

# Visualizing the Stock Market

Martin Wattenberg

Dow Jones & Co. (SmartMoney Magazine)

1755 Broadway

New York, NY 10019 USA

+1 212 830 9226

[mwattenberg@smartmoney.com](mailto:mwattenberg@smartmoney.com)

## ABSTRACT

We describe a new 2-dimensional visualization algorithm capable of presenting detailed information on hundreds of items while emphasizing overall patterns in the data. This display method, which builds on Shneiderman's treemap technique, makes use of both hierarchy and similarity information. We have implemented this display in the *SmartMoney Map of the Market*, a web page that reports current data on over 500 publicly traded companies.

## Keywords

Visualization, interactive graphics, treemap, investing

## INTRODUCTION

A key goal of financial journalism is to answer the question, "How is the market doing today?" What makes this question tough is that there are so many possible answers: on a given day, the market as a whole may be up, but technology stocks could be down--but Apple Computer's stock might be up. Summing up the market with one or two index values hides a lot of the action. But if you give all the details, as in the stock price pages of a newspaper, the result is hard to read and there is no way to spot overall trends. This paper describes a new interactive graphical display, related to Shneiderman's treemap diagram [3], that allows a user to track the performance of hundreds of stocks at once without losing sight of the bigger picture.

## PROBLEM

Using a treemap to provide a snapshot of stock market performance makes sense for several reasons. Hierarchical classifications of companies (into broad sectors and, within sectors, specific industries) are common. Companies also have a natural "size," namely their market capitalization. Finally, price performance can be indicated by color coding. However, previous approaches to using treemaps to display stock market data[2] present two problems.

First, the "slice-and-dice" layout method often creates a partition containing rectangles with extremely large aspect ratios. Figure 1, a slice-and-dice layout with two levels of hierarchy, clearly shows this effect.




Figure 1: Typical slice-and-dice layout



Figure 2: New layout scheme, described below

Rectangles with high aspect ratios are often hard to see, let alone compare in size. Extreme aspect ratios also make a zoom option problematic, since the zooming transformation must stretch one direction more than the other, which can confuse the user. Previous attempts to solve this problem, such as intentionally distorting rectangle sizes or eliminating the high-aspect-ratio rectangles[2,4], interfere with the accuracy of the diagram.

The second problem is that we wish to place similar companies and sectors next to each other. Doing so gives the user added cues to the relationships between companies and sectors, and increases the chance that when several similar companies experience related price movements, the

trend will appear as a contiguous region of similar colors. Unfortunately the standard treemap layout method uses only hierarchical information, and there is no canonical method to map the space of stocks onto a plane.

## DESIGN

Our design solves both of these problems: it includes a simple algorithm that produces a layout in which the vast majority of rectangles have reasonable aspect ratios. It also places companies with similar price histories near each other.

### Step 1: The initial layout algorithm

To avoid problems with high-aspect-ratio rectangles, we use an algorithm that creates a more fully 2-dimensional layout by employing both vertical and horizontal partitions at each level of hierarchy. Although certain sets of sizes make it impossible to avoid extreme aspect ratios, this new layout scheme seems to work well with real-world data. Figure 2 shows the results of the algorithm on the data from Figure 1.

### Step 2. Optimizing neighbor similarities

The resulting layout has many degrees of freedom, in the sense that there are many possible ways to rearrange the rectangles. We aim to choose the arrangement in which neighboring stocks are as similar as possible. Over the course of many trading days, this arrangement should lead to a set of the most coherent possible pictures. Moreover, using proximity as a cue to similarity adds another layer of information to the display.

To define a "similarity" measure we associate each stock (or industry or sector) with its percent price changes for each month of the past three years. Thus it is represented by a point in 36-dimensional space. We then define the distance between two items as their Euclidean distance in this space. Our goal is to find a layout in which, for each set of "siblings" in the hierarchy, the sum of the distances between all pairs of adjacent items is minimized. For most groups of sibling in our stock market visualization tool, there are few enough (under a million) possible arrangements that we can test them all by brute force; we use a simple stochastic method of estimating the minima for the remainder.

Although this method seems to work well in practice, it is just one of many possible algorithms. An interesting area of future research would be to compare the results of using other similarity measures and other algorithms for choosing an optimal layout.

### Step 3. Displaying market data

The resulting diagram is published as a Java applet on a web page, available at <http://smartmoney.com/marketmap/>

Each rectangle is filled with a color indicating the percent change of the stock since the previous market close. Following tradition in financial journalism and trading terminals, we use green to indicate upward movement and red to show downward movement. To provide maximum contrast we use black to indicate lack of movement. Since the applet is aimed at a mass audience, we also provide an alternate color scheme friendly to people with red/green colorblindness.

## INTERACTIVITY

To prevent text from obscuring the structure of the diagram, we leave out company labels. Instead we use a "brushing" interface, where moving the mouse over a company's region causes a small rectangle to appear with the company's name, ticker symbol, and a few statistics. Clicking on a company's region brings up a menu of choices for more research (pointing to other pages on the SmartMoney.com site) as well as options to zoom in or out. Since the aspect ratios of the component rectangles are relatively low, we zoom via uniform expansions, which are easy for users to understand.

## CONCLUSION

The technique described above is a simple and effective extension to the treemap display method. Because its layout scheme creates rectangles that are approximately square, objects are easy to spot, quick visual comparisons between areas are possible, and zooming can be achieved with uniform expansions. The added step of optimizing the similarity between neighboring objects makes it possible to spot trends and patterns in the data. The resulting diagram can give a comprehensive overview of a complex set of data at a glance.

## ACKNOWLEDGMENTS

We thank Marc Frons for numerous invaluable suggestions.

## REFERENCES

1. Johnson, B, and Shneiderman, B. *Tree-Maps: A Space-Filling Approach to the Visualization of Hierarchical Information Structures*. In Proceedings of IEEE Visualization '91 Conference, pp 284-291, 1991.
2. Jungmeister, W. and Turo, D. *Adapting Treemaps to Stock Portfolio Generalization*, Human-Computer Interaction Laboratory, Univ. of Maryland, College Park, MD, Sept., 1992, CS-TR-2996.
3. Shneiderman, B. *Tree Visualization with Treemaps: a 2-D Space-Filling Approach*, ACM Transactions on Graphics, vol. 11, no. 1, pp 92-99.
4. Turo, D. and Johnson, B. *Improving the visualization with treemaps: Design issues and experimentation*. Proceedings of Visualization '92, IEEE Computer Society Press, 1992, pp. 124-131.