

# Coordinated Highlighting in Context

## Bringing Multidimensional Connections to Light

Stephen Few, Perceptual Edge  
*Visual Business Intelligence Newsletter*  
 July, August, September 2010

A promising visual analysis technique was first proposed back in the late 1970s, which has since been researched fairly well, but has seldom been integrated into commercial software and never to its full potential. I'm referring to the technique called *brushing and linking*. Chances are, you've never heard of it—at least not by this name. If this is the case, I'd like to introduce you to a powerful way to spot and examine otherwise elusive meanings in quantitative data.

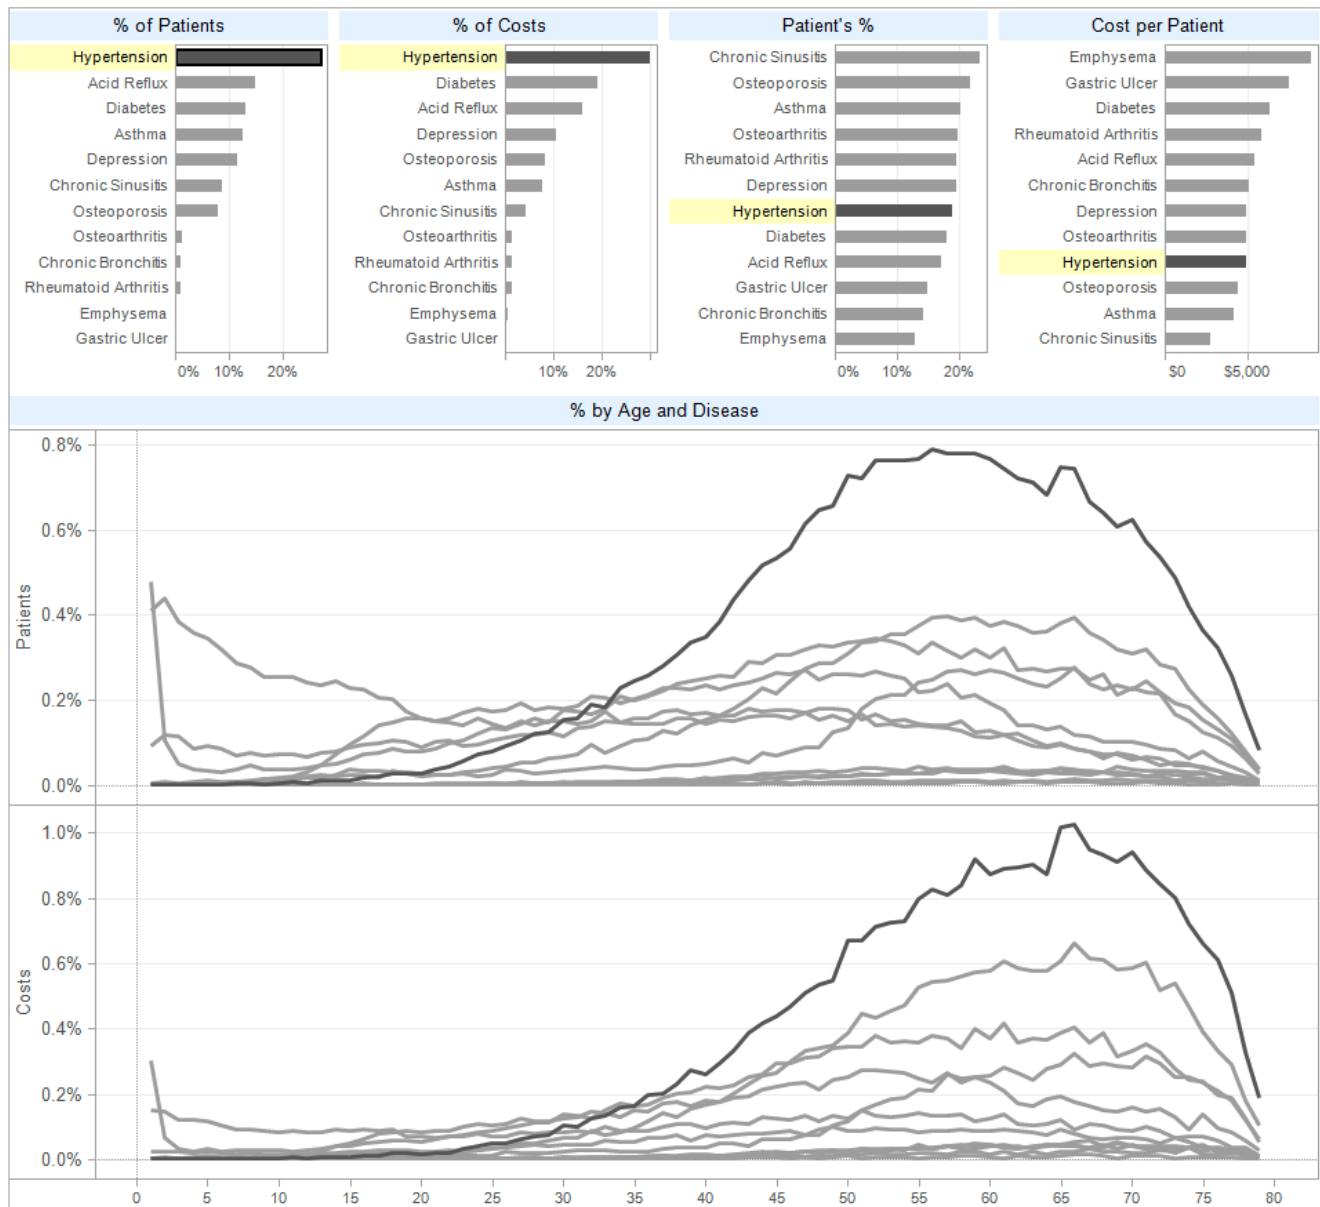
### Brushing and Linking Defined

*Brushing* is the simple act of directly selecting data in a chart. Imagine that your mouse is a paintbrush that you can use to select specific objects in a chart such as data points, bars, and lines by brushing across them or drawing an enclosure around them.



When we use charts to analyze data, there is usually no more direct or simpler way to select particular items that interest us than by grabbing them right there in the chart. This certainly beats having to construct a SQL query with Boolean operators. Direct data selection in a chart can be done without taking our eyes off the information, which allows us to keep our minds in the flow of analysis without interruption.

*Linking* takes advantage of the fact that the items we've selected in one chart are related to items in other charts by automatically selecting them wherever they appear. The link causes selected items to automatically and immediately exhibit the same response—usually highlighting—in all charts. In the following example of a coordinated multi-chart display of healthcare costs, I have brushed the “Hypertension” bar in the upper left to select it, which caused data related to hypertension to be automatically highlighted in all of the charts.




The process of data analysis requires a great deal of interaction with information to spot what's interesting and then tease out what it means. The fundamental activity of data analysis is comparison. Almost everything we do while analyzing data either involves comparison directly or is done to make comparisons possible. We might use a bar graph to compare a series of values to one another. We might sort the values from high to low to make that comparison easier. We might display the meandering nature of monthly sales in a line graph to compare trends and patterns of change through time among several sales regions. We might use a scatterplot to compare the costs of marketing products to their respective profits, which leads us to discover a lone product


that costs little to market but earns huge profits nonetheless. If that one outlier complicates our ability to examine the otherwise predominant correlation between costs and profits, we might filter that data point from the scatterplot to remove its skewing effect on the other data. All of these interactions are either comparisons or are done in support of comparisons. To these examples we can add brushing and linking as a unique and powerful tool that will expand our ability to make comparisons into otherwise undiscovered territory.

It's almost impossible to compare things unless they are simultaneously present in front of our eyes. One of the reasons that data visualization is so useful for analysis is its ability to put a great deal of data in front of our eyes at once in a way that we can understand and compare—far more than we could manage if we were relying on tables of numbers alone.

