical values to the right of the text labels, while the 12 month trend on the far right provides historical context for why the three flagged areas may be exhibiting problematic behavior.

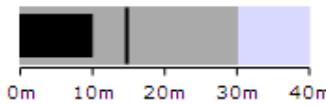


Plus/Minus indicators display both magnitude and direction from a predefined baseline for a given metric. Color can be used to provide context. In this example, the arrow is displayed in green, rather than red, to signify cost reductions from \$25M to \$10M. If we were displaying a drop in profits instead, red would obviously be the appropriate color.

The **Bullet** indicator is a compact and highly customizable visual dashboard building block.

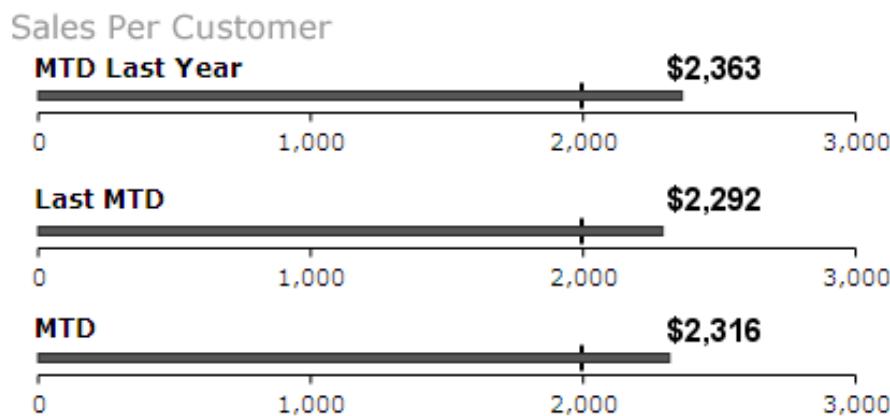


In its simplest form, the **Bullet** indicator displays a quantitative value and a horizontal scale. No additional context or threshold is displayed.



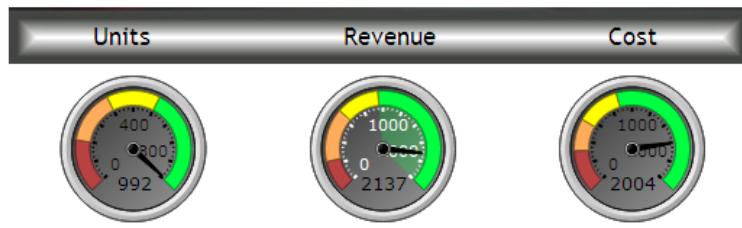
This version of the **Bullet** indicator adds additional context by displaying a threshold via the black vertical line, and also encodes a qualitative measure using color (satisfactory, unsatisfactory).

The compactness of **Bullet** indicators makes them an excellent choice for facilitating “at-a-glance” comparisons between related metrics. Note that the exact dollar value can be displayed as a text label when this information is required by the user. If this is not useful, then don’t clutter your dashboard with the extra labels.



Combining a basic **Bullet** layout with **Alerts** helps illustrate month-to-date auto sales and directs our attention to the month-over-month changes that are either up or down significantly. We could also have displayed this data equally effectively by incorporating the numeric columns directly into the **Bullet** indicator as in the “Sales Per Customer” example above, thereby eliminating those columns from the dashboard:

Make	Sales MTD	Last MTD	Diff	Change%
Acura	360,943	383,882	-22,938	-5.98 ▼
Audi	527,153	494,443	32,709	6.62 ▲
BMW	716,778	646,213	70,565	10.92 ▲
Buick	481,480	470,405	11,075	2.35
Cadillac	587,325	679,380	-92,055	-13.55 ▼
Chevrolet	420,864	396,578	24,286	6.12 ▲



Gauges and **speedometers** use a “needle” to indicate the current value of a metric. Red and green indicate problematic/preferred zones of performance. As with stacked **Bullet** indicators, **Gauges** and **Speedometers** make visual comparisons between related metrics easy. While eye-catching, these indicators do not make optimal use of dashboard space, so use them sparingly.

sparingly. This is probably one of the most overused dashboard indicators. Consider using one of the **Bullet** indicators whenever possible, especially when dashboard space is at a premium.



Racetrack lights are a line of 2 to 5 color-coded circles. Each circle denotes a discrete value based on its relative position in the line. Depending on context, this indicator can convey either a quantitative (customer satisfaction rating from 1 to 5) or qualitative (poor, okay, neutral, satisfactory, excellent) value.



