

Task-At-Hand Interface for Change Detection in Stock Market Data

Carmen Sanz Merino*
MicroStrategy Ibrica, Madrid, Spain

Mike Sips†
University of Konstanz, Germany

Daniel A. Keim‡
University of Konstanz, Germany

Christian Panse§
FGCZ - Uni|ETH Zurich, Switzerland

Robert Spence¶
Imperial College London, UK

ABSTRACT

Companies trading stocks need to store information on stock prices over specific time intervals, which results in very large databases. Large quantities of numerical data (thousands of records) are virtually impossible to understand quickly and require the use of a visual model, since that is the fastest way for a human brain to absorb those enormous collections of data. However, little work has been done on verifying which visualizations are more suitable to represent these data sets. Such work is of crucial importance, since it enables us to identify those useful visual models and, in addition, opens our minds to new research possibilities. This paper presents an empirical study of different visualizations, that have been employed for stock market data, by comparing the results obtained by all studied techniques in typical exploratory data analysis tasks. This work provides several research contributions to the design of advanced visual data exploration interfaces.

1. INTRODUCTION

Today's financial community is being bombarded by massive amounts of information from real-time data sources. The market consists of thousands of companies quoting in different stock markets in many countries around the world, with a continuously increasing number of private and public investors selling and buying stocks. Every few seconds hundreds of trades characterized by stock price and trading volume are accomplished and observed through seconds, minutes, hours, days, months, and years. The impressive stock ownership increase and the rapid and drastic movements of the stock market have obviously led to an increase in the investors' need for information about the market. All

*c.sanz@gmx.net

†sips@dbvis.inf.uni-konstanz.de

‡keim@dbvis.inf.uni-konstanz.de

§cp@fgcz.ethz.ch

¶r.spence@ic.ac.uk

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AVI '06, May 23–26, 2006, Venezia, Italy

Copyright 2006 ACM 1-59593-353-0/06/0005...\$5.00.

daily generated stored stock prices possess potentially important information, however, the manual analysis of these thousands of data sets exceeds the human capacities by far. Finding valuable details that reveal fine structures hidden in the data, however, is difficult. For data exploration to be more effective, it is important to include the human being in the data exploration process and combine their flexibility, creativity, and general knowledge to enable a better interpretation and evaluation of the probably plentiful results, and hence, take the most out of them.

Techniques such as bar charts, line graphs and tables have been employed in the past and are used very frequently for the representation of stock market data in the present. During the last few years, however, important changes have occurred in the field of information visualization, leading to a rapid growth in the development of new techniques and tools for supporting visual data exploration. Some of those have also been successfully applied to our area of interest. It is, however, difficult to judge the effectiveness of the achieved visualizations. In order to solve this problem and whatever the aim of the visualization, it is becoming increasingly necessary to perform evaluations of the corresponding tools or techniques in order to identify the most reliable systems, to determine if they are useful and usable, and to steer their development in the future.

A positive trend is that visual feedback plays an increasing role in the visual analysis of stock market data. New approaches have turned up, which attempt not only to make the visualization more effective by increasing the participation of the human, but to let the human interact with the visualization as well.

2. TASK-AT-HAND INTERFACE

Presenting data in an interactive, graphical form often fosters new insights, encouraging the formation and validation of new hypotheses to the end of better problem-solving and gaining deeper domain knowledge. Although a visualization is a powerful tool, there are two facts connected to it that limit its possibilities. These limitations are the human ability to understand the visualization and to distinguish displayed interesting patterns and the suitability of the visualization to represent the corresponding data sets properly. The goal is to identify techniques that do not just increase human insight and understanding but also lead to knowledge.