A single line in the following graph reveals the changing number of houses that sold each month throughout the year—the equivalent of twelve separate numbers in a table—in a way that can be perceived and understood at a glance as a single visual pattern.



This natural benefit of visualization can be magnified by placing several pictures of data on the screen at once. When examining monthly house sales, rather than looking at each zip code and each group of homes based on square footage one at a time, the following series of graphs, arranged as a visual crosstab, makes it possible to not only see each combination of zip code and house size, but to easily and rapidly compare them because they're all right there in front of our eyes.

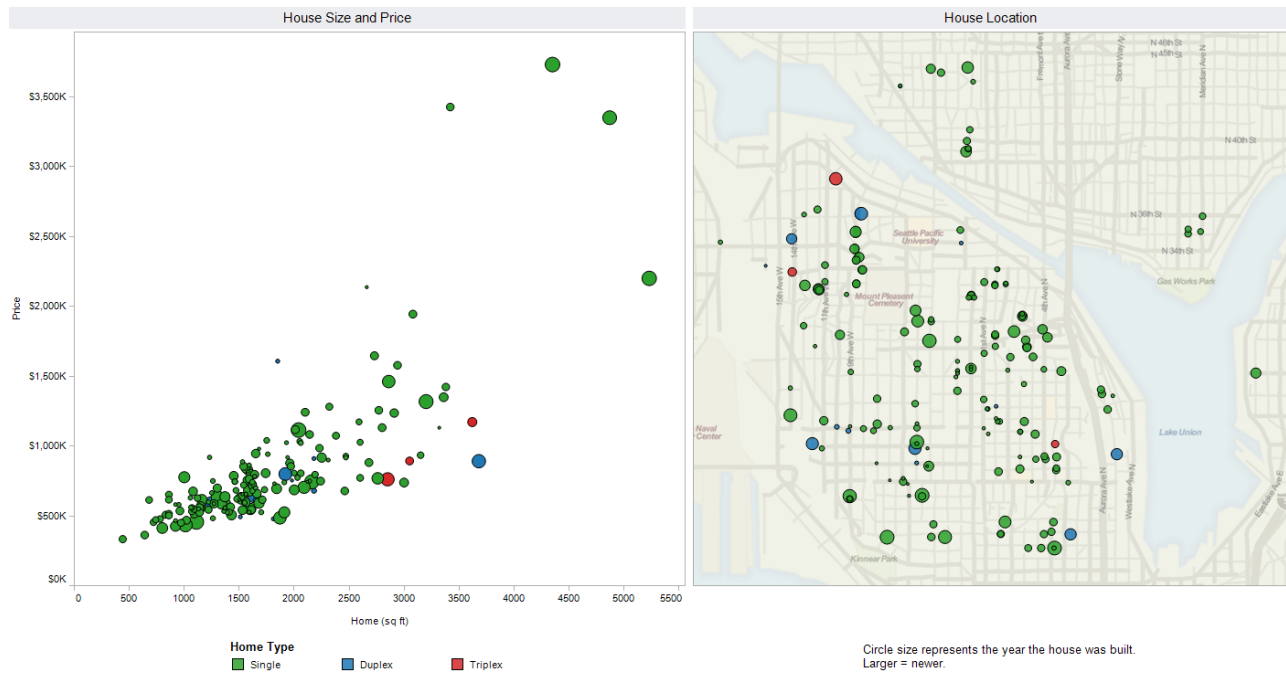


What we have here is a single view of information that has been broken into a series of related graphs in a way that extends the amount of data we can simultaneously examine and compare without becoming overwhelmed, because the information has been broken into perceptually manageable chunks. This is an example of what Edward Tufte called *small multiples* in his book *The Visual Display of Quantitative Information* way back in 1983.

Quite often during the course of analysis, we must examine information from multiple perspectives—not just a single view broken into chunks like the one above, but entirely different views, because the information must be examined multi-dimensionally to be understood.

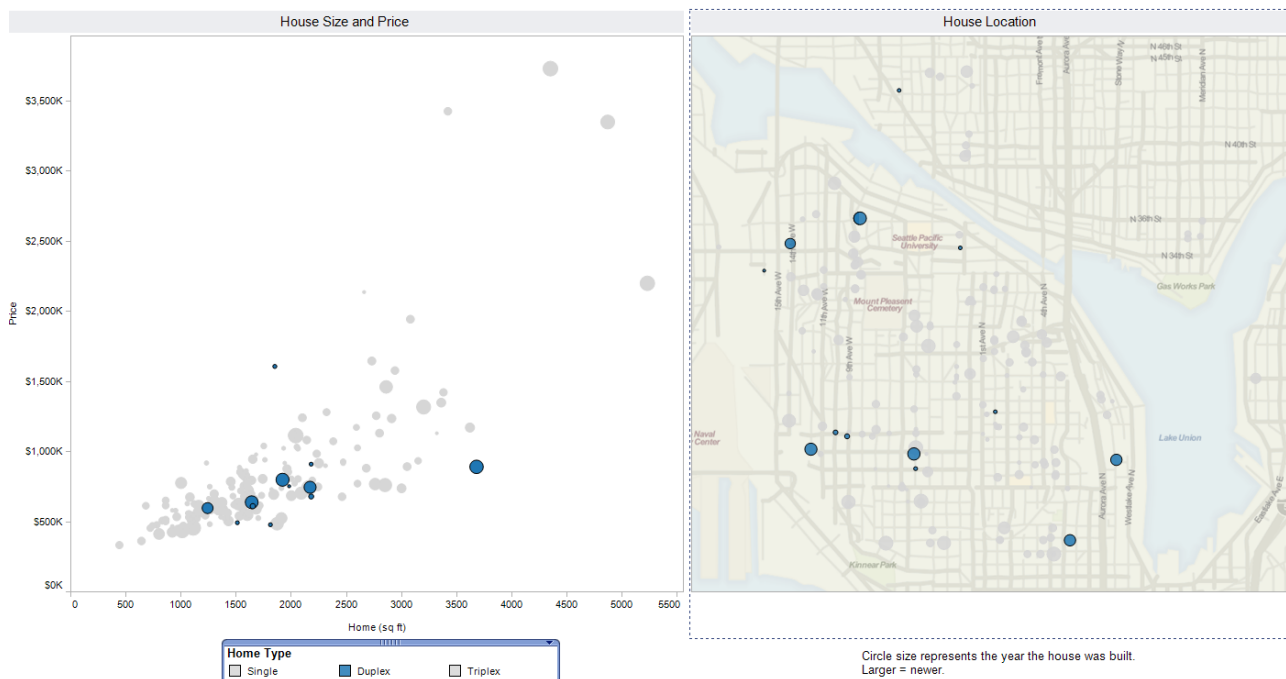


We could look at these different views and their distinct perspectives one at a time, but there are often discoveries that can only be made when those views are laid out in front of our eyes at the same time. In the following example, we see the same set of house sales displayed from two distinct but complementary perspectives—a scatterplot that features house sizes and prices and a map that features their locations. In both views, the sizes of the circles represent when the houses were built—small circles are old houses and the largest are the newest. The colors of the circles represent house types—green for single family, blue for duplex, and red for triplex. Even though these two charts display four variables each—size, price, age, and type in the scatterplot and longitude, latitude, age, and type on the map—by showing them both on the screen at once, we have created an even richer multidimensional workspace for analysis.



It wouldn't be possible to display the correlation between house sizes and prices as well as their geographical locations in a single graphical view—at least not in a manner that could actually be perceived, understood, and used for analysis. By placing a scatterplot and a map side by side, however, we can display multiple perspectives simultaneously without overcomplicating the information.


When multiple views are concurrently displayed and linked in a way that causes interaction with one to be automatically reflected in the others, information visualization researchers refer those views as *tightly coupled* or *coordinated*. For instance, because the two views in the example above are tied together in this manner, we could use a data filter to remove the same set of houses such as all duplexes (the blue circles) from each view through a single manipulation of the filter control. We could also highlight the duplexes (vs. single-family homes and triplexes) in both displays to make it easier to focus on them without losing awareness of the overall context (that is, the entire set of houses), as shown below.



