First publ. in: 11th International Conference Information Visualization : 4 - 6 July 2007, Zurich, Switzerland / ed. by Ebad Banissi ...
Los Alamitos : IEEE Computer Society, 2007, pp. 308-315

Visual Exploration and Discovery of Atypical Behavior in Financial Time Series Data using Two-Dimensional Colormaps

Hartmut Ziegler, Tilo Nietzschmann, Daniel A. Keim
University of Konstanz, Germany
{ziegler, nietzsch, keim}@inf.uni-konstanz.de

Abstract

This work describes two pixel-based paradigms for visual financial time series data analysis that allow analyzing assets in overview and in detail, and offer improved insights into the characteristics of assets compared to traditional visualization methods. We contribute a two-dimensional color coding scheme for inter-/intra-asset analysis that extends the two paradigms and supports the discovery of significant characteristics of assets, such as atypical "against-the-market"-behavior in form of exceptional stability in case of whole market losses, or the discovery of assets in a portfolio that - while still being profitable - notably underperform the market median. We apply our techniques on real world data sets, and focus on assets of the banks of Switzerland.

1 Introduction

The financial market is an important domain dealing with large amounts of complex time dependent data every day. Companies like Reuters currently provide 50.000 updates of data per second for the world wide stock market. Although it is possible to use algorithms to automatically analyze the data for making automated decisions, a deep understanding regarding the huge amount of data can only be achieved through information visualization and visual data mining. In financial analysis, the almost exclusively used techniques by analysts for decades are still traditional line and bar charts, and new visualization techniques are only slowly entering this domain. In this paper, we illustrate several approaches for visualizing financial time series data that have been proposed in the past, and describe two conceptual models for getting better insights into the development of an asset over a period of time. We propose a basic - yet highly effective - technique that reveals the characteristics of an asset in overview and detail, its performance compared to the whole market, and special atypical tendencies that are usually not evident at first glance and that can effectively be emphasized with our technique.

2 Motivation

The financial market with its large amount of data is a challenging and interesting field for information visualization. Especially the ten thousands of assets that can be analyzed for potential investments have a high demand for suitable visualization techniques, and should not only allow an overview and detailed view on the data of one asset, but also allow comparative views in order to evaluate the performance of an asset in comparison to other assets. Higher level functions, which result from a combination of the mentioned techniques, include detection of atypical behavior of assets under certain market conditions, like exceptional stability in case the whole financial market crashes, or detection of too slowly growing assets in a portfolio while other assets have much higher growth.

In contrast to this, financial analysts and investors still primarily work with easy to use traditional line charts and statistical measures for their daily work, which on the one hand supply the basic information, but on the other hand only offer a very limited insight into the real characteristics of an asset. While few statistical numbers are per se not sufficient to describe the real market characteristics of an asset, analysts that use line charts also have no possibility to compare an asset against thousands of other assets on the market at the same time. One could say about these analysts: they see some single words, but they will not see or understand the full sentence.

In this paper, we describe two pixel-based paradigms that allow to having an overview as well as a detailed view on the performance of an asset. We describe, how these two concepts can be used to analyze the growth of an asset over time, and how it can be used to compare the performance of an asset against the whole market or market sectors for all possible time intervals. Our visual approach supports the discovery of patterns, trends and correlations in different markets, and allows comparative screening in order to support the decision making process.

3 Related Work

Large time series data in financial markets offer many possibilities for developing improved visualization techniques in order to get more insight into the market behavior. However, Information Visualization is only slowly entering this domain, and analysts and investors still primarily focus on known techniques like line charts for making their decisions [11]. Several approaches in recent years have tried to visualize the financial market with maps, like treemaps [16], Neovisions RiskMaps [12], Smartmoneys 'Map of the Market' [15] or Self-Organizing Maps [1] [14] [4]. These techniques offer a good overview over the financial market, but do not allow any detailed analysis or comparison of single assets, neither analysis of effectively more than one interval of time and therefore cannot replace traditional line charts. Commercial products for visualizing financial time series data are GSphere by Gravity Investments [5], Market Topology [10], Portfolio Impact by High Tower Software [8], or StockVIS [13]. Other well known projects are TimeSearcher [6] or ThemeRiver [7]. One of the most recent approaches for visualizing financial time series data is based on pixel-based triangular matrices [9], which allow analysis of the development of an asset in overview and detail for one or many assets. Except the last, none of the other approaches allows the visualization of large amounts of time periods, and therefore do not provide sufficient information for an analyst for making reasonable investment decisions.