A **Slider** indicator displays the selected value in a box that is positioned along the continuum from right to left depending on the value. The color gradient of the slider bar also varies to maximize visual impact.

The following intentionally redundant example shows some of the ways that these indicators can be used to encode a particular metric and the tradeoffs between various indicators:

- Color coded text can be as effective as graphical indicators in conveying qualitative information or calling attention to problematic data points
- A Slider is more effective than a simple tabular display of those values because it positions those values within a range
- While the simple Bullet indicator seems less cluttered than the Slider, if end users require the actual numeric value then the simple Bullet will fall short
- If the relative position or category of a data entity matters, then the Racetrack indicator will need to be used. If not, use a simple Alert instead and save dashboard space.



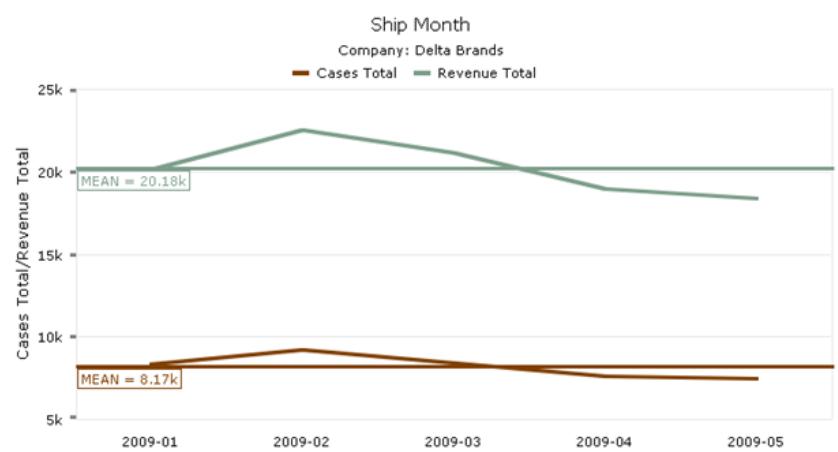
Next week, we'll explore basic charts to see how they can be used most effectively in a dashboard environment.

Part 2 - Basic Charts

This category includes **Line**, **Bar**, **Point**, **Linear Bubble**, and **Pie** charts. These charts are simple to construct and very easy for users to comprehend, because they are used to display exactly one data element alongside one or more quantitative metrics describing that element. The **Bar** chart below displays the Product Brand for a supply chain company and two metrics describing that dimension, Total Cases Sold and Total Revenue. Depending on how much dashboard space is available, several additional metrics could have been displayed. Although this bar chart is displayed in a 3D format, it adds no additional information content but may slightly increase the visual impact of your dashboard. Don't use fancy display options just because your BI tool provides them, ask yourself what, if any, additional benefit these options provide.



While **Line** and **Point** charts might appear to be interchangeable with **Bar** charts, these two charts are most effective when used to display time or date based dimensions. As with **Bar** charts, multiple quantitative metrics can be displayed in a single chart, provided that the chart does not get too crowded. In this example, the mean statistic is superimposed over each metric. This can be useful for alerting dashboard users when a particular metric has deviated too far from expected behavior, which can be either good or bad, depending on the metric. If statistics such as mean, median or regression lines are requested by end users or help to convey important information that is not immediately obvious by looking at the chart, then display it. If not, then don't clutter your charts with it. Most importantly, make sure the statistic is valid - displaying a regression line through only 3 points is not only statistically invalid, it implies a trending relationship that probably does not exist. Don't mislead your users by displaying erroneous information.



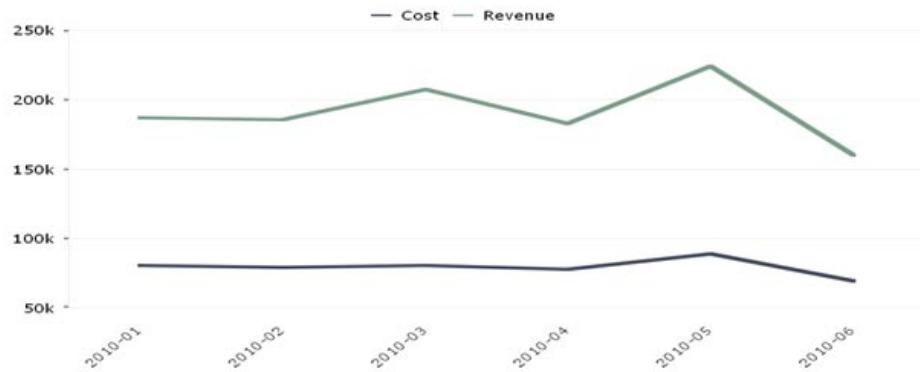
A **Point** chart is simply a **Line** chart without the connecting lines. **Point** charts are useful for displaying a single quantitative metric, but become confusing when multiple metrics are displayed. The **Point** chart below would be easier to understand as a **Line** chart, because the Cost and Revenue metrics would be clearly differentiated:



The **Spline** option is probably one of the most misused formatting options found in today's dashboard environments. Intended to make boring line charts more visually appealing by turning lines into curves, the spline option should rarely if ever be used, since it implies that the underlying data can be modeled by a polynomial relationship. This error is similar to the regression line mentioned above, and should be used only when a polynomial relationship truly exists.

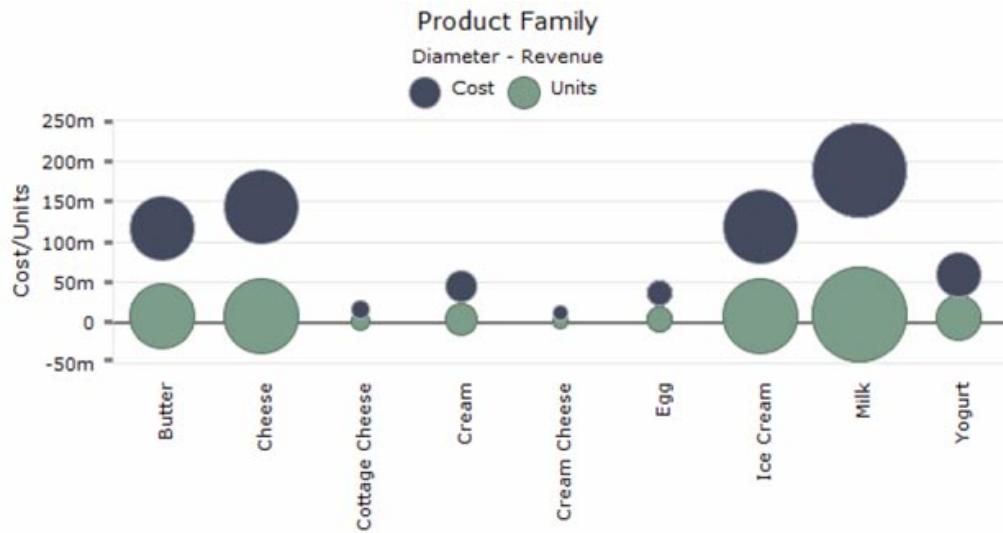


With their limited display real estate, dashboards present developers with the dual challenges of maximizing the information content displayed while avoiding clutter and end-user confusion. One way to meet this challenge is to display multiple data series in a single graph. But pay close attention to scaling. The Revenue metric has relatively little variance compared to the Cost metric, so using a single vertical scale makes the Revenue metric appear as a nearly flat line.



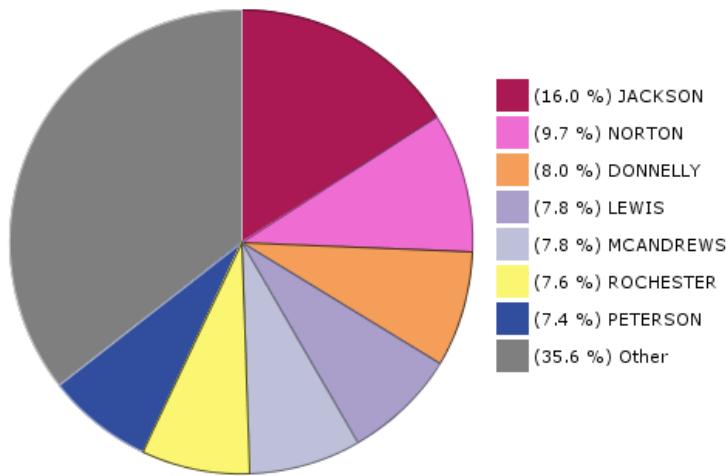
By scaling Cost separately on the right side of the chart and using a smaller chart increment the variability in this series is brought to the user's attention.