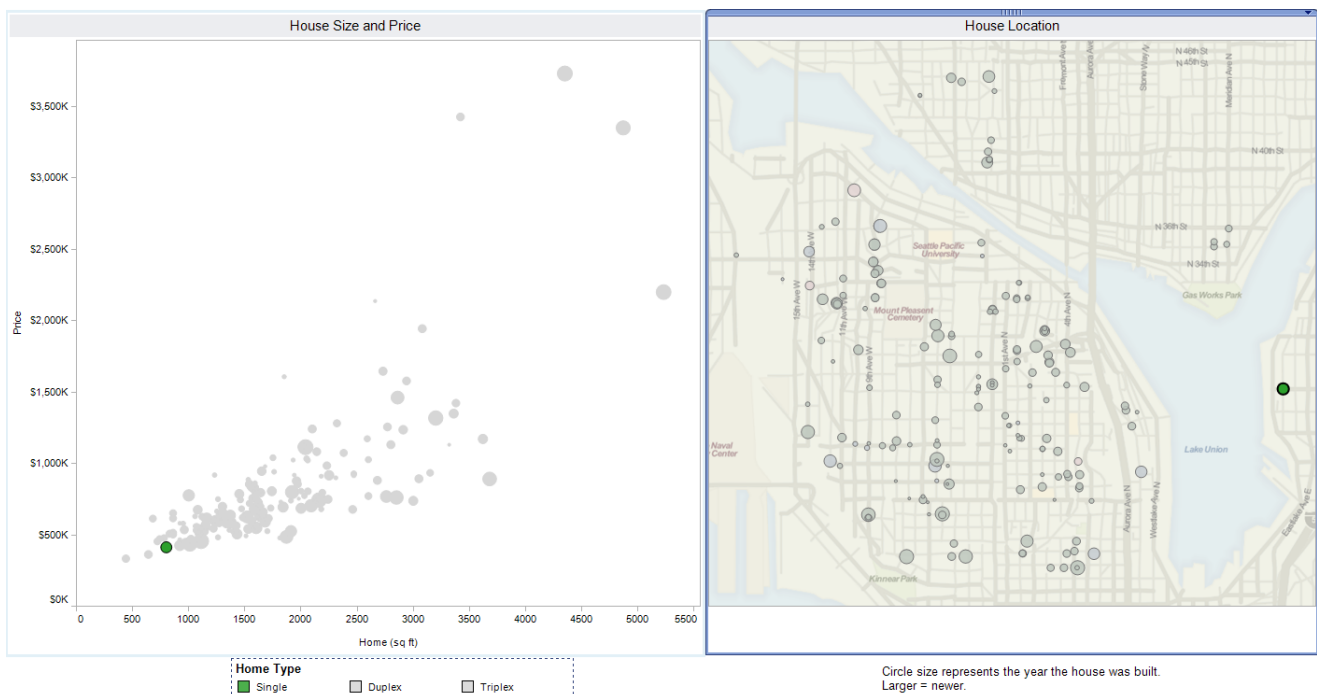


As you might imagine, the ability to examine and compare multiple views of information on the screen at same time is one thing, which in and of itself is quite useful, but the ability to interact with those views in a coordinated manner through a single interaction raises the possibility of discovery to a whole new level. Brushing and linking is one such way to interact with multiple views in a coordinated manner—one that can enable unique and powerful insights.

Let's see how brushing and linking can bring these house sales to life by flushing out facts that might otherwise hide in the shadows. On the map below, notice the lone house in the east (the right) of Lake Union (marked with a red arrow).

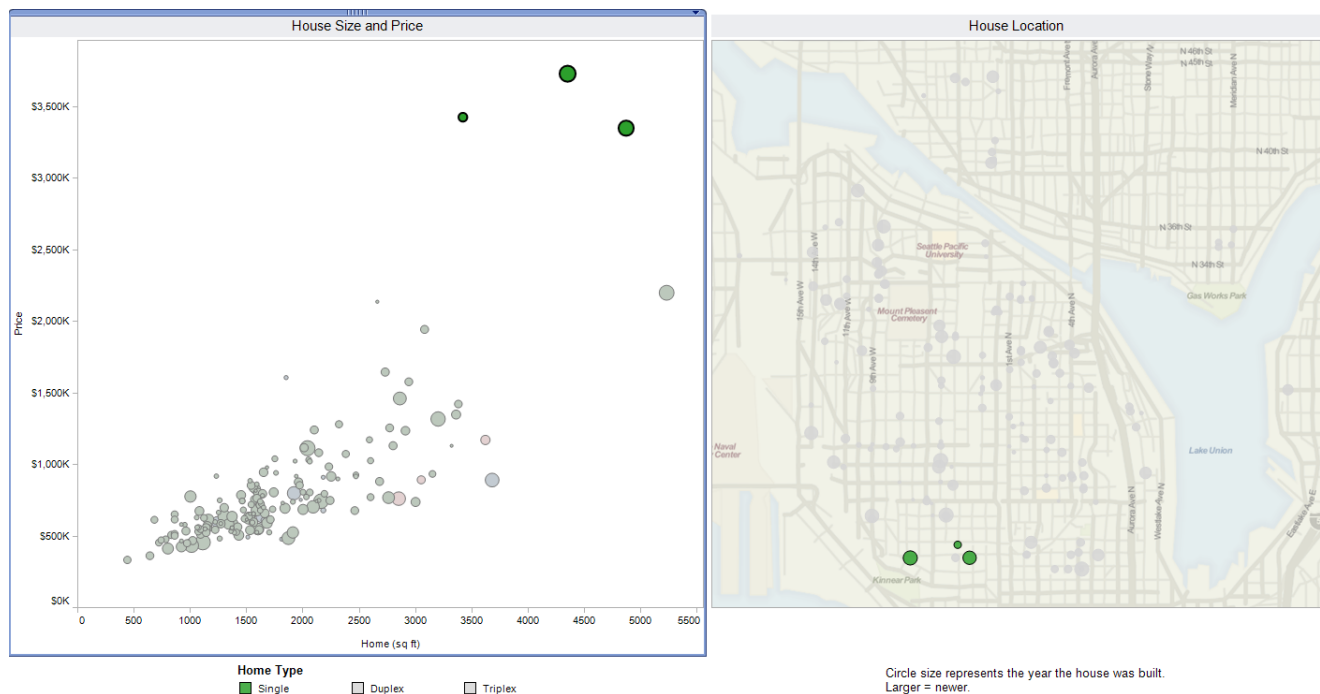


We know its location and the fact that it's a single family house by its green color, but what about its size and price? To see these other attributes, we need to find that house in the scatterplot, but we certainly don't want to hover over the data points one at a time to access the identity of each house individually until we finally stumble onto it. We could easily find it by brushing it on the map to highlight it. When I did this, the appearance of the display automatically changed to what's shown below.



Not only has the selected house been highlighted on the map by having its border thickened and by dimming the appearance of all other houses to reduce distraction, but the same highlighting effect has been automatically applied to that house in the scatterplot. This simple act has made a meaningful connection visible and understandable. If this example, consisting of only two views, doesn't excite you, hold on, because you'll soon see that the potential of brushing and linking really shines in examples that are more complex.

For the moment, however, let's keep things simple and try something new. Notice the three high-priced houses in the top right section of the scatterplot. Let's find out where they are on the map by using the mouse to select and highlight them in the scatterplot. The result of this action appears below. As you can see on the map, all three of these houses are in the same neighborhood.



Before going on to even more interesting examples, let's think a bit more about multi-view displays and the occasions when they're useful.

## Benefits of Multi-View Analytical Displays

We shouldn't use multi-view displays gratuitously, simply because we can. We should use them when they give us the ability to make sense of data in ways that we couldn't in any other way or we couldn't otherwise do as easily and quickly. Young, Valero-Mara, and Friendly have clearly explained why we must view data in multiple ways.

