

# Interacting with Stock Market Data in a Virtual Environment

Keith Nesbitt

Department of Computer Science and Software Engineering,  
University of Sydney, Sydney, NSW. 2006. Australia  
[knesbitt@cs.newcastle.edu.au](mailto:knesbitt@cs.newcastle.edu.au)

**Abstract.** Virtual Environment technology enables new styles of user interfaces that provide multi-sensory interactions. For example, interfaces can be designed which immerse the user in a 3D space and provide multi-sensory feedback. Many information spaces are multivariate, large and abstract in nature. It has been a goal of Virtual Environments to widen the human to computer bandwidth and so assist in the interpretation of these spaces by providing models that allow the user to interact 'naturally'. One goal for this interaction may be to uncover useful patterns within the data. This paper describes a Virtual Environment system called the "Workbench" and explains three models of stock market data that have been developed for this environment. The aim of this work is to provide models that allow analysts to explore for new trading patterns in the stock market data. Some early results of this work are discussed.

## 1 Introduction

Virtual Environment technology provides a new style of human-computer interface, the primary goal of which is to significantly increase the communication bandwidth between human and computer.

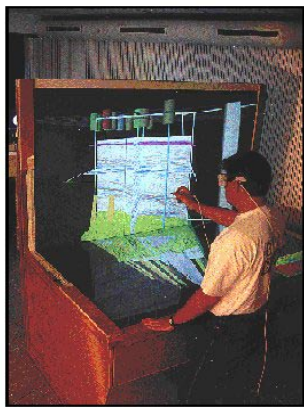
Virtual Environments attempt to create a natural way of interacting with computers using the human body and all its senses. In Virtual Environments users do not operate computer applications via an interface, rather people participate, perform tasks and experience activities within a computer generated world. The idea is to immerse a person in an environment that allows natural interaction and participation in order to perform tasks.

Many different types of Virtual Environment systems have been built and the technology has been applied in a wide range of fields [1,2]. This paper begins by describing a "Workbench" environment developed at the German National Research Center for Information Technology (GMD) [3, 4].

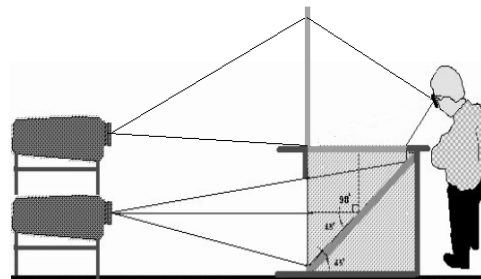
Like many modern industries, stock market analysis is characterized by an increase in the size of the data sets available. This data is large and multivariate. Analysts and

traders attempt to make profitable trades by determining relationships within the data. The domain of 'Technical Analysis' focuses on market activity to determine the balance of supply and demand for a financial instrument. This assists traders in making assessments on probabilities and risks about likely market directions. Estimating the direction and size of price movements from patterns in the data is useful for trading over various time frames. While many traditional techniques have developed to trade patterns within this data, finding new rules or patterns in stock market data may lead to new and more profitable trading systems. Section 3 of this paper provides an overview of the field of "Technical Analysis".

Virtual Environments like the "Workbench" offer new ways of presenting and exploring abstract data. With appropriate models it is possible to efficiently utilize our human capability for pattern recognition. Section 4 describes three new models that have been developed using traditional stock-market data. These models have been developed and demonstrated on the "Workbench" Virtual Environment. The paper concludes with a discussion of the early results, associated work and future directions.



**Fig. 1.** The Workbench at the German National Research Center for Information Technology. (Photo courtesy of GMD)



**Fig. 2.** The layout of an L-shaped Workbench showing how two projected images are combined to create a single displayed model.