When dashboard space is at a premium, an additional quantitative metric can be displayed on a chart without using a separate data series. In this **Linear Bubble** chart, the Revenue metric is displayed by varying the diameter of the bubbles, rather than cluttering the chart with a third set of Revenue bubbles. Be sure that your chart legend makes users aware that the bubble diameter is being used for this purpose.



Pie Charts are a compact tool for displaying data elements that have relatively few components and only one quantitative metric. Pie charts require that the component values sum to 100% and that no component is negative. By using a “catch-all” category labeled Other, the smallest (and likely least significant) data points are consolidated. This avoids graph clutter and draws the user's attention to the

important data points on the pie chart. Visualization gurus suggest that the maximum number of data points to display on a pie chart is between 7 and 10. While this seems intuitive in most cases, if the distribution of data is severely skewed, such as 80%, 10%, 5%, 1%, 1%, 1%, 1% it makes more sense to consolidate the 1% data points into an “Other” category.

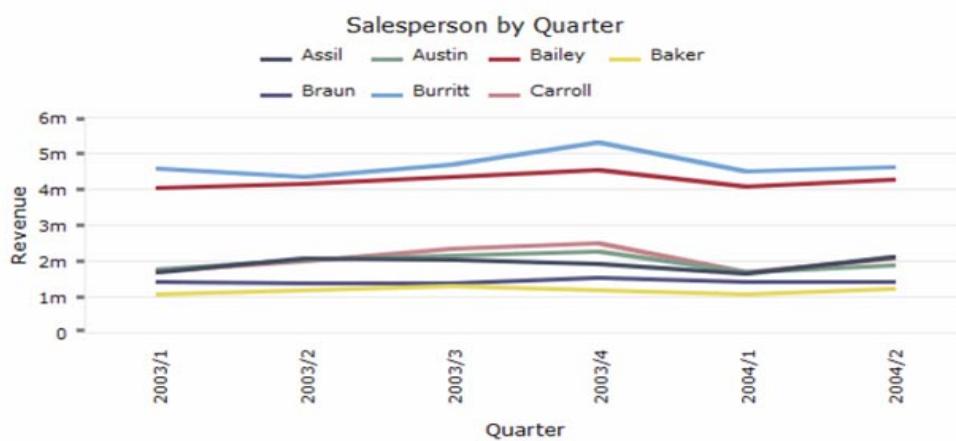
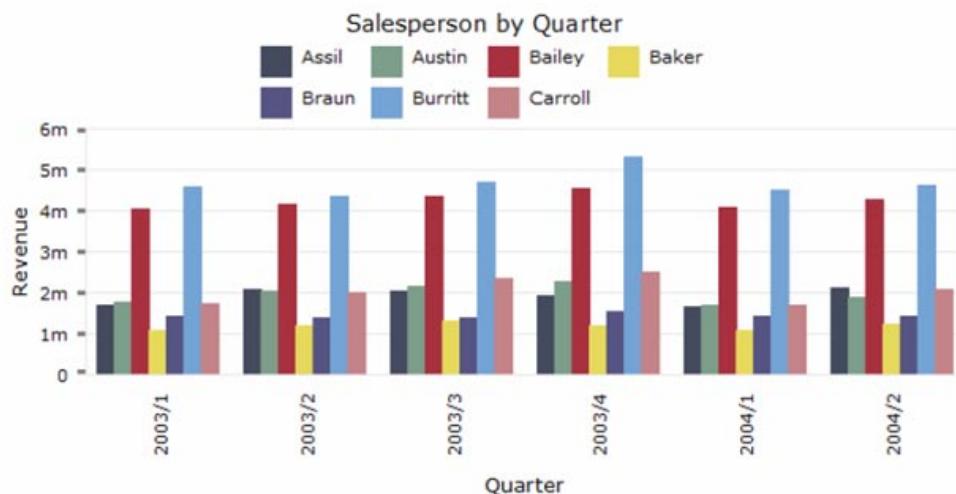


This illustrates another potential problem to plan for when selecting dashboard chart types: high variability in the underlying data series. What happens when your prototype data resembles the 8 relatively similar categories shown in the chart above, but varies from update to update, at times resembling a severely skewed distribution? Dashboard designers obviously can't anticipate every nuance of their future data feeds, but enough existing data should be run through your dashboard during the design phase to see what your charts look like at the extremes.