*It is usually the case that a single display cannot show you everything that is interesting about a set of data. Even for the simplest data, we will often see different features of interest in different displays, if only because most graphics are designed to help us explore one (or at most a few) features. Thus, it is usually necessary to examine several different graphics in order to get the full picture of the problem at hand.*

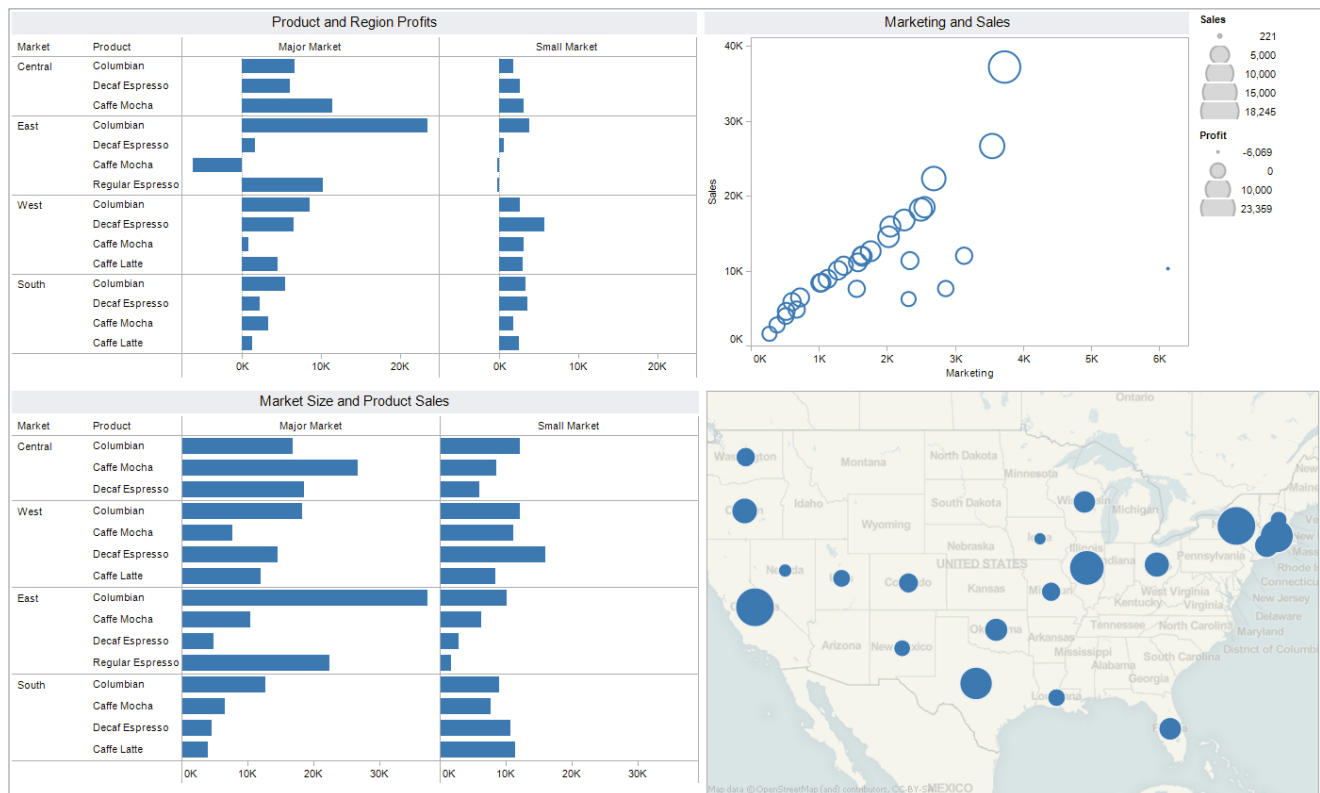
*(Visual Statistics: Seeing Data with Dynamic Interactive Graphics, Forrest W. Young, Pedro M. Valero-Mara, and Michael Friendly, John Wiley and Sons, Inc., 2006, p. 133)*

But why not use one view at a time in sequence, rather than placing several on the screen at once? When we want to compare information that's contained in multiple views, we must place them on the screen at the same time because of a limitation that's built into our brains. Working memory, which is where we temporarily store chunks of information during the active process of thinking, is surprisingly limited. According to the latest research, we can only store about three chunks of visual information in working memory at once. One

of the natural benefits of displaying quantitative data visually is the fact that several numbers, which might otherwise appear in a table, can often be perceived as a single visual pattern, such as a line, which could be stored as a single chunk in memory. When properly displaying information graphically, we can squeeze more information into those three available slots in memory, but we can't expand the number of slots. Good visual analysis software augments working memory by allowing us to place several coordinated views on the screen at once, which in effect uses the screen for storage external to our brains. Anything that's in front of our eyes can be rapidly swapped in and out of working memory as needed. If we were forced to examine those same views sequentially, one at a time, however, we could only compare something that appears in the current view to three chunks of information at most that appeared in previous views. Analyzing data in this manner takes forever and produces few discoveries.

Multi-view displays can be used to solve various problems. The following five occasions are particularly useful and common:

1. The number of variables that we need to compare exceeds what can be effectively displayed in a single view.
2. Our ability to examine the data would be undermined by visual clutter if displayed in a single graph.
3. We would be perceptually and cognitively overwhelmed using a single graph to perform the task at hand.
4. We must compare various levels of summary and detail (for example, the same year's worth of data expressed quarterly, monthly, and weekly).
5. We must examine the data from multiple perspectives that require different types of graphs (line graphs, bar graphs, scatterplots, maps, etc.).

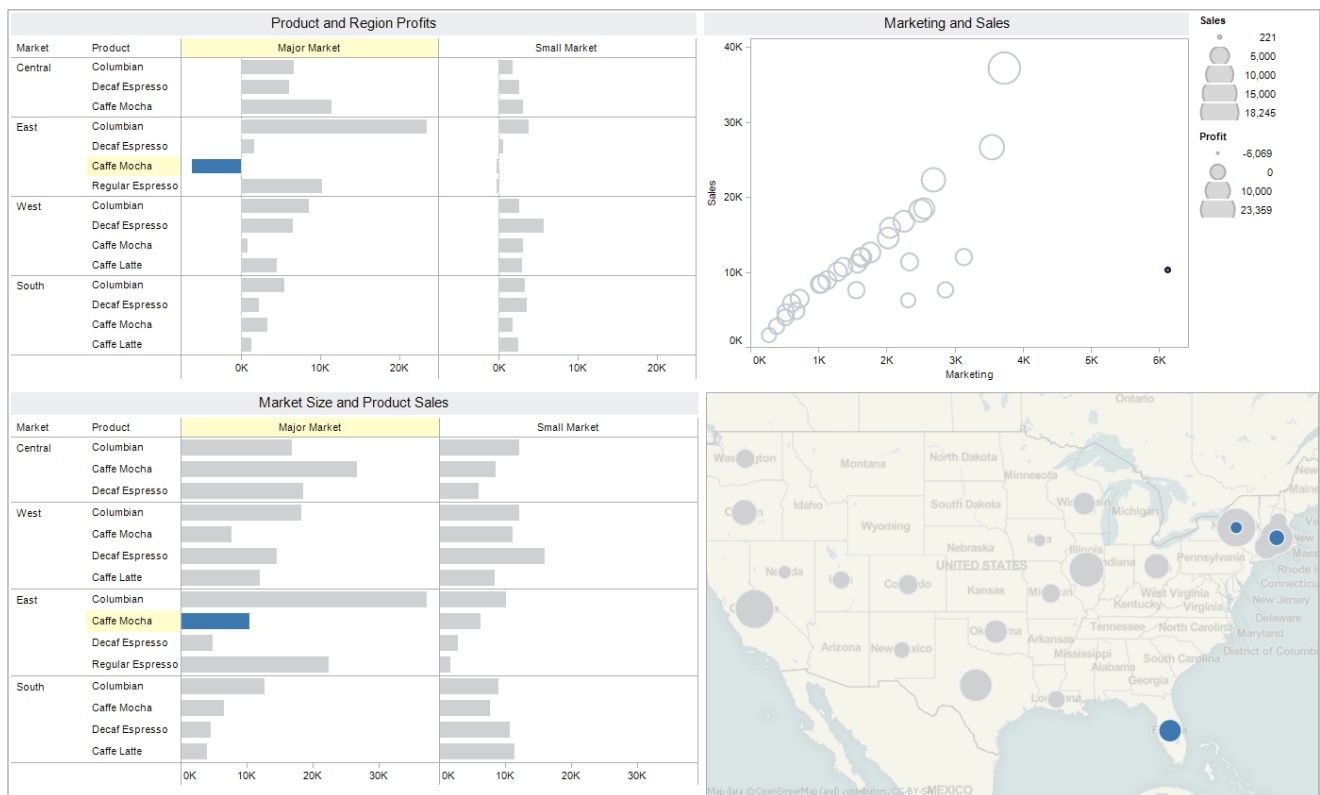


This example illustrates all of these occasions.