## 2 The Virtual Workbench

The concept of the "Workbench" originated from research by Wolfgang Krüger at the German National Research Center for Information Technology (GMD) [3]. The Workbench enables a user to perceive and interact with a 3-D image that appears to float just above a table (Fig. 1). The computer generated image is projected onto a mirror beneath the workbench, where it is reflected upwards to an horizontal rear-projection screen which forms the table top surface (Fig. 2). The original design used only a single flat projection surface. However, this arrangement allows only models

that have minimal height to be displayed. This design was augmented with a second rear-projection screen at the back of the bench. This creates an L-shaped surface and allows models to be displayed in a more vertical orientation. The extra projector provides a further advantage as it also increases the brightness and resolution of the display.

The image (typically 1024 x 768 pixels) for both the horizontal and vertical surface are generated by a high-end Silicon Graphics Onyx workstation, equipped with Infinite Reality Graphics hardware. The workstation also receives information from an electromagnetic tracker unit, which provides position and orientation of the user's head and hands in the Workbench's virtual workspace. To perceive a 3-D stereo image, the user must wear liquid crystal shutter glasses synchronized with the workstation's graphic output. The workstation generates a separate image for each eye, and alternates display of each image in synchronization with the liquid crystal shutters.

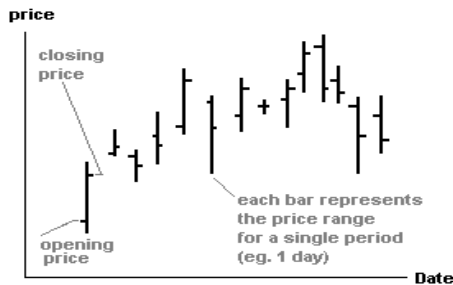
The distinguishing features of a workbench versus a traditional workstation based application are that:

1. The user perceives models in three dimensions.
2. The view of the data is controlled by the user's head position. As the user moves his/her viewpoint, the data is displayed as if seen from this position.
3. The user can interact directly with the virtual objects displayed above the tabletop. Models can be selected, rotated, translated and zoomed using virtual tools. These virtual tools are associated with a physical prop such as a pen with a selection button.

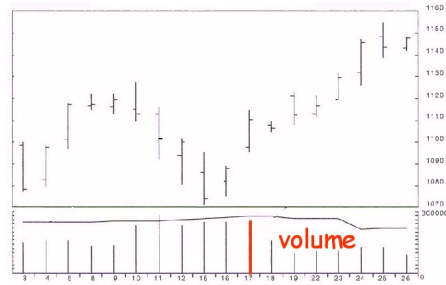
### **3 Technical Analysis**

'Technical analysis' is defined as "the study of behavior of market participants, as reflected in price, volume and open interest for a financial market, in order to identify stages in the development of price trends" [5]. Users of Technical Analysis seek to make profitable trades by studying market activity to determine the balance of supply and demand of a financial instrument. This field originated with Dow Theory and has developed to the extent that a number of different techniques now exist to assist with trading across different time periods.

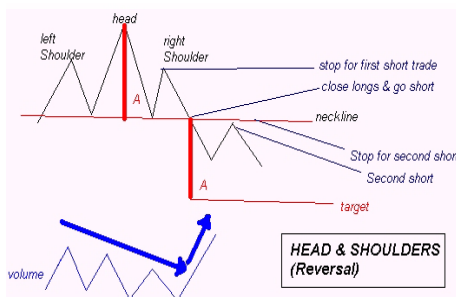
Technical Analysis is sometimes called 'charting' and, as the name suggests, often involves inspection of charts. The charts typically show price on the vertical axis and time on the horizontal. Price for a single time period is shown as a vertical line, or bar, which is drawn from the minimum price to the maximum price for the period. The period bar is augmented with ticks showing opening and closing price (Fig. 3). A time period represented may be a very short period, of the order of minutes, or longer periods such as a day, a week, months or years. The chosen period length reflects the trading strategy, longer periods if the emphasis is on long term trading or shorter periods for trading short-term market trends.