2.1 Interface Design

According to the conclusions arrived at in the next sec-

Type of Action	Task
Identify	Identify the best period of the stock "T" (AT&T Corp.)
Locate	Locate the stock "IBM" and the approx. price it reached in May 1999
Distinguish	Select the stock or stocks you would invest in right after the displayed period, if you had \$3000
Categorize	Do you see different classifications in the displayed data?
Cluster	Choose two stocks that showed a similar behavior during the same period
Rank	How are the different stocks in the image ordered?
Compare	Compare the performance of the stocks "KO" (Coca-Cola) and "MRK" (Merck)
Associate	When did most of the stocks reach their highest price in their history?
Correlate	Could the performance of Intel have had an influence on the performance of Microsoft?

tions, a combination of different techniques is in this case required, since no single technique can support all tasks-at-hand scenarios separately. Our purpose is thus to design an interface, which takes advantage of the strengths of the visualization techniques, which proved to be the most suitable ones during the evaluation, and to overcome their weaknesses by combining them with each other by linking and brushing methods. The resulting system should be able to support the end-user in the decision-making process by primarily enabling the performance of all tasks-at-hand, that lead to the discovery of knowledge extracted from the data and providing better than or similar results to the ones obtained by automatic methods.

Two different combinations of visualization techniques have been thus selected for this application, in order to fulfil not only our but also the users' requirements for the two identified stock market data types, respectively. Our resulting design consists of two main windows placed next to each other and joined by linking mechanisms. The first one displays just historical market data (see figure 5) and the second one shows current intra-day stock information (see figure 6).

2.2 Tasks from the Community

The elemental tasks were chosen from the taxonomy of Zhou and Feiner [9]. The full taxonomy contains about 50 tasks. Although a number of tasks and visualization techniques are available for stock market data, their effectiveness presents many challenges. This makes it necessary to prune the task tree. Evaluating the considered visualization techniques based on human integrated tasks-at-hand is the goal of this evaluation, since only in that way, the weaknesses and strengths of each technique for our data sets can be identified. The tasks were generated and selected based on (a) discussions with people from the visual analytics and stock market area and (b) whether or not stock market analyzes are interesting and would naturally occur in the analysis of stock market data.

3. CHOOSING TECHNIQUES

Since the evaluation of every single visualization is not the aim of this study, we attempt to choose in this section one technique of each category identified in [4] for the final evaluation. Only those visualization techniques that fulfill the criteria below will be considered for the final evaluation. These three criteria are their suitability for the representation of large data sets, their suitability for the comparison of multiple data sets, and the advantage they take of the available space on the computer screen.

3.1 Classical techniques

Classical Techniques used for the visualization of time series data, in our case, stock market data, vary from bar charts, line graphs, point charts, pie charts to tables. These techniques are easy to understand and work well when the development or the tendencies are to be visualized over the time. *Line or Stock Charts* will be chosen for the evaluation (see figure 1(a)).

3.2 Geometrically-transformed visualizations


Geometrically-transformed visualizations, like parallel coordinates or pyramid representations, do quite clearly show if there are clusters in the data, what helps find the desired patterns. Parallel coordinates [2] are apparently very useful in exploring large data sets but it is still difficult to achieve a clear and helpful representation, if many dimensions and many data items are to be visualized. The visualization turns out completely different, too, if the mapping of the data to the different axes changes, that is, the right ordering of the data axes plays an important role and is decisive for finding the right clusters. Nevertheless, in this evaluation we will analyze this technique, i.e. *Parallel Coordinates* (see figure 1(b)).

3.3 Iconic displays

Iconic displays map attribute values of each data record to the features of the icons. An interesting technique belonging to this category is the so called heatmap. Heatmap [7] is a two-dimensional visualization and monitoring software, that was specially developed for the finance world. It enables the interactive monitoring of several variables, such as finance indicators and their relationships. However, it is not appropriate for very large data sets. Chernoff Faces [1] have also been used for the representation of stock prices, however, this technique is not appropriate when many stocks are visualized on the screen. The color icon [6] display enables the representation of large data sets, the comparison of the displayed data sets and takes good advantage of the computer screen. Therefore, the technique *Color-Icon Display* is chosen for the evaluation (see figure 1(c)).