1. There is no way that we could examine all of these variables in a single graph.
2. Even if a single graph could handle all the variables in a meaningful way, it would appear much too cluttered.
3. By breaking this story down into smaller sub-stories, the information is chunked into perceptually-sized bites.
4. On the map, I'm concerned with a state-level summary, but in the other views I'm concerned with individual products by region and market type (major and small markets).
5. No one type of graph could support the many perspectives that I need to see.

When we apply brushing and linking to this display as well, what's already enlightening is extended dramatically. For instance, you have perhaps noticed the tiny outlier on the right side of the scatterplot. What is that and how does it relate to the entire story that's told in this multi-view display? The answer becomes readily apparent when we brush the outlier to see where it lives in each of the charts.



The story that is told by this highlighted display goes something like this:

1. This outlier with excessively high marketing expenses and relatively low revenues represents sales of Caffe Mocha in the East.
2. Sales of Caffe Mocha in the East are losing money—a little among Small Market customers but a lot among Major Market customers (see the upper left-hand graph).
3. Although this product represents only a portion of sales in the Massachusetts and New York, it constitutes all sales in Florida.

The simple act of brushing one little data point in the scatterplot brought this story to light and expressed it simply.

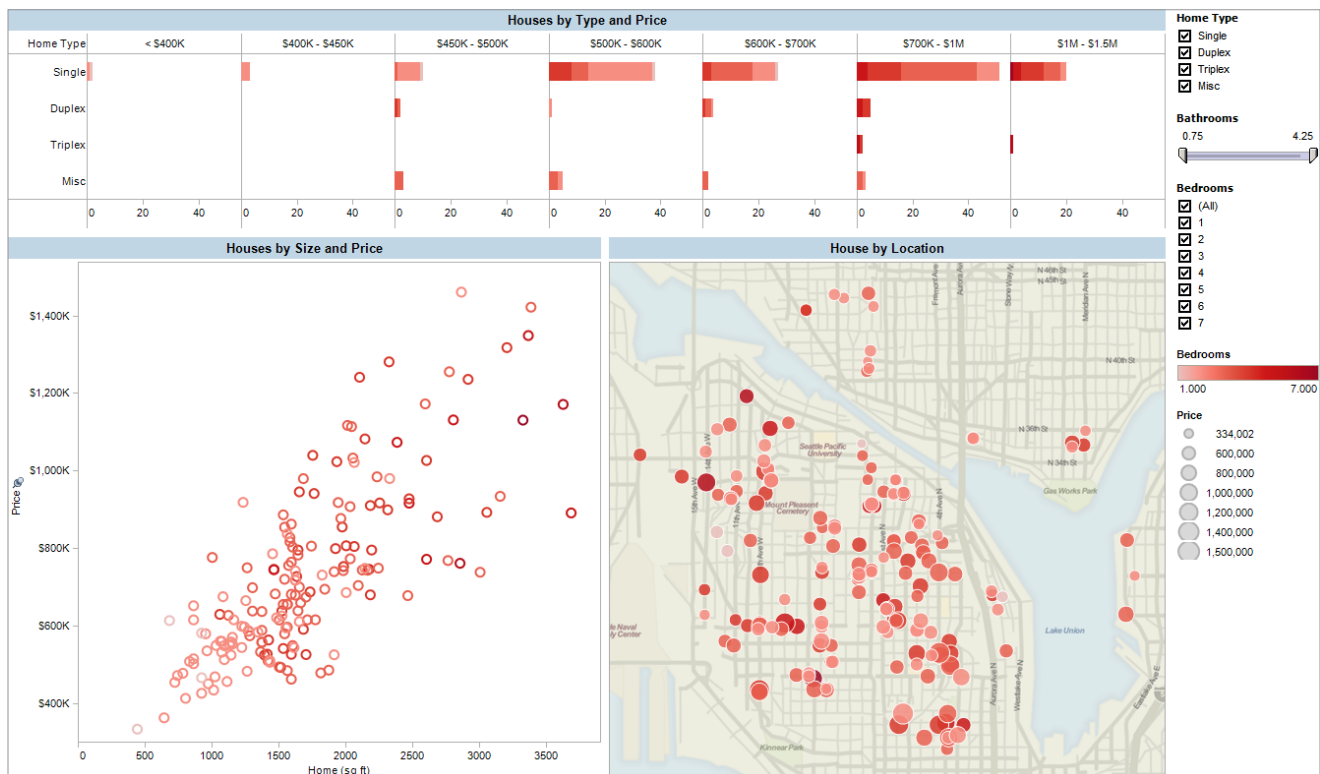
## Designing Multi-View Analytical Displays

Despite their potential benefit, the simultaneous display of multiple views is not a panacea. A multi-view display can sometimes tax the mind to a counter-productive degree. There is no simple rule of thumb such as “more than three simultaneous views are too many.” The point at which a multi-view display becomes too complex is determined by the nature of the data, the number of variables and values in each view, the nature of the task, the expertise of the person examining the data, and last but perhaps most important, the design of the display. A single graph with few variables and meager data could overwhelm our brains if it is poorly designed, while a well-designed display that incorporates many graphs could be perfectly understandable.

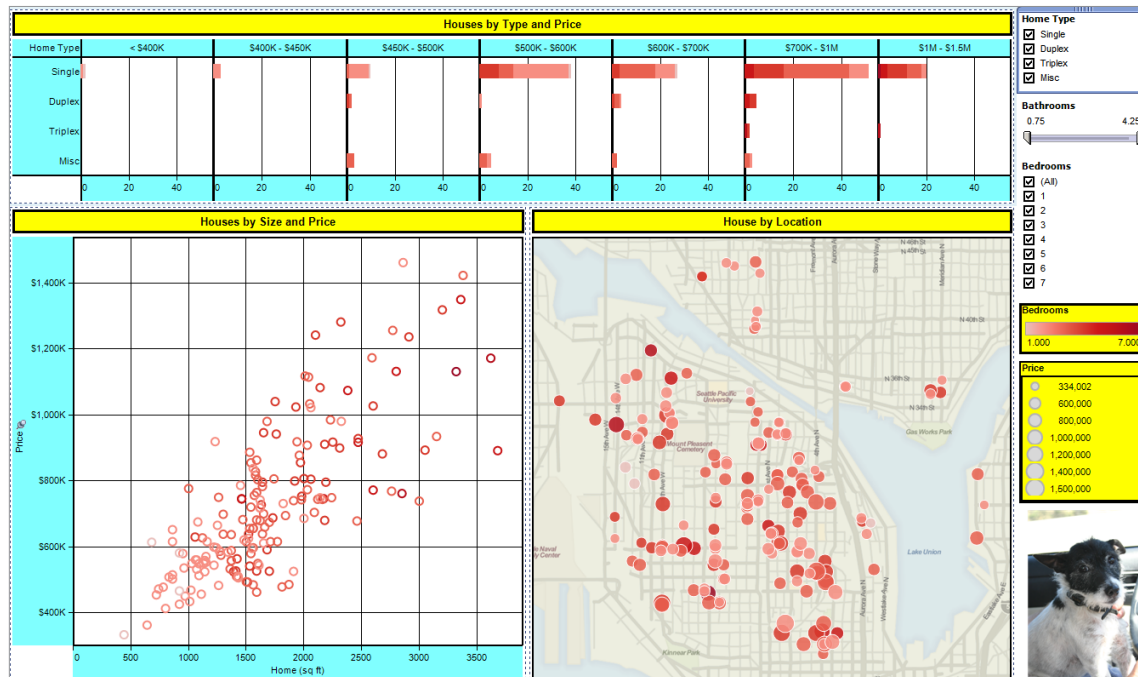
Two general guidelines for the visual design of multi-view displays are especially helpful:

- Keep non-data ink to a minimum.
- Maintain consistency among views wherever possible, except when doing so might suggest connections that don't actually exist in the data.