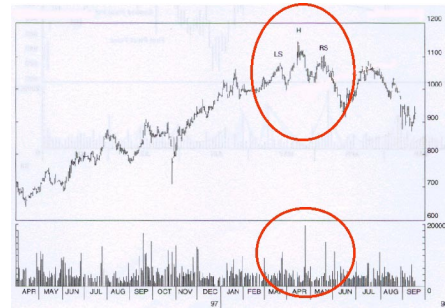
**Fig. 3.** A traditional daily bar chart.



**Fig. 4.** A bar chart with a volume histogram.



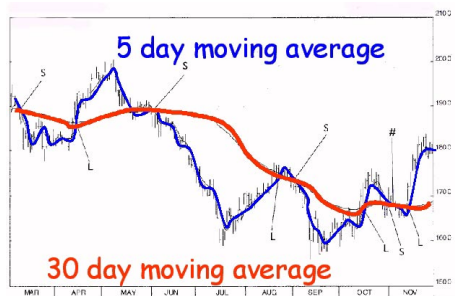
**Fig. 5a.** Trading strategies for the head and shoulder pattern.



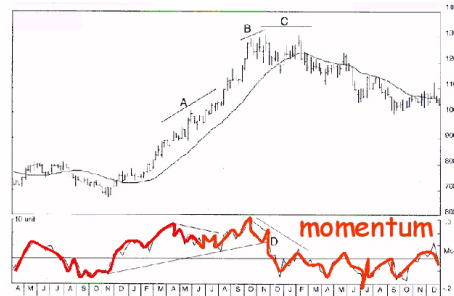
**Fig. 5b.** A head and shoulders pattern shown in the bar chart of price and volume data.

Charting techniques rely on the well-understood concept of a two-dimensional abstract space to present relationships. Price is used for the vertical axis and time is represented along the horizontal. Choosing time as one axis creates a time series of the data and each data bar can be compared to another in some temporal ordering. Other simple rules and relationships within this space are also understood. The height of the bar represents the variation of price for the period. The price from one period to the next can be compared by the placement of the bar along the vertical axis. Larger structural patterns can also be found such as the upward progression of consecutive bars which represents an up trend in prices. Many other patterns have been identified which characterize turning points of such trends and are useful for the trader to identify [12]. An example of a trend reversal pattern is shown in Fig. 5a and Fig. 5b.

These charts may also be augmented by a volume histogram that shows the volume of trades for each period at the base of the chart (Fig. 4). A number of derived indicators are also used to assist in analysis. These include a curve showing the moving average of closing price for consecutive periods. Moving averages help filter out short term fluctuations in price and provide information about longer-term trends (Fig. 6). Taking the difference in closing price between periods is the basis of another set of indicators. They provide an indication of the relative movement of price and are known as momentum indicators (Fig. 7). Like moving averages they can be derived for a range of different time steps.



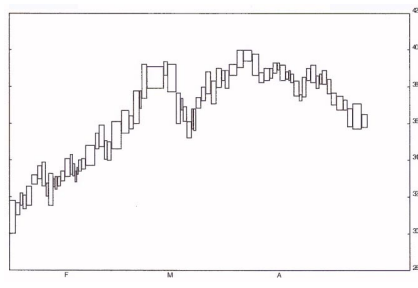
**Fig. 6.** Daily bar chart with two moving averages.



**Fig. 7.** Daily bar chart with momentum indicator



**Fig. 8.** Candlestick charting techniques use metaphors to describe patterns. Shown here is the "3 black crows".



**Fig. 9.** Equivolume charting where trading volume is represented by width of the price box.

While bar charts are the most frequently used charting techniques, there are also a number of specialised visualizations that have been successfully applied. Candlestick charts [6] are similar to bar charts but were developed independently in Japan. The bars or candles have a 'body' which is defined to be between they open and close price. The candle 'wicks' extend beyond the body to show maximum and minimum price. Black candles indicate price has fallen from open and close of the market. White candles occur when price has risen during the trading period. Metaphors are characteristically used to describe useful patterns such as "dark cloud cover" or "three black crows" (Fig. 8).