3.4 Dense pixel displays

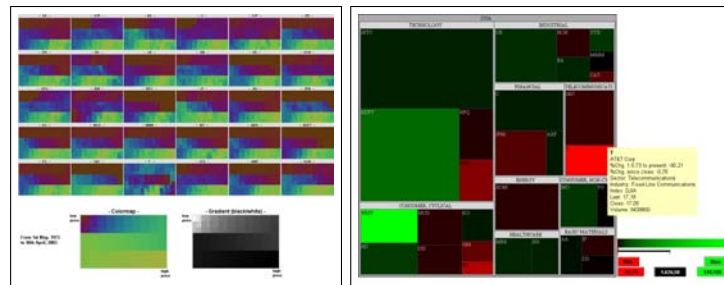
Dense pixel displays map each dimension value to a colored pixel and group the pixels belonging to each dimension in adjacent areas. In general, these displays use one pixel per data value, allowing the visualization of the largest amount of data possible on current displays (up to about 1.000.000 data values), since they make full use of the available place on the screen. These pixels have to be arranged in an appropriate way, to provide detailed information on local correlations, dependencies, and hot spots. We will choose from this category the technique *Recursive Pattern* [5] (see fig-



(a) Standard 2D/3D Displays

(b) Geometrically-transformed visualizations

(c) Iconic displays



(d) Dense pixel displays

(e) Stacked display

Figure 1: Chosen visualization techniques for the user study showing the Dow Jones Industrial Average over 30 years (all figures have the same screen resolution)

ure 1(d)), since this technique allows a better understandable representation of our time attribute in its different taxonomies, i.e. in seconds, hours, days and so on.

3.5 Stacked display techniques

Stacked display techniques are tailored to present data partitioned in a hierarchical fashion. In case of multidimensional data, the data dimensions to be used for partitioning the data and building the hierarchy have to be selected appropriately. Treemaps [3] belong to this category and consist of a two-dimensional space-filling approach in which each node is a rectangle whose area is proportional to some attribute such as node size. The usefulness of the resulting visualization largely depends on the data distribution of the outer nodes and, therefore, the dimensions which are used for defining the outer node system have to be selected carefully. In addition, the readability, aspect, and change ratios of the final visualization strongly depend on the algorithm chosen for the TreeMap.

4. EMPIRICAL EVALUATION


80 participants volunteered for the empirical evaluation. The users were mostly recruited within the university community, undergraduate and graduate students from the *Faculty of Computer and Information Science* at the University of Konstanz and external professionals from the visual analytics community and stock market area. Some participants were invited by email to take part in this evaluation. Web address and password were only provided to the users who accepted to take part in this study.

4.1 Experimental Design


All techniques are tested with both data sets, that is, with historical DJIA (Dow Jones Industrial Average) and intra day DAX (Deutscher AktienIndex) prices, so that the suitability of the different techniques to each data type can be also identified. The online evaluation has been divided into three sections, which are a *Background Questionnaire*, that will help us get to know each participant, the *Evaluation Tasks*, consisting of questions to be answered after observing the visualization, and a *Feedback Questionnaire*, in which the participant's opinion regarding the visualization and the evaluation is asked. The user can just proceed to the next task, when the previous one has already been completed. The number of tasks to perform is also known by the user. In that way, we try to avoid that the user feels stress or pressed for time. No knowledge about or experience with the visualization is needed in order to successfully solve the tasks.

4.2 Procedure

All our experiments have the same basic form. The participants are randomly assigned to one visualization. They were informed of the purpose of the evaluation and all test sections the evaluation consisted of, of their approximate length, and that they would have access to a brief introduction of the visualization technique, if required. In this way, each participant would know what they should expect during the test. For the evaluation of each technique, users were asked to fill out a questionnaire. The questions were grouped into nine categories and were aimed at testing the



(a) Historical Data – DJIA (from May 1st, 1973 to April 30th, 2003)



(b) Real-time Data – DAX (September 11th, 2001)

Figure 2: Ratio of Success to Failure vs. Time to Completion – The ideal visualization technique needs no time to complete all tasks-at-hand without any error. In the figure for our Historical Data the Recursive Pattern technique seems to be the best one and for the used Real-time Data the Charts are the favorite ones.

corresponding visualization against the tasks-at-hand. The same questionnaire with the same order of the same tasks was provided for the evaluation of each visualization. After all tasks were performed, participants were required to fill out a Feedback Questionnaire, so that they could inform us of the problems or difficulties.