When combining multiple views on a screen, because we're adding more visual content for our eyes to process, it's important to eliminate all non-data content that isn't necessary to support the display of data in a meaningfully way. Axes on graphs are a common example of non-data content that meaningfully supports the display of data by clearly indicating the boundaries of the plot area where the data appears and by providing straight lines to which tick marks and labels can be attached, which tell us how to interpret the data. Wherever non-data content such as axis lines are included, they should only be visible enough to do their job effectively, but not so visible that they compete with or distract from the data. In the following example of a multi-view display, notice how easy it is to focus on the data without distraction.



Now notice the difference when unnecessary non-data ink is included and overly salient.



I bet you can't help looking mostly at the photo of my adorable dog Mangia, can you? Unnecessary and overly salient non-data ink should be avoided even in a single graph, but in a multi-view display like this, the practice is critical. Visual analysis tools that keep non-data ink to a minimum as a design default make it easy to follow this useful practice. This is important, because you don't want to waste time arranging and formatting the display to eliminate non-data ink; you want to remain in the flow of analysis without distraction or interruption.

## Two Responses to Brushing—Highlighting or Filtering

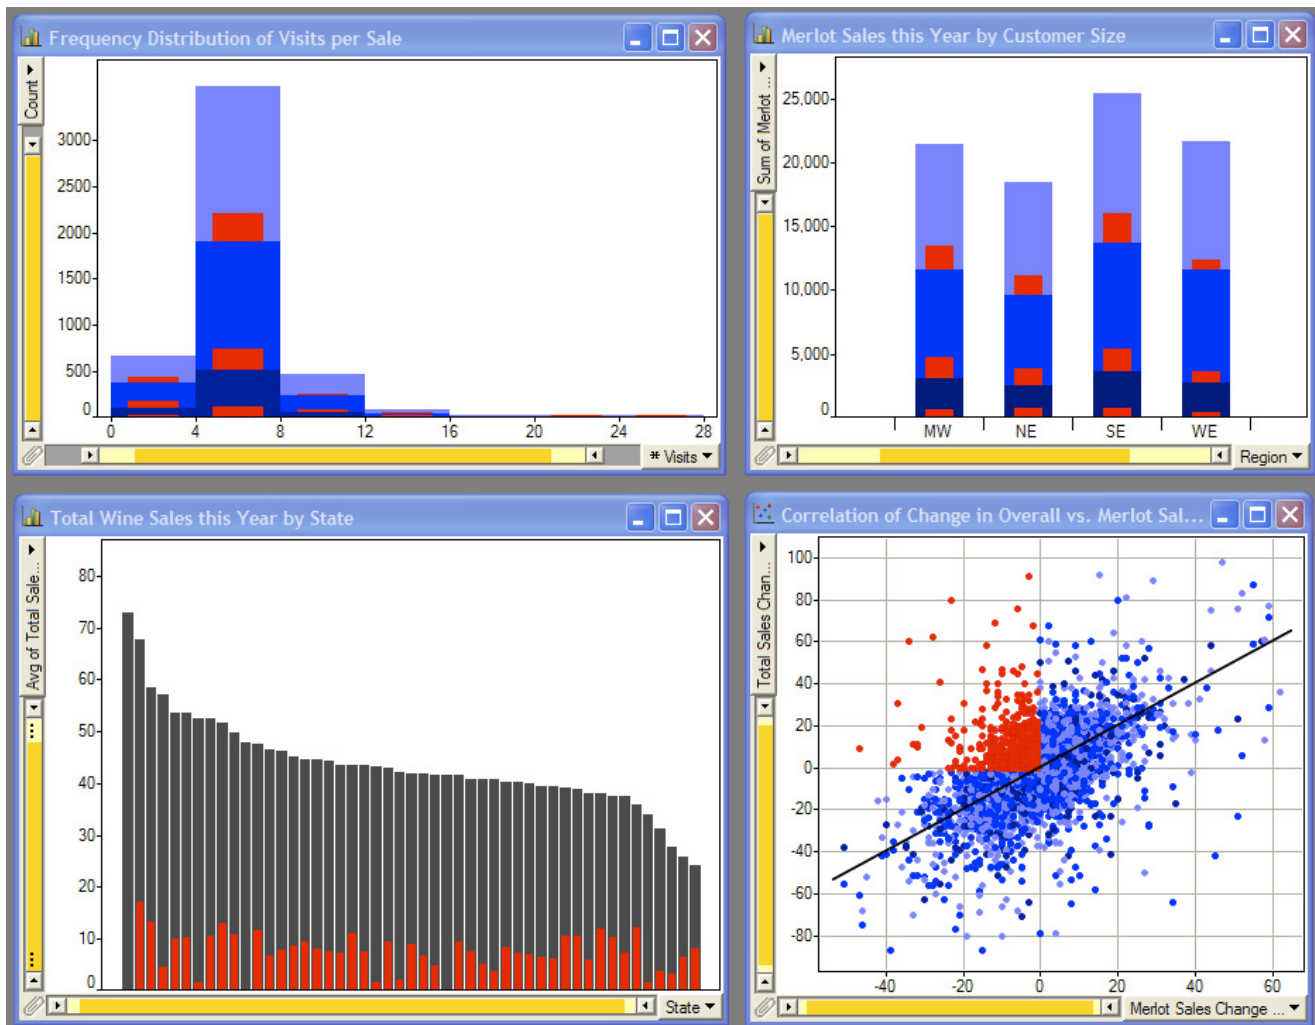
Highlighting is the usual response to brushing, but it's not the only response that's potentially useful. Another is *filtering*. Just as it's convenient to select information directly in a chart when you wish to highlight it, it's also often convenient to select a subset of information in this way when you want to focus on it by removing everything else from the display or when you want to focus on everything else by removing what you've selected. All three responses—brushing to (1) highlight what's selected, (2) remove everything that isn't selected, or (3) remove what is selected—should be easy to choose and switch between as needed during the course of analysis.

When is one response preferable to the others? For what reasons is one more useful than the others in a given situation? Each response has its strengths and weaknesses. Once you understand what they are, the choice becomes easy.

Response	Strength	Weakness	When to Use
Highlight	We do not lose sight of the data set as a whole, for unselected data remains visible, although visually deemphasized.	Even though unselected data has been visually deemphasized, it is still visible and can therefore be slightly distracting.	When it is helpful to see what we've selected in the context of the whole, and we don't want the display's appearance to change due to rescaling of quantitative axes, this response is preferable.
Filter	Only information that we wish to focus on remains visible, which allows us to view it without distraction.	We can no longer see how the information that we've chosen to view relates to the data set as a whole.	When we only wish to consider a selected subset of data as clearly as possible without distraction or influence from other data, this response is preferable.