4 Pixel-based Concepts for Visualizing Financial Time Series Data

4.1 The Growth Matrix Approach

Regarding financial time series data, as the financial market focuses on profits, analysts and investors are primarily interested in growth rates and not absolute values of assets. A growth rate can be defined as the growth of an asset in a time interval from the time of purchase (starting point) to the point of sale (end point). As line charts do not reflect the growth rates directly and do not allow intuitive perception of all possible growth rate intervals at a glance, a different approach for visualizing growth rate intervals is required.

As all possible combinations of "time of purchase" and "time of sale" for an asset should be covered by our visualization, the one-dimensional time series data is transformed into a two-dimensional spectrum with two axes in a Cartesian coordinate system, mapping "time of purchase" on the x-axis and "time of sale" on the y-axis, thus spanning a triangle covering all possible $n * (n-1)/2$ time subintervals of an asset (see Figure 2). Each time subinterval is represented by one pixel, with the result that showing all time intervals in one single view is possible at the same time. This technique is called *Growth Matrix* [9].

As both axes of the 2D-coordinate system are used to represent the time intervals, the color is used to indicate and encode the value of the growth rate at each position *(x,y)*. The color coding scheme forms a smooth transition from shades of red (losses) over beige (unchanged growth) to green (for profits), see Figure 1.

Figure 1: The color map covers a range from shades of red (negative growth, loss) to beige (no growth) to green (positive growth, profit)

For generating our images, we use a real-world database that contains data of 12.000 funds which consists of various types of assets (bonds, stocks, property) and is sampled on a monthly basis from January 1991 to March 2005. The 171 months therefore create a maximum of 14535 possible combinations of *time of purchase* and *time of sale*, each with its individual growth rate. For illustrating our examples in this paper, we have selected assets from banks in Zurich (Switzerland).

Two examples of such pixel-based layouts can be found in Figure 2. The example on the left illustrates the intra-asset analysis of a bond with a nearly continuous steady growth over time, accordingly the positive growth rates of all possible subintervals form a completely green triangle illustrating the positive growth for all possible combinations of time of purchase and time of sale. The image on the right illustrates the intra-asset analysis of an asset that consists primarily of technology shares, with significant growth rates and profits if purchased in the early years and sold in 2000, and high losses after 2000 during the breakdown of the "new economy" in the so-called "dot-com" crisis.

The examples reveal typical growth rate patterns, identifying strong and weak regions where an asset has been profitable, mediocre or lossy. They also allow analysis of the short-term volatility along the diagonal, or long-term behavior of the assets for up to eleven years holding time the further we move to the upper left corner of the image.

The *Growth Matrix* therefore can be defined as a function $GM : (Integer, Integer) \rightarrow Color$, mapping discrete 2D coordinates (i, j) to color codes taken appropriately from a color space. Regarding growth rate normalization, we define two metrics. The first one is called *growth index*, and is a function $gi^A(i, j)$ based on relative growth

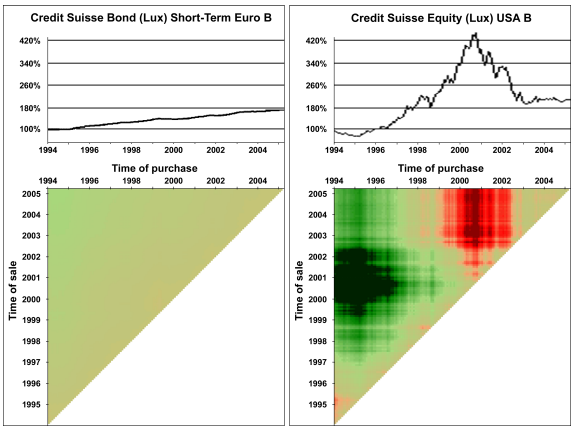

Figure 2: Two examples of Growth Matrices with prototypical patterns: The image on the left illustrates a bond with a continuous growth, resulting in a completely green triangle as all possible combinations of times of purchase and times of sale have a positive growth. The image on the right shows a highly volatile fund that consists primarily of technology stocks, with high profits towards the year 2000 and high losses after the crash of the new economy during the "dot-com" crisis, with corresponding green and red regions in the Growth Matrix triangle for the affected combinations of times of purchase and times of sale representing the absolute profits and losses.

of the price of an asset A within the interval $[i, j]$:

$$gi^A(i,j) = \frac{v_j^A}{v_i^A}, \quad i,j \in T, \quad i < j,$$

where i and j are the indices of time stamps corresponding to valid start and endpoints in a global time interval T as defined by the respective time series. v_i^A and v_j^A are the observed prices of asset A at time indices i and j. Then, $gi^A(i,j) \in [0, \infty)$, also referred to as *growth coefficient*, is the relative increase (decrease) of the time dependent variable v in the interval $[i, j]$.