5. EXPERIMENTAL RESULTS

5.1 Time to Completion

The time needed for the performance of the tasks was measured for every user and for each visualization technique.

5.1.1 Historical Data Visualizations

The results of the analysis of time to completion with respect to the visualization method are shown in figure 2. First, the significant differences between the parallel coordinates to all other visualization methods is not significant with respect to time to completion. The ANOVA shows us $p > 0.06$. Second, there are significant differences among all remaining visualization methods with respect to time to completion. ANOVA shows $p < 0.02$ for this comparison. ANOVA indicated that users spent much more time with the *Charts* than with the other techniques, followed by *Recursive Pattern*.

5.1.2 Real-Time Data Visualizations

The results of the analysis of time to completion with respect to the visualization method are shown in figure 2. There are no significant differences among the visualization methods with respect to time to completion in case of real

time data. ANOVA shows $p > 0.1$ for this comparison. ANOVA indicated that users spent nearly the same time with all techniques, that is, it is not possible to define a ranking.

5.2 Correctness of Answers

Ratio of Success to Failure (Error Rate) er depends on the amount of tasks completed correctly and the amount of tasks performed incorrectly.

5.2.1 Historical Data Visualizations

Figure 2 shows the results of the analysis of error rates with respect to the visualization methods. There are only significant differences among *Parallel Coordinates* and all remaining visualization methods with respect to the error rate. The ANOVA shows us $p < 0.001$. ANOVA indicated that users created much more incorrect answers with the *Parallel Coordinates* than with other techniques.

5.2.2 Real-Time Data Visualizations

The results of the analysis of error rates with respect to the visualization method are shown in figure 2. First, there are significant differences among *Chart* and all other visualization methods with respect to the error rate. The ANOVA shows us $p < 0.02$. Second, there are significant differences among *ColorIcon* and the *TreeMap*, *Recursive Pattern* and *Parallel Coordinates* techniques with respect to the error rate. The ANOVA shows us $p < 0.03$. ANOVA indicated that *Chart* performs best in case of real-time data streams, followed by *Color Icons* and *TreeMap*, *Recursive Pattern* and *Parallel Coordinates* with nearly the same mean error

		Visualization Techniques				
Tasks Evaluation Results (HD)		Charts	Parallel Coordinates	Color-Icon Display	Recursive Pattern	Treemap
T a s k s	Identify	+	-	--	+	--
	Locate	+	--	--	--	--
	Distinguish	✓	--	--	✓	✓
	Categorize	+	-	++	✓	++
	Cluster	✓	--	--	++	-
	Rank	++	++	++	++	++
	Compare	+	--	--	+	-
	Associate	++	✓	✓	++	--
	Correlate	✓	--	--	-	--

Legend: -- = Poor; - = Needs Improvement; ✓ = Average; + = Good; ++ = Excellent

Figure 3: Tasks Evaluation Results – Historical Stock Data – This table summarizes the results obtained by all five visualizations. The gray fields refer to the best performing techniques in the corresponding task.

rate.

5.3 Task Performance

For the categorization of the results in the five classes represented, a method was employed that just takes into account the normalized number of subjects (NNoS) who completed the corresponding tasks in estimated mean time correctly. Following this method, “Poor” was considered the NNoS between 0.0 and 0.1, “Needs Improvement” was called the NNoS between 0.11 and 0.3, “Average” refers to the NNoS between 0.31 and 0.5, “Good” was the category named for the NNoS between 0.51 and 0.7, and the title “Excellent” was awarded to the NNoS between 0.71 and 1.0.