## Highlighting Methods

To highlight a subset of information in a graphical display, we must make it look different from everything else—different enough that it stands out. Simply changing the color of selected items could do the job, but only if those items stand out distinctly from the rest in a pre-attentive way—that is, a way that can be distinguished without conscious effort. Using color to highlight works only if that color is not used elsewhere and the overall effect is not visually assaulting. In the following example, a bright red has been used to highlight selected information. Its contrast to the dark gray bars and the three shades of blue that appear elsewhere is sufficient to highlight the selected items. The combination of red and blue is a bit jarring to the senses, but not to a dysfunctional degree.



If more colors had been used in this display or colors were used that didn't provide sufficient contrast to red, however, this approach would not have worked as well. In the following version of the same display, I've changed the colors in the upper left stacked bar graph to three shades of orange and those in the scatterplot to three shades of green.




This version doesn't work as well. For instance, notice that red against the darkest shade of orange in the upper left graph lacks sufficient contrast for effortless perception. Also, with so many different colors, the overall aesthetic is less pleasing to the eye, especially the Christmas colors in the scatterplot. Although the use of a unique color can work for highlighting, this approach sometimes falls short.


Rather than using a unique visual attribute, such as the color red in the examples above, information can be highlighted in context more simply and effectively by either by increasing the visual salience of the selected items (bars, lines, data points, etc.) or by decreasing the salience of all other items. Both approaches can be effective, but require finesse to get it right.

Prior to highlighting, all data items in a display should be salient enough to be easily seen and distinguished. If selected items are then made more salient—for instance, by increasing color intensity (that is, making colors darker, brighter, or more fully saturated)—because the unselected items are still as salient as they were before, they might demand more attention than necessary, competing to some degree with the highlighted items. This problem is avoided if we take the opposite approach of decreasing the salience of all that isn't selected. This approach also has its challenges, however. If we decrease the salience of unselected items too much, we might lose our ability to discern meaningful characteristics of the data. In the following display, which currently lacks highlighting, the four colors of data points, which separate them into four regions in the scatterplot, are easy to distinguish.





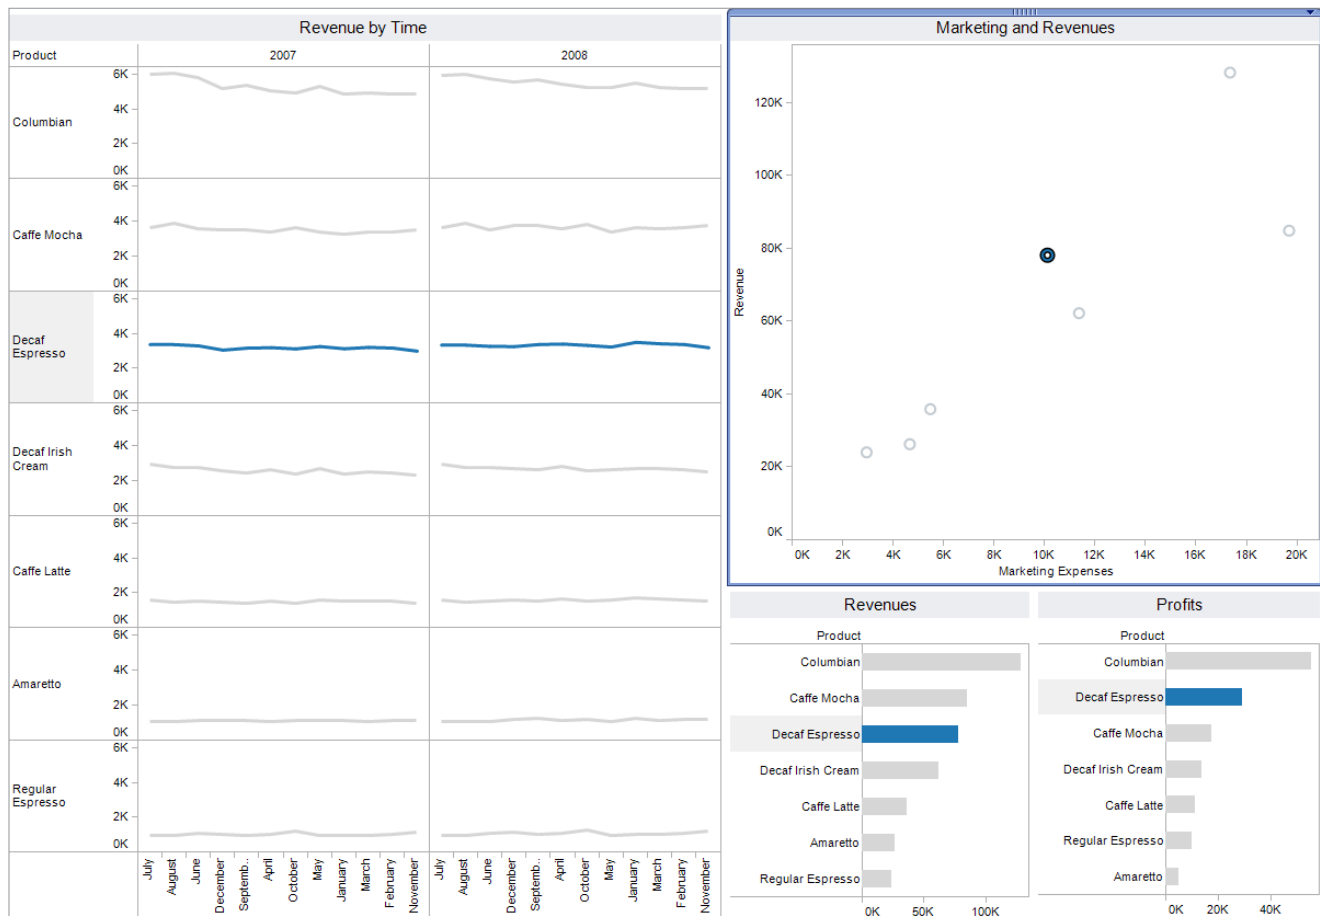
In the following example, however, after brushing the bar for the West region, notice that, although the colors of the non-highlighted bars are still possible to distinguish, those same colors in the scatterplot are no longer distinguishable.



When objects are small, such as these data points, there simply isn't enough color for their differences to be perceived if the colors are pale. Especially with small objects, it is difficult to achieve the balance between reducing the salience of unselected items enough to make the selected items stand out in contrast without reducing it so much that useful information is lost.


## Highlighting One-to-One vs. Proportional Relationships

Sometimes each of the coordinated views in a display features the same entity at the same level, such as time at the month level. For example, a collection of views for analyzing sales could each focus on products from a different perspective. In the example below, each row of lines in the series of line graphs, each data point in the scatterplot, and each bar in the two bar graphs features a specific product. By selecting a particular product in any of the graphs, such as “decaf espresso,” we could highlight it in each of the graphs, as illustrated below.




It's easy to understand the connection between the items that are highlighted in each view when they exhibit a one-to-one relationship as they do here. The relationship might not be as obvious, however, when a one-to-many relationship exists, but the opportunity for insights can grow along with the complexity.

When the relationship between the item that's brushed in one chart and one that's highlighted in another chart is one to many, you can either highlight the entire item in the secondary chart, even though what you've selected only represents a portion of it, or you can highlight only that portion. In the following example, we have a series of bar graphs: the first four show the relative number of sales orders associated with categorical items (regions, states, customer segments, and product categories) and the final two are histograms, each of which displays the frequency distribution—one for sales revenues and one for profits.



In the example below, I've brushed the "Home Office" customer segment (highlighted in dark blue), which has automatically caused sales associated with home office customers to be proportionally highlighted in each of the graphs. We can gain a great deal of insight by seeing the proportions of each item that are associated with home office customers. For example, we can see in the region graph that the portion of sales associated with home office customers is not uniform across the West, South, East, and Central regions. In contrast, the proportion of home office sales appears to be fairly consistent across the three product categories. If every bar in every graph that had home office customers were highlighted in its entirety instead of proportionally, with the exception of a few bars at the extremes in the Sales and Profit histograms, every single bar would be entirely highlighted, resulting in little insight. It's hard to imagine a situation when proportional highlighting isn't significantly superior.