The *Growth Matrix* so far only reflects the characteristics of the performance of an asset, and does not allow to compare the funds one to each other which would be even more interesting for investors because for a specified time frame a different fund might have been the better investment with a higher growth rate and return. To allow a relative inter-asset comparison of how good a fund has performed compared to the whole market, for *each* of the 14535 positions *(x,y)* of the matrix we can compute the growth rate for *all* 12.000 funds from our database for the corresponding time interval, create a sorted list of 12.000 elements for each pixel, and determine the ranking of the asset which is then linearly mapped across the spectrum of our color map, with the worst and best performing assets in the database on the opposite ends of the color coding scheme.

We can therefore consider the second growth metric as *rank index* as a function $ri^A(i, j, DB)$, defined as the $[0, 1]$-normalized rank of a given growth coefficient with regard to all corresponding growth coefficients in a database of time series DB. Formally:

$$ri^A = \frac{1}{|DB| - 1} \sum_{A' \in DB \wedge A' \neq A} \begin{cases} 1 & gi^{A'}(i,j) < gi^A(i,j) \\ 0 & \text{otherwise} \end{cases}$$

This descriptor allows to compare the performance of a fund against the whole market for all possible time intervals, Figure 3 illustrates an example. For creating a space-efficient view on both absolute and relative-to-market performance, the *Growth Matrix* for inter-asset comparison is placed in the large triangle, and the *Growth Matrix* for the growth rate is placed in the small triangle in the lower right corner.

We perceive by the small green triangle that while the bond never creates any absolute loss, in the times before the year 2000 the relative comparison shows large red areas, indicating that other funds would have been a better investment than the bond because the share market had higher growth rates compared to bonds. Vice versa, we can derive from the green areas between 2000 and 2002 that in times where the stock market crashed, a bond with a moderate positive growth has outperformed most other assets in the database and would have been a good investment.

It is also possible not to match against all 12.000 funds, but to match against a specific type of asset (bonds, equities, property), market sector (technology, energy, metals), market region (USA, Asian, Europe) or any other selection in order to compare a subset of funds one to each other. As can be identified on the images, our technique creates unique fingerprints for each asset, revealing local and global asset fund characteristics with prototypical patterns for visual analysis and exploration.

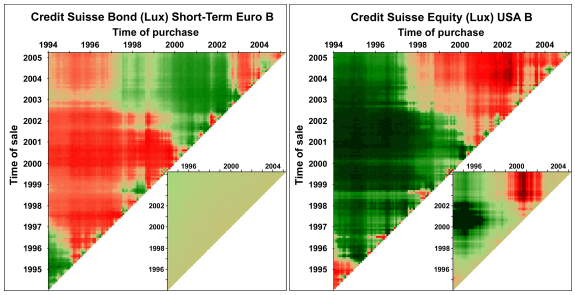

Figure 3: This image compares the performance of the two assets from Figure 2 against the market. The small triangles illustrate the absolute intra-asset analysis, the large triangles a comparative (relative) inter-asset analysis against a database of 12.000 funds.

4.2 The Performance Matrix Approach

While the *Growth Matrix* offers many new insights into the characteristics of an asset compared to line charts, it is often difficult to give intuitive answers to several specific interests of analysts and investors. One of the main interests for financial investments is the long-term performance of an asset, especially the 1, 3 and 5 years holding time. These statistical values have the disadvantage that they only describe 3 out of several thousand possible time intervals, and do not reflect any important behavior of the asset that has taken place between the three time intervals. It is even possible that the "dot-com" crisis with a duration of nearly 2 years is completely missed out by the three statistical values because it is completely located between the time intervals [17]. This incompleteness is inappropriate for a reasonable financial investment. The three performance values can also be derived from the *Growth Matrix*, but are difficult to read as the "time of sale" on the y-axis and the holding period (time of sale minus time of purchase) on the diagonal line form a 45 degree angle.