5.3.1 Historical Data Visualizations

According to the results obtained in this section, *Recursive Pattern* is the technique which enabled the most correct answers in the Identify task. The second best technique for this task seems to be *Charts*, which did achieve very similar results. As for the Locate task, *Charts* is the only technique, which did get good results, while the other techniques did not enable any answer. The techniques *Parallel Coordinates*, *Recursive Pattern* and *Treemap* did allow the location of the stock being asked for, which was the first part of this task. However, none of these three techniques enabled the user to find the needed concrete price information of the mentioned stock. *Recursive Pattern* is the technique, which led to the best results in the Distinguish task, followed by the *Charts* and the *Treemap*. The *Color-Icon Display* and the *Treemap* are the techniques that delivered the best results for the Categorize task. These two techniques are also the only ones that grouped and represented the different stocks according to the industry sector they belong to. *Recursive Pattern* is the only technique that led to excellent results for the task referring to finding stocks with similar behavior (cluster task). The second best technique for this task

was *Charts*, however, its performance was just average. The order of the different data sets (Rank task) was mostly correctly identified at all five visualizations. *Recursive Pattern* and *Charts* did also enable a good comparison of stocks in the Compare task, while the other three techniques mostly delivered poor results. As for the Association task, *Charts* is the only technique that led to excellent results, followed by the *Recursive Pattern* which also performed well. Finally, in the Correlate task is the *Chart*, that performed the best, again followed by the *Recursive Pattern*. The other three techniques performed very badly, since, according to the participants, they did not enable this kind of task at all.

5.3.2 Real-Time Data Visualizations

According to the results obtained in the Identify task, *Charts*, *Recursive Pattern* and *Treemap* are the techniques which led to the best results, although the achieved results are not that good. The visual dynamics of the *Treemap* may have prevented the users from focusing on the required stocks. As for the Locate task, *Charts* and *Parallel Coordinates* led to excellent results. The results obtained with the other three techniques are, however, very poor. In *Recursive Pattern*, the requested information was probably located on the display (first part of the task), but the concrete price information was not found. *Charts* and *Recursive Pattern* are the tasks, which best enabled the completion of the Distinguish task, although the results were just average. The other three techniques were considered rather inappropriate for this task. *Color-Icon Display*, *Treemap* and *Charts* enabled an excellent categorization (Categorize task). The other two techniques did not clearly show the categories to which the different data sets belonged. The results obtained in the Cluster task are not good in any case, but *Charts*, *Parallel Coordinates* and *Recursive Pattern* led to the best ones. The order of the different data sets (Rank task) was

		Visualization Techniques				
Tasks Evaluation Results (RD)		Charts	Parallel Coordinates	Color-Icon Display	Recursive Pattern	Treemap
T a s k s	Identify	-	--	--	-	✓
	Locate	++	++	--	--	--
	Distinguish	✓	--	--	✓	-
	Categorize	++	-	++	+	++
	Cluster	-	-	--	-	-
	Rank	++	++	++	++	+
	Compare	+	✓	--	-	-
	Associate	++	+	+	++	+
	Correlate	+	-	--	+	✓

Legend: -- = Poor; - = Needs Improvement; ✓ = Average; + = Good; ++ = Excellent

Figure 4: Tasks Evaluation Results – Real-Time Stock Data – This table summarizes the results obtained by all five visualizations. The gray fields refer to the best performing techniques in the corresponding task.

mostly correctly identified. *Treemap* is the only technique which led to some misunderstandings in this task, probably due to the dynamics of the visualization. *Charts* did also enable a good comparison of stocks in the Compare task, while the other four techniques delivered partly average and partly poor results. As for the Association task, the results obtained are generally not bad. *Charts* and *Recursive Pattern* led to excellent results, while the results delivered by the other three techniques were good and average. Finally, *Chart* and *Recursive Pattern* enabled a good Correlate task, whereas the other three techniques performed badly.

6. STOCK MARKET ANALYZER

It is important to present the user a *general picture or overview* of the entire data set, in order for the user to understand the whole scope of the data instead of being forced to use just a narrow view of it. Such a representation should allow the application of interactive *zooming and filter* methods, so that interesting subsets can be identified and the focus on one or more of them is possible. Besides, offering *details of the data*, the user can access to, if desired, in order to analyze the found patterns, should be also enabled. Such presentation method must also provide evidence to support such visual exploration process, called *Information Seeking Mantra* [8], so that the user's confidence in the visualization tool is complete.

6.1 Historical Stock Market Data

The visualization techniques, which have been selected to be employed for the representation of historical stock market data, are, as already discussed in section 4, *Charts* and *Recursive Pattern*. These two techniques have proved to be suitable for the representation of large data sets as well as of the time dimension of the displayed data.