As you can imagine, it is much easier for software developers to program a visual analysis tool to highlight items in their entirety rather than proportionally. Getting the tool to support proportional highlighting can get a bit tricky; it forces the tool to be a lot smarter about the underlying data.

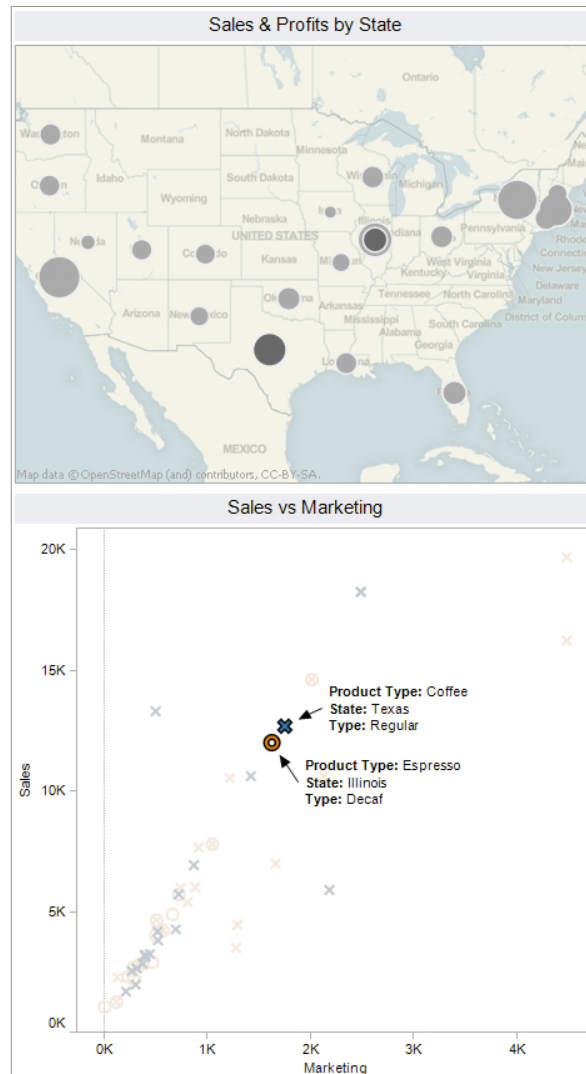
Consider the following example. Each of the four charts displays information at a different level of granularity:

Bar graph	Sales by product and market
Line graph	Sales by product and month
Map	Sales and profits by state
Scatterplot	Sales and marketing expenses by state, product type, and type (decaf or regular)

When items are brushed in a particular chart, to what level of granularity should the highlighting occur in the other charts?




Although this question might cause the software engineer who's responsible for developing the functionality to struggle with an array of possibilities, conceptually the answer is quite simple: the granularity of the items that were brushed should determine the granularity on which proportional highlighting is based. If data points in the scatterplot are brushed, then only those portions of items in the other charts that correspond to the selected cross-section of state, product type, and type should be highlighted. In the following example, I've brushed two data points, which I've also labeled to identify the cross-section of state, product type, and type (regular vs. decaf) that each represents. In the map, we can now see that those two items are associated with all sales in Texas (the entire bubble is highlighted), but only a portion of sales in Illinois (only a portion of the bubble is highlighted).



## Proportional Highlighting Encoding Methods

We need to consider one more issue: what visual means should we use to represent a portion of an object? In some cases the answer is simple. For example, to show a portion of a bar in a bar graph, we can highlight only that portion of its length, as follows:




But how do we encode portions of lines in a line graph or data points in a scatterplot? The solution is not as straightforward in cases like these.




I'll propose a general solution, but do so tentatively, because it has not been thoroughly worked out and proven. Portions of values should be encoded using the same visual attribute that encodes the whole. In the case of the bar graph above, length is the visual attribute that the bars use to encode values, so lesser lengths have been used to encode the portions. So far, so good. Will this approach work when attributes other than length are used to encode values? In the following table, I've listed the visual attributes that are commonly used to encode values in graphs, along with the specific objects that we use to encode them.

Role	Visual Attribute	Graphical Objects
Primary	Length	Bars
	2-D position	Points and lines
Secondary	Area	Bubbles (data points that vary in size)
	Color intensity	Fill colors on heatmaps


Let's start with quantitative encodings that use 2-D position. The points along a line in a line graph encode values based on their 2-D position. In the following example, horizontal positions encode categorical items, in this case months, and vertical positions encode quantitative values.



If we use lower 2-D positions to highlight a portion of these values, it might look like this.




Salience of the full values has been reduced to a lighter green line and the highlighted portion retains the original salience of the darker green. As you can see, this works quite well when you have but a single line. You can imagine, however, that it might get a bit messy with multiple lines. Here's an example with three sets of values, including a full and proportional values for each.




Not ideal, perhaps, but a better approach might not exist.

Now let's consider proportional highlighting of data points in scatterplots. Data points are placed in particular horizontal and vertical positions to encode two quantitative values, as illustrated below. In this example, marketing expenses and sales revenues are being compared, the four colors represent geographical regions (East, West, North, and South), and each data point represents a particular combination of product and state.




Now imagine that in another chart we brush a particular customer that happens to be located in California, which is in the West (orange) region. As a result, the data points associated with California and products that this customer purchased should indicate the portion of sales revenues and marketing expenses that were associated with this particular customer. Using lesser horizontal and vertical distances from zero to encode these portions could be displayed in a manner similar to this.




Once again, this could become quite cluttered and cease to work if too many items are highlighted, but this approach could otherwise do the job fairly well.

What about the secondary visual attributes that we use for quantitative encodings—area and color intensity? Portions of areas, especially when bubbles are used, are fairly easy to handle. A portion of a circle can be easily encoded as a smaller interior circle, such as I've done in the following example.




We're still stuck with the fact that visual perception is not good at comparing areas, but when we've already chosen an area-based encoding because it's the best that we can do, we are not increasing this problem by encoding the portion as a smaller interior area.

The use of color intensity for encoding values in the form of a heatmap presents a greater challenge to proportional highlighting. Consider the following heatmap arranged as a matrix, which uses variations of red to encode ranges of value for four quantitative variables.



Following the principle that I proposed earlier, we would attempt to alter the color of a cell in the matrix to display the value associated with the relevant portion of the full value, but doing this alone would cause the full value to be lost. So, we can't change the color of an entire cell, but we can change the color of a portion of the cell. Assuming that this heatmap displays summary values for a particular year and we wish to see how the first quarter alone compares to the year as a whole for a particular product—let's say "YAL056W." If we have a line chart that displays data by product and quarter, we could brush the first quarter for this product alone to highlight it in the heatmap. The results could be displayed in the following manner, using both length and color intensity to proportionally encode quarter 1's values. Because only a portion of each cell has been used to encode the selected quarter, the original value for the year as a whole has been preserved for comparison.



Much more work should be done to improve these techniques for proportional highlighting, and even with that work they'll never reach perfection due to limitations that are built into visual perception, but perfection isn't needed to bring the insights that can be revealed through proportional highlighting to light.

## Conclusion

This article is my attempt to provide an overview of brushing and linking—coordinated highlighting in context—expressed in practical terms. It is by no means comprehensive in scope, and little that I've written here is original. What it does, however, is bring together the concepts and practices that we must understand and follow to more fully tap into the benefits of this powerful visualization technique. I hope I've achieved this goal and have also stimulated renewed attempts to extend the effectiveness of brushing and linking to enhanced levels of data sensemaking. Even though brushing and linking might seem like an old, overly studied, and boring topic to information visualization researchers, we've only begun to tap its potential.

---

## Discuss this Article

Share your thoughts about this article by visiting the [Coordinated Highlighting in Context](#) 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\*](#), speaks frequently at conferences, and teaches in the MBA program at the University of California, Berkeley. 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](http://www.perceptualedge.com). Between articles, you can read Stephen's thoughts on the industry in his [blog](#).