A conceptual approach to overcome this is a two-dimensional rectangular approach, as sketched in Figure 4. It allows simultaneous comparison of 4500 combinations of times of sale and holding periods in one space efficient pixel-based visualization, forming a *Performance Matrix PM* with 90 degree angles and giving a much deeper, more reliable and easier to interpret insight into the long-term characteristics of an asset. In order to illustrate the concept, the statistical measures $PM(2005, 1y)$,$PM(2005, 3y)$ and $PM(2005, 5y)$ for 1, 3 and 5 year performance are denoted in Figure 4.

The performance for each combination of holding period and time of sale can be described as a metric function $PM^A(s, h)$:

$$PM^A(s,h) = \frac{V^A(s)}{V^A(s-h)} \quad s, s-h \in T$$

where s is the time of sale, $s-h$ is the time of purchase of an asset A taken from the set of all possible points of time T, and h is the holding period. We use the same color map as in Figure 1 for colorizing the values.

Accordingly, we can use the rectangular model both on the absolute and relative analysis of an asset. Figure 5 illustrates the concept applied on the two already known examples from Figure 3. While the bond (left) with its continuous moderate growth rate is a stable and risk-free long-term investment for all holding periods, the technology fund (right) has been a good long-term investment until the year 2000, but even to the year 2002 if the asset has be bought before 1998. This also explains the diagonal structures that can be seen in the right image. Positive or

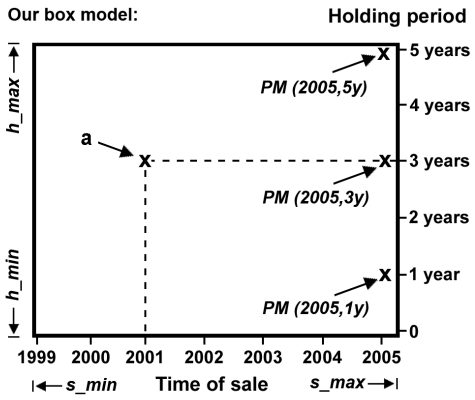

Figure 4: This figure shows the rectangular *Performance Matrix* model for analyzing characteristics of long-term investments. This pixel-based approach visualizes 4500 combinations of time of sale and holding periods in a single view by colorizing the growth rate for the corresponding time intervals. If an asset has been a good 3- or 5-year long-term investment can be analyzed by the corresponding horizontal line. The point "a" for example denotes the 3-year-performance of an asset with a time of sale in January 2001.

negative impacts *i* on the x-axis create diagonal effects, because the pixel at position *x* + $\triangle x$ for the time of sale and a holding time of $\triangle x$ refers to the same point of the impact *x* for calculating the corresponding growth rate. The short-term volatility of the asset can be recognized by the quickly changing red and green pixels directly along the x-axis.

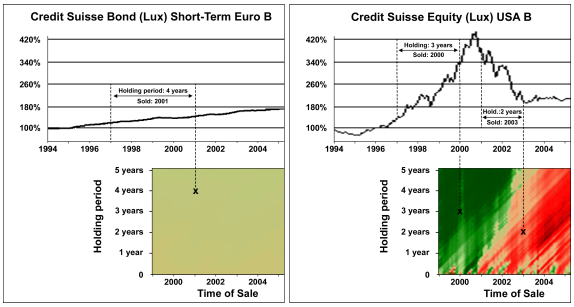

Figure 5: These images illustrate the same assets as seen in Figure 3 with focus on long-term investments. Instead of time of purchase / time of sale, now time of sale / holding time form the two axes of the matrix. The long-term capability for a specified holding period can easily be derived from the corresponding horizontal line.