#### 4 Interactive Stock Market Models

Three new stock market models have been developed and they are now described. The first is a simple extension of traditional bar charts into three dimensions called the '3-D Bar Chart'. The second model, the 'Moving Average Surface' uses a series of moving averages created from price data to create a surface. The final model allows for real-time monitoring of 'bids' and 'asks' during market trading as is called the 'Depth of Market Landscape'.

#### **4.1 The 3-Dimensional Bar Chart**

The first investigations into 3D spaces enhanced a normal time series of price bars (bar chart) with volume in the third dimension (Fig. 10). Volume is often used to confirm price signals such as a trading range breakout. Specialized techniques such as 'Equivolume' [5] have been developed to include volume explicitly in a 2-D bar chart. However, this dramatically changes the way time is represented as each bar's width is no longer uniform but varies with trading volume for each period (Fig. 9).

It is more typical to chart volume as a histogram separately below the price bars (Fig. 4, Fig. 5b). While this allows price to be observed in relation to volume it requires moving the eyes back and forth between volume and price when trying to distinguish a correlation between the two variables of price and volume. It is a simple matter to extend price bars in the third dimension by mapping volume to depth. The user can then simply rotate the chart in the Virtual Environment and so explicitly compare trends in volume and price (Fig. 11).

#### **4.2 The Moving Average Surface**

More complicated spatial structures such as a surface of moving averages are also possible. In some technical analysis tasks it is useful to smooth out fluctuations which occur in price at each time step. 'Moving averages' [5] can be used to do this. Closing price is simply averaged over some number of time periods. Some trading systems rely on the intersection of different moving averages to signal the beginning and end of trends. For example a one, fifteen and thirty day moving average may be calculated and plotted. Signals are generated by where these moving average lines cross. The choice of how many time periods to include in the averages can vary, creating somewhat arbitrary signals. It is desirable to analyze a wide range of moving averages to choose appropriate signals for trading a particular instrument. A two-dimensional display of multiple moving averages soon becomes crowded if more than a few curves are plotted. Occlusion makes it hard to distinguish between curves (Fig. 12).

To overcome this problem a surface of moving averages was constructed with traditional price bars positioned at the center of the surface. This surface is constructed by joining together a number of strips. Each strip represents a different moving average curve. These moving average strips are joined to create a continuous surface. The surface is reflected about the central axis so that each edge represents a moving average of 30 days. As the surface moves towards the central bar chart the number of days in the moving average is reduced, from 30 to 29, to 28 and so on until a 1 day moving average is placed adjacent to the central bar chart (Fig. 13, 14, 15). By definition the closing price on these central price bars corresponds to a one-day moving average.

Signals generated from this view of the data still need to be clarified. However, trading signals are often only treated as indicators for action rather than absolute rules.

What this model provides is another way of looking at an indicator like moving average, allowing an analyst to consider a range of possible values for the parameters involved in its calculation. After viewing this model it was suggested that the area above or below curve might be useful to consider as a 'new' type of trading indicator.

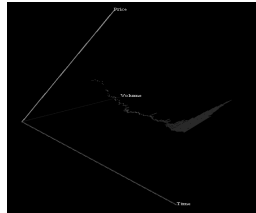
To assist in the analysis of this 'new' indicator a variation of this moving average landscape was created. Using the same series of moving averages from 1 to 30 days. Price bars are extended in one axis to create boxes that cover the width of the moving average surface (Fig. 16, 17, 18). Looking directly at one edge of the model shows a typical bar chart time series with the edge of the moving average surface seen as a line plotted through the closing price of each bar. By rotating the model the user can see price compared against the edge of the plane which represents the 30-day moving average.