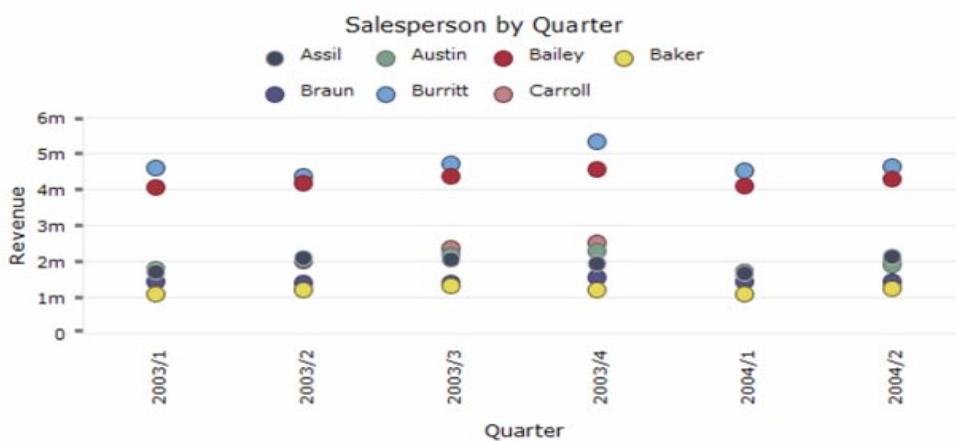
In next week's Dashboard Insight article, we'll discuss which charts to use for effectively displaying two data elements and their associated metrics.

Part 3 - Displaying Two Data Elements in a Single Chart

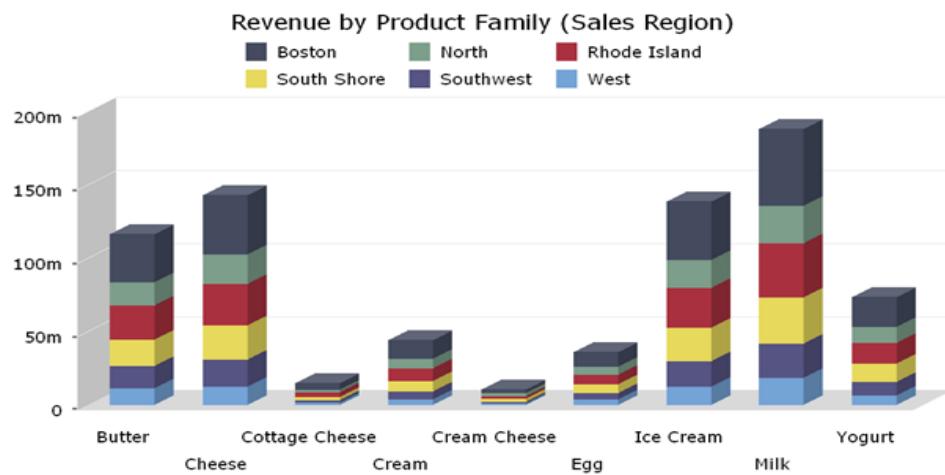
Basic charts such as **Bar**, **Line**, **Point**, and **Bubble** charts can be adapted to display the intersection of two data elements. The following three charts display revenue for the dual elements of fiscal Quarter and Salesperson. If your end user is more interested in intra-quarter comparisons between sales people, the **Bar** chart is the clearest representation. On the other hand, the **Line** chart does a better job of illustrating quarter to quarter revenue variations for each Salesperson. The **Point** chart is the least desirable option, because several data points are hidden underneath others. The **Point** chart does however provide the clearest display of Salesperson ranking within each Quarter.



A **Bar Stack** chart is another way to encode two data elements in a single graph. In this example,