5 Improvement by 2-Dimensional Color Map

5.1 Motivation and Conception

Both conceptual models in the previous chapter are well suited for evaluating global asset returns and for comparing the performance of funds for a large amount of time intervals among a large database. However, for certain constellations and to answer important questions, the two separate visualizations for absolute growth and relative ranking are not the optimal choice. To give two examples, a fund can have absolute profits, but can still perform worse than the market average making a different asset much more attractive (for example, see the red region in Figure 3 (left) for the bond between 1997 and 2000 where other assets have outperformed the bond even though the bond had continuous absolute profits). In contrast to that, there are also situations where almost the whole stock market has losses, and the investor is interested in finding assets for a long-term investment that tend to have the most stable behavior in such situations. In order to quickly find answers to these questions, the investor has to analyze both *Growth Matrices* for intra-asset and inter-asset performances for a large amount of assets, and create a mental overlay of the two triangles for each asset in order to derive which of the red areas in the inter-asset *Growth Matrices* have a corresponding green area in the intra-asset *Growth Matrices*, and vice versa. This is impossible to accomplish when trying to analyze a set of hundred funds, because the analyst has to keep each overlay image in mind in order to compare it to the others. In order to facilitate the decision making process and to reduce the mental overhead, the problem of combining the intra-asset and inter-asset *Growth Matrices* can be solved by using a two-dimensional color map, visualizing all four possible combinations of intra-asset/inter-asset and positive/negative growth/rank with a separate color in a single *Growth Matrix* (see Figure 6).

The color map is defined by a two-dimensional grid, with inter-asset performance on the x-axis and intra-asset performance on the y-axis. Therefore the color map is tiled into four regions with a different color representing each region, and a color in the middle considered as ”neutral” which originates from the middle of the 1D color map (Figure 1 in chapter 4.1). The colors form smooth transitions from one region to the other and towards the middle. Several color models have been explored to create the two-dimensional color map (complementary colors, CMYK color model, and others) in order to find suitable colors for the four corners, as well as several algorithms to compute the color transitions from region to region, but as cognitive recognition and intuitive associative correlation of the colors to the meaning of the corresponding corner is not a mathematical but a perceptual problem, we selected the colors of the four corners manually with respect to an intuitive recognition of the corresponding meaning: On the horizontal split axis, in the upper half we chose ’positive’ colors like ”yellow” and ”green” for absolute profits, and ’negative’ colors ”red” and ”blue” in the lower half representing the absolute losses. Accordingly, we split up the vertical axis, but as the region for ”better inter-asset performance” should be filled with a ’positive’ color and the lower right corner also represents absolute losses at the same time, it would be contradictory to fill this corner with a ’positive’ color like yellow (accordingly for the upper left corner). As a conclusion, we decided to set the focus on the intra-asset performance, as an absolute profit is more attractive to the investor than having an absolute loss yet having one asset that only had the smallest loss. Therefore, we use ”yellow” in the upper left corner. These colors have also been chosen because they maintain differentiable color transitions between the four regions with respect to the whole color map. The colors for the corners and the center as well as various transition algorithms can also be modified with a color chooser tool according to user preferences if the user prefers different colors.

While appearing to be complicated if applied on the assets for the first time, after getting used to the new color map it allows an instant visual analysis of the characteristics of an asset within few seconds.

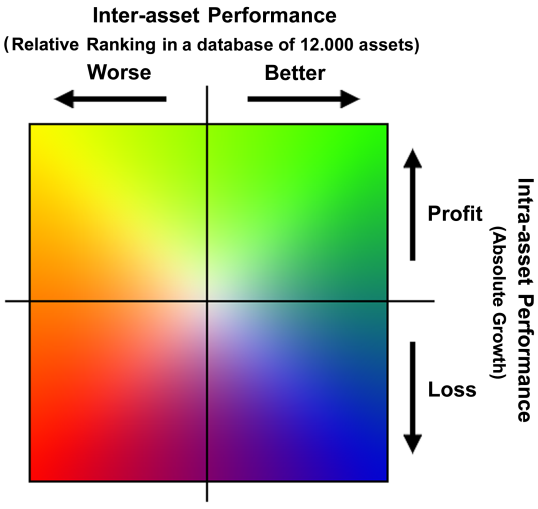

Figure 6: Our two-dimensional color map is split into four regions, representing all four combinations of positive/negative and intra-asset(absolute)/inter-asset(relative) performance values of assets. The colors for the corners and for the neutral center of the map have been chosen manually according to the semantic correlation that the colors intuitively have (like green for profits and red for loss).

5.2 Application

Applied to real-world examples, Figure 7 illustrates the effect of this idea if used in conjunction with the previous assets in Figure 3 in chapter 4.1. Figure 7 (left) now clearly reveals large areas in orange color, where the bond - while still being profitable - has significantly underperformed the market median. The technology fund (right) shows a similar structure as in Figure 3, the red area is surrounded by an orange area which reflects the time intervals where the technology funds has been a bad investment compared to other assets due to the high losses in this market sector, but has still been profitable in the long term regarding the time interval between purchase and sale.