As for the information that should be available, the end-user should have access to actual stock price data in order

to be able to determine the stock price evolution in the past and in the present, in addition to dividends analysis of the stocks, and minimum and maximum prices reached during specified time periods. *Technical analysis* indicators, such as Moving Average Convergence Divergence, Relative Strength Index or resistance and support levels, as well as information about the overall economy and individual companies, and industry and financial conditions to enable the *fundamental analysis* among other things should be also available, since they have been identified by our users as valuable information to have access to. Stock volume information should be also provided, since it has been also considered a key way to measure stock supply and demand, and an indicator of a new price trend. Figure 5 shows the graphical layout of the resulting design for the historical stock market data.

The *Recursive Pattern* visualization provides an overview of the available data. These data are categorized according to the industry sector the stocks belong to. For the layout displayed in figure 5, we part from the idea that from that overview, just two industry sectors, i.e. the telecommunications and technology sectors, have been selected for a more detailed view. This window is linked with the *Charts* representation, so that each selection in the overview can be also highlighted on the *Charts* representation and vice versa, like figure 5 shows. This way, the same data can be observed from a different point of view. The selection of additional stocks, of all stocks of the same sector, or of a stock index can also lead to their representation on the same *Chart* to compare their performance against each other better. The number of stocks that can be displayed on the same *Chart* is limited to four, however, additional charts can be also launched.

Moreover, these two displays are also linked to the two top right windows offering stocks information and news, respectively. That is, when a time period of a stock is selected on any of the displays, the corresponding stock information

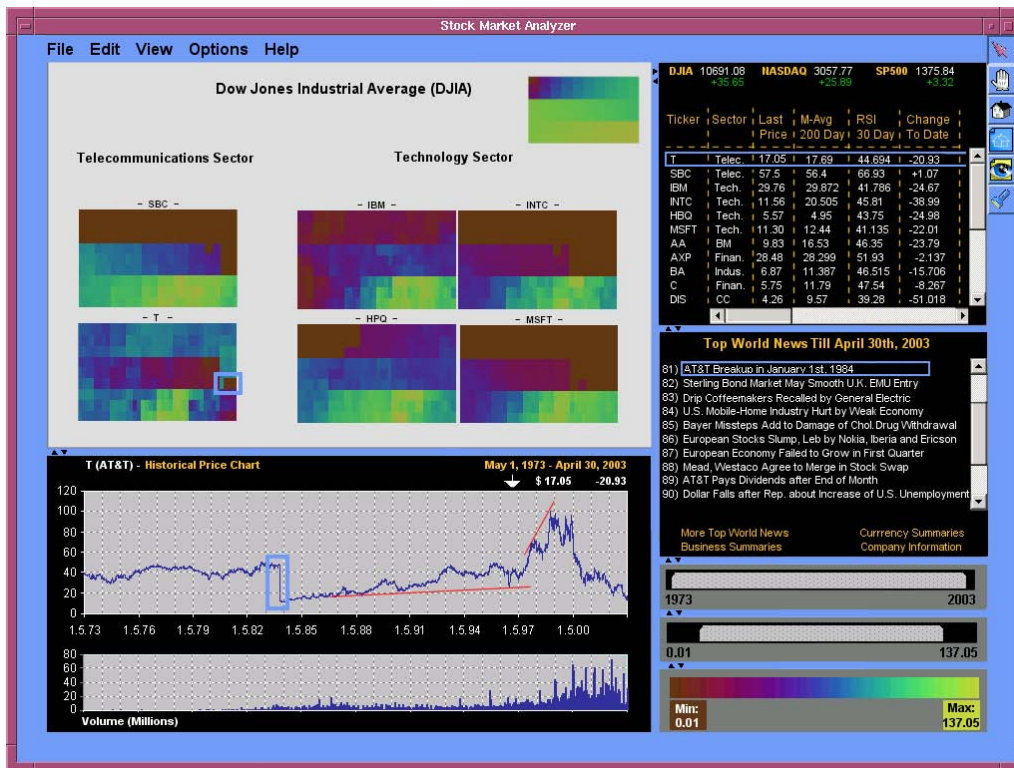


Figure 5: *Stock Market Analyzer – Historical Stock Market Data:* The first main window of the *Stock Market Analyzer* is displayed. When selecting a determined time period of the stock T on the *Recursive Pattern*, the corresponding *Chart* is displayed and the same period is blue marked on that *Chart*, and the corresponding data are highlighted in the windows for stock information and news.