While exploring this model it was found that in an up trend the price bars are predominantly above the plane and in down trends they are predominantly below. This has suggested a new way to evaluate trends by considering the extent that price bars lie above or below the moving average plane.

### **4.3 The Depth-of-Market Landscape**

Traders of financial instruments often have very different trading time frames. Long term traders for example will have a strategy of entering the market as it begins a primary up trend. This minimizes transaction costs as few trades are made and profits can result from both dividends and from the general upward trend of prices over time. The time frame may be months or years. Short-term traders on the other hand attempt to trade much shorter fluctuations in prices for profit. Here the time frame is of week or days. Both the 3D Bar Chart and the Moving Average Surface extend traditional pattern analysis techniques of charted stock market data. They can be useful for examining trends over different time frames by altering the period of the price bar. Some traders are interested in very short term trading opportunities that may result from fluctuations in market prices over five or 10 minutes.

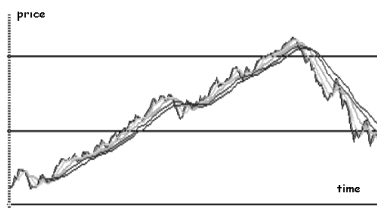
The 'Depth of Market landscape' was developed to explore the potential of trading opportunities that occur in very short time frames. In particular, the depth-of-market landscape allows the user to explore for new patterns in depth of market data that the short-term trader could exploit for profit.



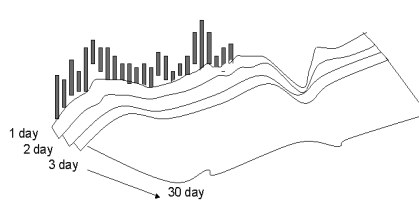
**Fig. 10.** A Bar Chart in 3D, showing price, time and volume.



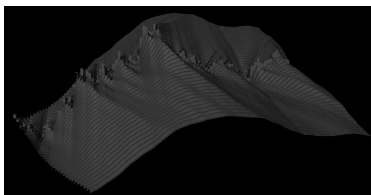
**Fig. 11.** Images showing a user rotating the 3D graph model on the 'Workbench' at GMD.



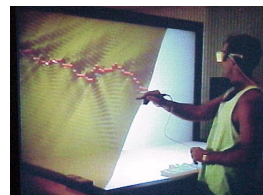
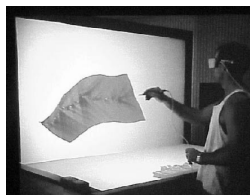
**Fig. 12.** Displaying multiple moving averages in 2D.



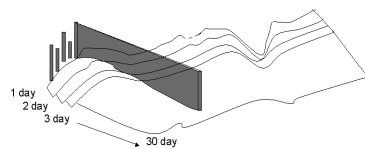
**Fig. 13.** Constructing a surface of multiple moving averages.



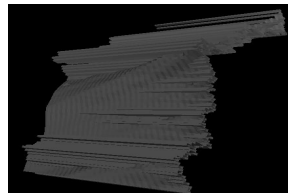
**Fig. 14.** The moving average surface.



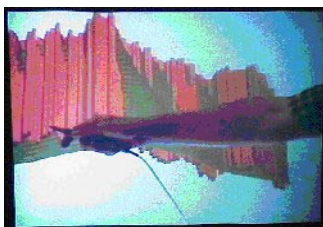
**Fig. 15.** Zooming in on the Workbench.



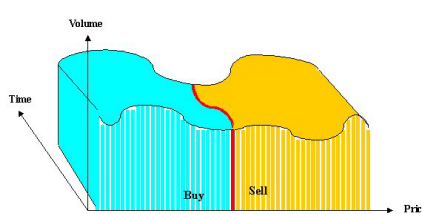
**Fig. 16.** Extending the price bars to intersect with the moving average plane.