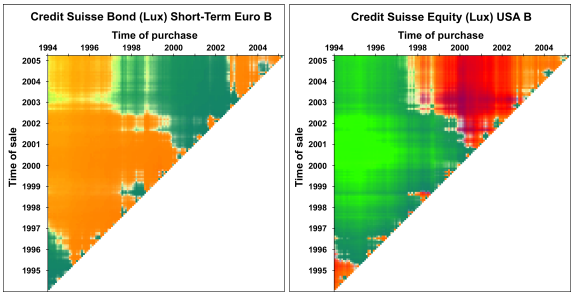

Figure 7: Our two-dimensional color map applied on the assets from Figure 3. The asset on the left instantly reveals by the orange regions where it has performed weak compared to the whole market, while still having an absolute growth. The color map instantly reveals that this asset is growing slower than the rest of the market, and a different fund would have been the better investment in this time frame.

Accordingly, this idea can also be applied on the rectangular Performance Matrix from chapter 4.2. for evaluating the long-term performance of assets (see Figure 8). If getting familiar with the four colors of the color map, an investor instantly sees within seconds that the bond (left) is composed of the two colors which represent absolute profits, and there are time frames where other investments have been better a better choice (orange color) and times where this asset was optimal (green color).

Apparently, it could be anticipated that assets that generate absolute profits also tend to have a better relative performance ranking compared to other assets and vice versa, so due to this correlation using a two-dimensional color map like in Figure 6 the colors that are displayed in the visualization should group along the diagonal axis from red to green. In fact, it could already be seen in Figure 7 with the strong orange color that assets with continuous solid profits like bonds often underperform the market during the boom times of the share market, and therefore the color spectrum does not only consist of the colors along

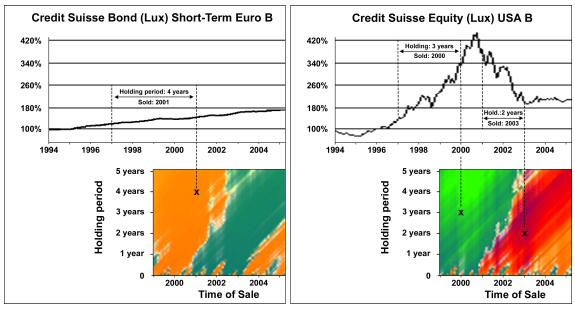

Figure 8: Our two-dimensional color map applied on our Performance Matrices from Figure 5. With little practice, an analyst instantly sees by the yellow and green color in the left image that there have been no losses for the investment, although there have been other assets that have had higher growth and higher return for the orange regions of the matrix. The nearly optimal choice here would have been to have the asset on the right side from the beginning up to 2001, and then have the left asset from 2001 to 2005, as the combined rectangle of these two assets and holding times results in an almost completely green Performance Matrix.

the diagonal axis but also consists of colors that originate from orthogonal regions of the color map.

This can be seen in Figure 10 which illustrates six assets from Swiss banks, giving an instant impression of the real characteristics of a fund for all possible time intervals and which has not been perceptible before while using separate *Growth Matrices* for inter-asset and intra-asset comparisons. It reveals deviations from market specific general patterns at one glance, and allows to identify structural similarities and uncorrelated ”against the market”-behaviors by yellow and blue color instantly, which is difficult to recognize by having only traditional line charts. The blue color, for example, indicates that in case of an absolute loss, the fund has still performed as one of the best assets in the database, indicating an extraordinary stability of the fund towards a general negative market tendency.

With respect to the two-dimensional color map, an optimal fund would consist of mostly green color, meaning absolute intra-sset growth over all time intervals while still outperforming most other assets on the market. By visually screening hundreds of assets from Swiss banks in the database on a large screen powerwall display, with this suitable visualization technique it only takes few seconds for the human eye to derive the assets that almost fulfill these conditions. Therefore, this visualization technique can greatly improve the exploration and speed up the decision making process for the financial analyst. Figure 9 illustrates the best performing fund in the comparison of all Swiss assets in our database.