is highlighted on the first window and the news published during that time period affecting the corresponding stock are marked on the second one, enabling an exploration of the corresponding data and knowledge about the causes of the stock reaction during that time. Selecting a stock on the first window or a piece of news on the second one leads to more detailed information about the stock performance and the news, respectively. The two slides underneath the news window enable the selection of the time period that should be displayed on all windows and the price range of the stocks that should be part of the display, offering in that way a filter method to facilitate the visual data exploration. Finally, the colormap at the bottom right corner specifies again the price range of the stocks displayed and the color mapping of this price range employed in the *Recursive Pattern* window.

In addition, selecting a stock on any of the windows also leads to the display of the current intra-day stock information on the second main window, which is described below.

6.2 Real-Time Stock Market Data

For the representation of intra-day stock market data, *Charts* and *Treemaps* have been selected, since the combination of both techniques enables the performance of all above mentioned tasks. The same mechanisms and instruments already described in the previous section are also applied to these techniques, to provide the end-user with very valuable information for the decision-making process. Figure 6 shows the resulting design of the intra-day data main

window, which is linked to the historical data main window, as already commented in the previous section. The *Treemap* visualization provides an overview of the available data. These data are categorized according to the industry sector and to the index the stocks belong to. This window is linked to the *Charts* representation, so that each selection on this overview is also highlighted on the *Charts* view, like figure 6 shows, in order to enable the observation of the same data with different techniques at the same time. The fact that, all stocks are displayed alphabetically ordered within each of their corresponding industry sectors facilitates the user interpretation of the showed data, since the minimal impact on data changing is ensured.

The actual stock data and a brief current news summary are placed on the top of the corresponding window at real-time, however, when selecting a specified time period during the day or a determined stock on any of the displays, the corresponding stock information is highlighted on the first window and the news published during that time period affecting the corresponding stock are marked on the second one, enabling the visual exploration of the corresponding data and knowledge about the causes of the stock reaction during that time of the day. The selection of a stock on the first window or of news on the second one leads to more detailed information about the stock performance and a broad summary of the corresponding news, respectively. The two slides underneath the news window enable the selection of the time during the day that should be displayed on all win-