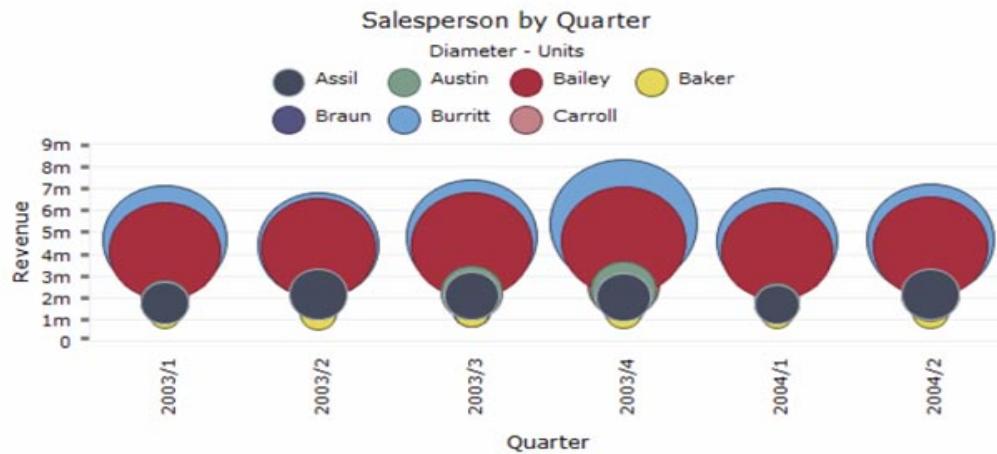


Product Family is represented by the labels on the x-axis while the colored columns represent the Sales Region. Staggering the x-axis labels instead of displaying them at an angle makes the graph far more readable. At a glance, dashboard users can see that the biggest revenue generators are Butter, Cheese, Ice Cream and Milk, while Cottage Cheese and Cream Cheese are underperformers.



But what if we are more interested in gaining a clearer picture of the revenue each Sales Region contributes to the individual Product Families? A variation called a **Percentage Bar Stack** chart can be used. The y-axis now describes percentage of revenue rather than absolute dollars amounts. Knowing which view of the data your users are interested in helps to choose the right chart style.

Just like **Linear Bubble** charts, the **Bubble Cross-plot** uses the bubble diameter to display an additional quantitative metric, in this case Units Sold.



In next week's edition of Dashboard Insight, we'll present some additional charts that you can incorporate into your dashboards to maximize the information content of certain types of data.

Part 4 - Special Purpose Charts

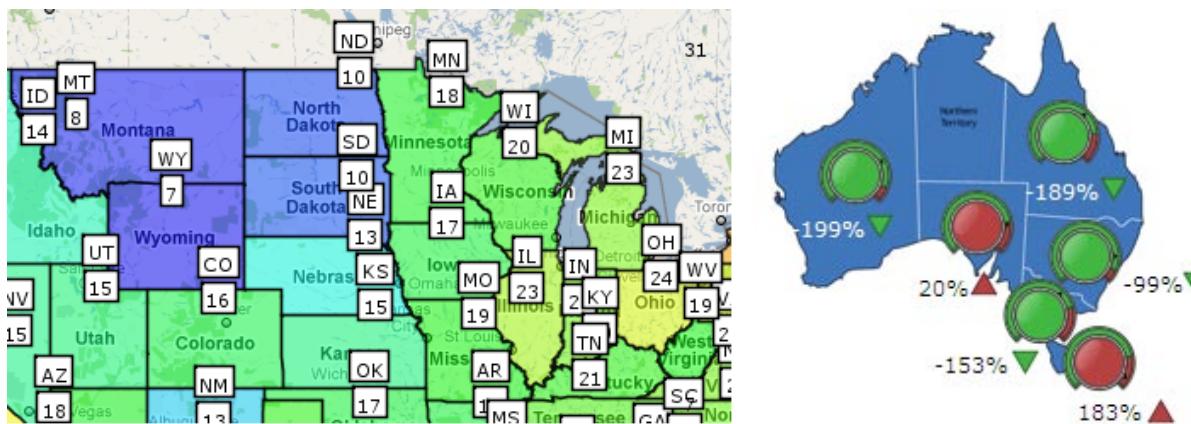
Now let's look at a few special purpose charts found in today's dashboard environments. These charts are very effective for displaying certain types of data and can set your dashboards apart from the crowd when applied correctly.