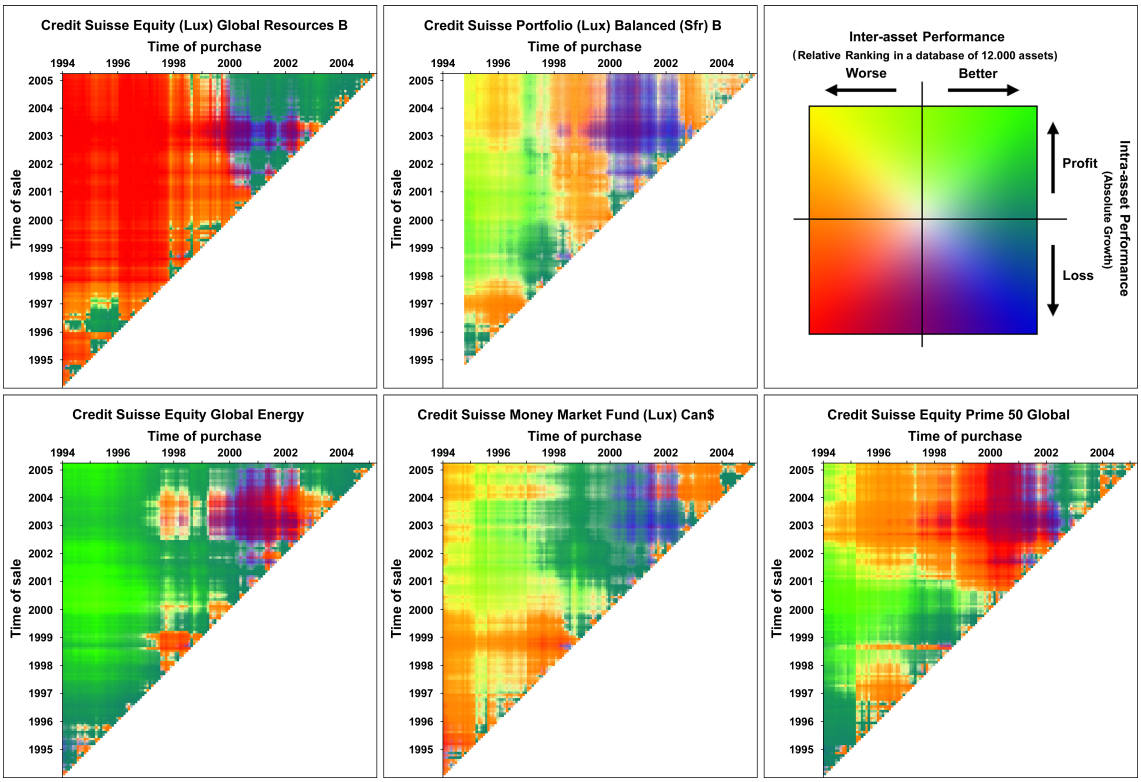

Figure 10: These images illustrate six typical assets from Swiss banks. It can be perceived that each asset not only generates its own unique fingerprint, but also comprises colors from orthogonal regions of the color map (see Figure 6) that are not arranged in a spectrum along the diagonal between red to green. This new visual approach reveals which assets behave different than most other assets in the database. The images shows large regions in yellow and blue, emphasizing regions where assets have made absolute profits while still being outperformed by most other assets and therefore other assets would have been a better investment (yellow), as well as regions where assets made absolute losses while still performing better than most other assets on the market which illustrates an exceptional stable behavior of an asset towards general negative market tendencies (blue).

6 Future Work

In addition to further technical improvements of this project, a usability study that evaluates the usefulness of this technique is planned. Due to the fact that this is a special visualization technique developed to support financial experts, the candidates for the usability study will also be taken from this group.

Conclusion

In this paper, we have described two pixel-based paradigms for visual financial time series data analysis that allow analyzing assets in overview one to each other and into smallest detail. Compared to traditional line charts, our visualization technique offers powerful insight into the characteristics of assets in one view, by covering all possible combinations of time of purchase / time of sale, time of sale / holding time, and inter-asset analysis by comparing the performance value for each pixel of an asset with the performance values of 12.000 other funds in the same time intervals. We have presented a two-dimensional color coding approach for inter/intra-asset analysis that can be applied to both pixel-based paradigms, and which offers significantly improved insights into the characteristics of an asset by revealing uncorrelated atypical ”against-the-market”-behaviors in form of exceptional stability against general market losses, or general low relative performance with respect to the overall market, which are difficult to perceive by using separate visualizations for intra-asset and inter-asset performance or by using previous visualization techniques.