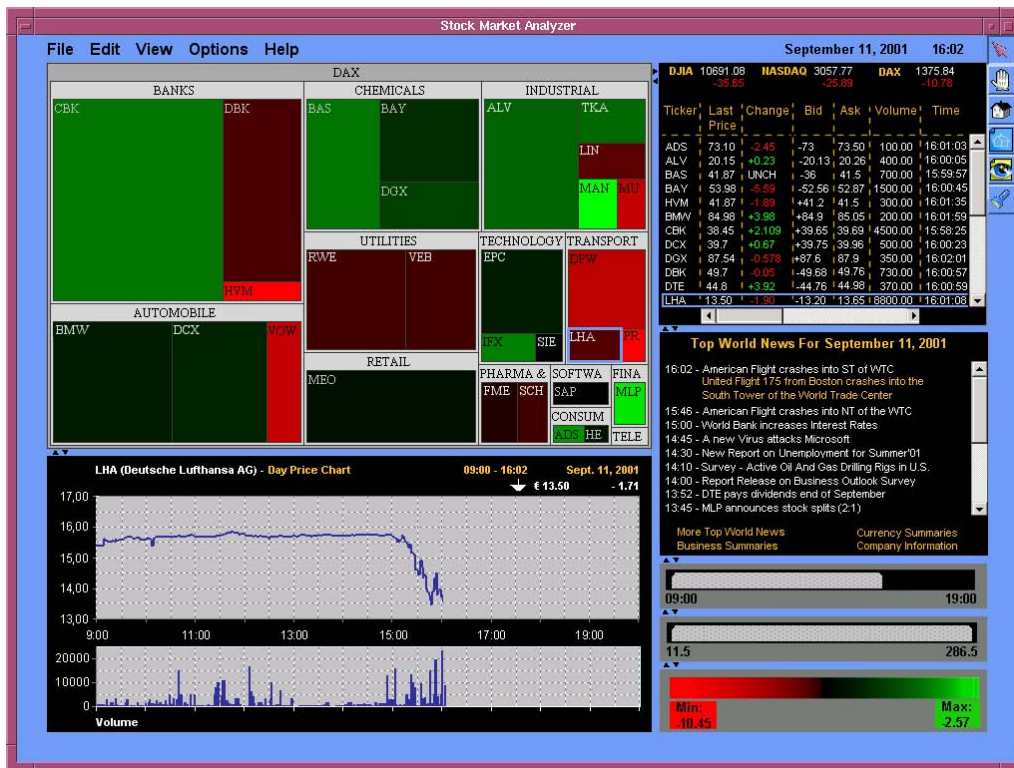


Figure 6: Stock Market Analyzer – Real-Time Stock Market Data: The second main window of the *Stock Market Analyzer* is displayed. When selecting the stock LHA on the *Treemap* overview, the LHA chart is displayed on the bottom left window and the corresponding information, if available, is highlighted in the other windows. The news window shows real-time news on the top.

dows and the price range of the stocks that should be part of the display (same features as in section 6.1).

Such an application would demand reliable, efficient, expedient database systems as the back-end, linked to a real-time or 10-15 minute delayed data feed, which could be obtained from Reuters or Bloomberg.

7. SUMMARY AND CONCLUSION

In this work, five different visualization techniques have been evaluated with regard to their suitability for two different data sets belonging to the above mentioned data types, as well as their support for our different tasks-at-hand. The result of our empirical study indicates that no visualization technique addresses all sections separately. According to the evaluation results the most appropriate techniques for historical data are *Charts* and *Recursive Pattern*. The empirical study confirms our observation that pixel techniques can lead to very good results in finding patterns and exploring large time series data sets in comparison to the other studied techniques (see figure 2(a)). As for real-time data, *Charts* and *Treemap*, after applying a different sorting algorithm, have been considered as the two most appropriate techniques for intra day data (see figure 2(b)).

8. REFERENCES

[1] CHERNOFF, H. The use of faces to represent points in k-dimensional space graphically. *Journal Amer. Statistical Association* 68 (1973), 361–368.

[2] INSELBERG, A., AND DIMSDALE, B. Parallel coordinates: a tool for visualizing multi-dimensional geometry. In *Proceedings of the 1st conference on Visualization '90* (1990), IEEE Computer Society Press, pp. 361–378.

[3] JOHNSON, B., AND SHNEIDERMAN, B. Treemaps: A space-filling approach to the visualization of hierarchical information. In *Proc. Visualization '91 Conf* (1991), pp. 284–291.

[4] KEIM, D. A. Information visualization and visual data mining. In *IEEE Transactions on Visualization and Computer Graphics* (January–March 2002), vol. 7.

[5] KEIM, D. A., KRIEGEL, H.-P., AND ANKERST, M. Recursive pattern: A technique for visualizing very large amounts of data. In *Proc. Visualization 95, Atlanta, GA* (1995), pp. 279–286.

[6] LEVKOWITZ, H. Color icons: Merging color and texture perception for integrated visualization of multiple parameters. In *Proc. Visualization 91, San Diego, CA* (1991), pp. 22–25.

[7] SHNEIDERMAN, B., AND WATTENBERG, M. Ordered treemap layouts. In *Proceedings of the IEEE Symposium on Information Visualization 2001 (INFOVIS'01)* (2001), IEEE Computer Society, p. 73.

[8] SHNEIDERMAN, B. The eyes have it: A task by data type taxonomy for information visualization. *Visual Languages* (1996).

[9] ZHOU, M., AND FEINER, S. Visual task characterization for automated visual discourse synthesis. In *Proceedings of the Conference on Human Factors in Computing Systems 1998, CHI'98* (1998), ACM Press, pp. 392–399.