A **Calendar Chart** is a compact and intuitively simple vehicle for displaying time based data. This provides visibility into cyclic events, and facilitates easy periodic comparisons, such as month to month or week to week. With calendar charts, users can quickly spot trends and answer questions such as:

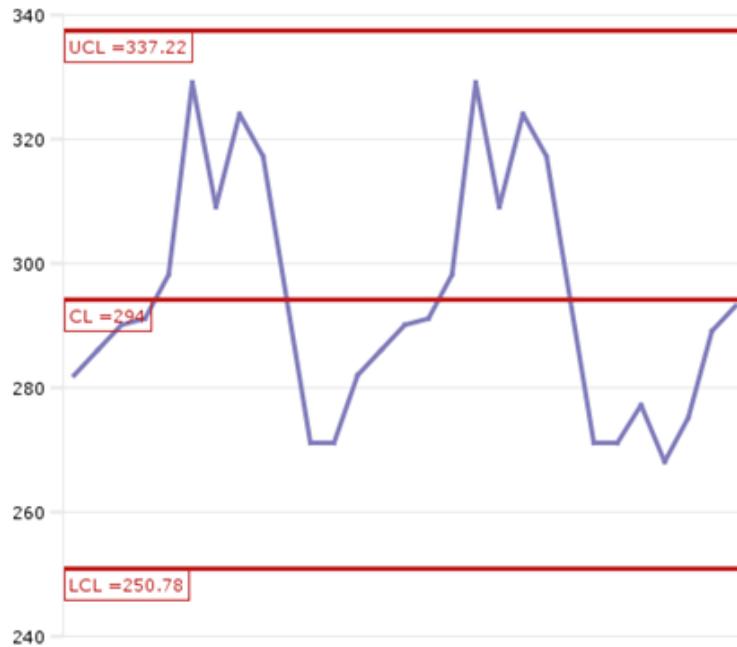
- Is revenue higher at the start or end of a month?
- Do costs tend to rise or fall on Fridays?
- Are these metrics impacted by seasonal fluctuations?



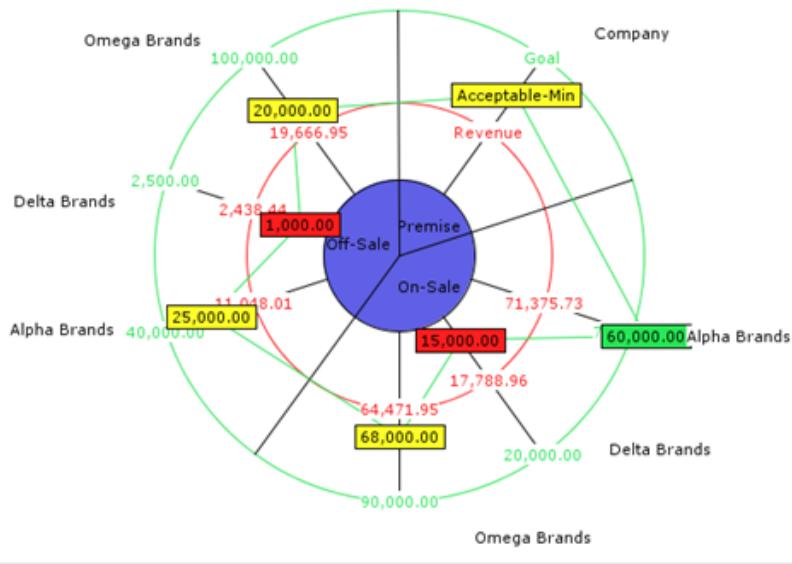
Maps provide a great way to display geographic data dimensions. A combination of text labels, dashboard indicators, and color coding are used to display quantitative metrics such as per capita income ranking and median level of education for the State dimension:



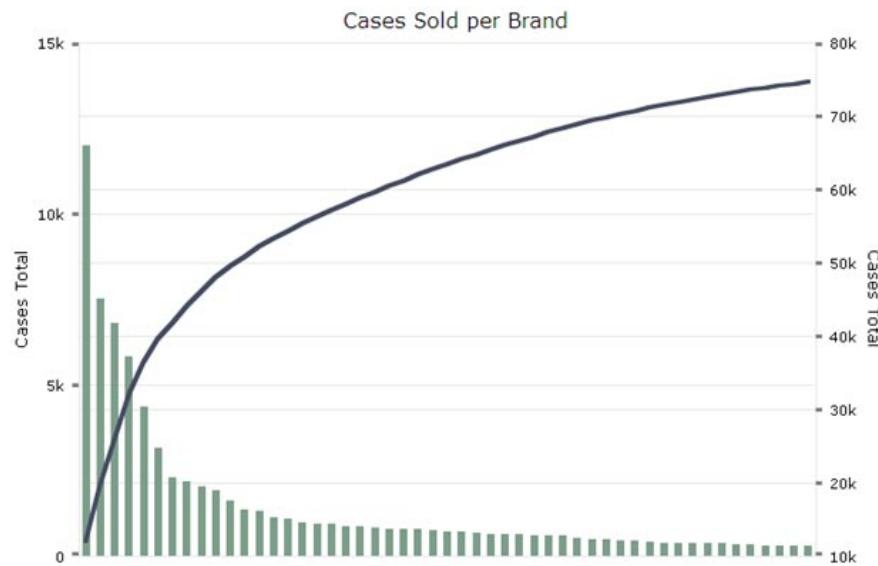
Statistical Process Control charts are used extensively in manufacturing and healthcare to track key metrics such as defects on an assembly line or hospital-acquired infections. Although several variations of control charts exist, most control charts display upper and lower control limits as well as the current time series of data.