Our methods show how modern visualization techniques can effectively extend the knowledge discovery and decision making process in financial markets for analysts and investors in one single view, with valuable insights that usually cannot be easily derived with any traditional visualization method.

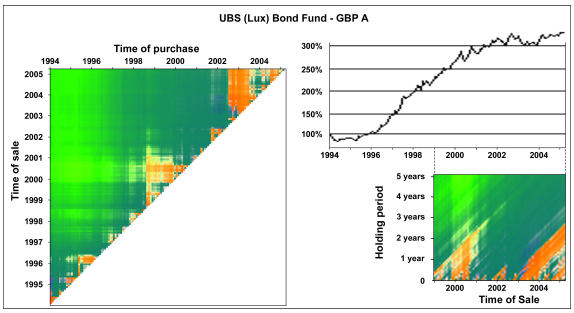

Figure 9: This fund represents one of the best performing candidates of Swiss assets in our database. Nearly all possible time intervals show a clear green color, which reflects that with respect to our two-dimensional color map both intra-asset growth and inter-asset ranking have been very positive. In fact, this fund is not only the best performing Swiss asset, but represents one of the best funds in our complete database of 12.000 assets.

Acknowledgements

This work was funded by the German Research Foundation (DFG) under grant GK-1042 "Explorative Analysis and Visualization of Large Information spaces".

References

[1] Barbro Back, Mikko Irjala, Kaisa Sere, Hannu Vanharanta, *Competitive Financial Benchmarking Using Self-Organizing Maps*. Department of Computer Science, Series A No 169, Abo, 1995

[2] Deutsches Aktieninstitut e.V.: Return triangles, url: http://www.dai.de/, 2005

[3] Tim Dwyer, David R. Gallagher. *Visualizing changes in fund manager holdings in two and a half dimensions.* In *Information Visualization 3, 4 (2004), 227-244*

[4] Guido J. Deboeck, Tuevo K. Kohonen. *Visual Explorations in Finance with Self Organizing Maps*, Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1998

[5] Gravity Investments, Inc.: GSphere, url: http://www.gravityinvestments.com/product_tour.htm

[6] Hochheiser, Harry; Interactive Graphical Querying of Time Series and Linear Sequence Data Sets, PhD. Dissertation, University of Maryland, Dept. of Computer Science, May 2003

[7] Havre, Susan, Hetzler, Elizabeth G., Nowell Lucy T.: ThemeRiver: Visualizing Theme Changes over Time, Proceedings of the Infovis 2000, p. 115-124

[8] High Tower Software, Inc.: Portfolio Impact, http://www.portfolioimpact.com/

[9] Daniel A. Keim, Tilo Nietzschmann, Norman Schelwies, Jörn Schneidewind, Tobias Schreck and Hartmut Ziegler. *A Spectral Visualization System for Analyzing Financial Time Series Data.* In *Proceeding EuroVis 2006: Eurographics/IEEE-VGTC Symposium on Visualization (EUROVIS'06)*, Lisbon, Portugal, May 2006

[10] Market Topology: url: http://www.market-topology.com

[11] John Murhpy (Ed.): *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*, Prentice Hall Press, 1999.

[12] NeoVision Hypersystems Inc.: HeatMaps and RiskMaps, url: http://www.neovision.com

[13] Edward Parrish. *StockVis: An Internet-Based System for Visualizing Stock Market Data. Master Thesis*, UC Santa Cruz, Department of Computer Science, 2000

[14] Kresimir Semunic, *Visualizing of Stock Market Charts.* In *Proceedings from the 11th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision 2003* (2003), Plzen-Bory (CZ), 2003

[15] Smart Money: 'Map of the Market', http://wwww.smartmoney.com/marketmap/ , 2005

[16] Martin Wattenberg. *Visualizing the stock market.* In *CHI '99: CHI'99 extended abstracts on Human factors in computing systems*, New York, NY, USA, 1999. ACM Press, pp. 188-189.

[17] Hartmut Ziegler, Tilo Nietzschmann, Daniel A. Keim. *Relevance Driven Visualization of Financial Performance Measures*, In *Proceedings of the EuroVis2007: Eurographics/IEEE-VGTC Symposium on Visualization (EUROVIS'07)*, Norrköping, Sweden, May 2007