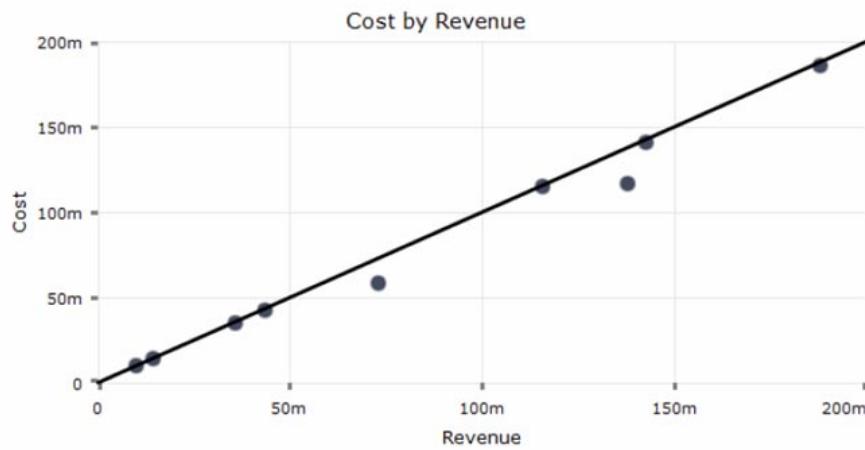
In **Radar**, or **Spider** charts, data is plotted radially from a center point. Each data point has a vertical component, namely, its distance from the center of the chart. Radar graphs are useful for displaying cyclical data such as daily staffing levels or monthly revenue totals. Dashboard designers can overlay goals and thresholds to provide context for the quantitative metric that is displayed in the Radar Chart. Depending on the dashboard environment, Area, Line, and Point formats are typically supported for Radar charts.



Pareto charts are constructed by creating a sorted bar chart and overlaying a cumulative total line. This chart provides a visual interpretation of the contribution made by the “top N” values of the charted dimension. In the following graph of beverage brands, the pareto chart illustrates that the top 6 brands were responsible for 40,000 cases sold, or more than half of total sales.



In some cases, the absolute value of a data dimension’s quantitative metric is less important than its value relative to other numbers. For example, height vs. weight, cost vs. revenue, or profit vs. sales. For these situations a **Scatter Plot** is ideal. The diagonal line provides context for the displayed data points. In this case, more points below the diagonal line is preferable to seeing more points above the diagonal line:



Summary

Hopefully this series of Dashboard Insight articles has provided some chart and indicator selection guidelines to consider before launching your next dashboard project. But rather than rigidly adhering to these guidelines, remember to experiment with scaling, orientation, and other charting options by iteratively developing dashboard prototypes. Most importantly, don't work in a vacuum - involve your end users early and often in the design process to make sure your dashboard uses graphs in the most visually impactful manner.

Sponsored by



About Dimensional Insight

Dimensional Insight, www.dimins.com, is a leading privately held company specializing in the development of business intelligence solutions. Founded in 1989, Dimensional Insight has more than 2,600 customers and over 50 partners throughout the Americas, Europe, Asia, Africa and the Middle East.

Dimensional Insight offers a comprehensive set of BI capabilities ranging from data integration to sophisticated reporting, analytics, information portals and dashboards. Follow DI on Twitter http://www.twitter.com/DI_tweet.

Frank Smietana
Corporate Business Intelligence Strategist



About Dashboard Insight

Dashboard Insight (DI) is the paramount online resource for professionals in the business intelligence (BI) and data visualization industries. The DI community bridges the gap between business people (analysts, executives, other managers, dashboard users) and development/IT staff, making it the definitive forum for advancing the digital dashboard industry.

Steve Bogdon
Dashboard Insight
www.dashboardinsight.com