**Fig. 17.** The model showing the moving average surface and intersecting price bars.



**Fig. 18.** Interacting with the moving average surface and extended price bars using the workbench & pen stylus.



**Fig. 19.** The Depth of Market landscape, which consists of a series of surface strips representing a volume and price histogram at different time steps.



The "*Depth of Market*" refers to the number of buyers and sellers currently trying to trade a particular financial instrument. A financial instrument may be something like a company share or future contract. The current selling price for an instrument can be considered as a balance between the price buyers will pay and the price sellers will accept. A buyer makes a "*bid*" to purchase a specified volume of an instrument. At the same time sellers try to sell a certain volume of an instrument for which they "*ask*" a particular price. The balance of buyers "*bids*" and sellers "*asks*" determine the state of the current market. Often there may be a difference in the buying and selling price and this difference is known as the "*spread*".

Current displays of depth of market data usually present the "*bids*" and "*asks*" in a simple table format that orders the list of bids and asks by price. The last selling price is also displayed indicating at what price the last trade was made. The landscape consists of a set of strips at each time step. Each strip has three components that represent a volume and price histogram for "*bids*", "*asks*" and also "*trades*". As time changes a new strip is added to represent the depth of market at that time frame (Fig. 19). This makes the depth-of-market landscape a 3-D model of a surface that evolves over time. It is expected that patterns may occur both in the static spatial structure and the evolution of the surface over time.

The landscape has the natural analogy of hills and valleys in the real world. The landscape evolves with time as the balance of buyers and sellers changes. These changes create waves that move on the surface of the landscape and can indicate changing trends in the short-term market. This is better understood if we consider some simple scenarios. Where there is a high volume of both buyers and sellers which is symmetric about the last sale price we expect price to remain fairly static as buyers and seller exchange trades. This may represent a price point about which distribution or accumulation is occurring. If, however, there were a valley between a peak of buyers and a peak of sellers this would indicate a market spread. Over time we could see this evolve into different situations. There could be no change in the market in which case we would expect few trades. If the peak of sellers moves towards the buyers we may expect prices to be driven down. Or alternatively the buyers may move their bids towards the available sellers and this may drive prices upward.

Explorations with this model are still continuing and early results are encouraging, though it requires a very dynamic market - that is with many frequent trades so the landscape can evolve at an 'interesting' rate. Another useful way to exploit this model may be for real-time monitoring of a market.

## **5 Discussion**

Like many new emerging computer technologies much hype and speculation has surrounded the value and application of Virtual Environments. Realistically, everyday use of these environments for applications such as technical analysis is not likely in

the short term. High cost, many useability issues and the lack of commercial software make it infeasible for rapid adoption of these environments. A shorter-term possibility is the use of such environments to investigate the discovery of new relationships in abstract data, such as that provided by the stock market. In such cases the potential reward may offset the risks against success.

Metaphors that providing totally new ways of exploring financial data may help reveal patterns that have not previously been understood. This may in turn create new and unique trading opportunities. Many systems have been developed which use algorithmic or heuristic rules for trading the market based on price trends. Once new patterns or rules are discovered the opportunity then exists to incorporate them into such automatic trading systems.

Further work needs to be done in developing and testing these models. It has been shown that a number of new opportunities for interpreting financial data can be provided by Virtual Environments. Early feedback from users confirms that these models are intuitive and easy to understand. A 'new' pattern - the area above and below the moving average surface is indicated for further investigation. The Depth of Market landscape requires further work to improve the look of the model and determine its usefulness. However, early indications from user feedback are encouraging.

## **6 Conclusion**

New opportunities for developing multi-sensory human-based tools have been made possible with new user-interface technologies. Virtual Environments immerse a person in a computer interface that allows natural interaction and participation within that environment to perform tasks. Interaction and perception in the real world is based on the use of multiple senses. We use our eyes and ears, the sense of touch and smell to perform activities. Virtual Environments attempt to mimic interactions in the real world and has seen the development of interfaces that support interaction for many of the human senses. While most activity has centered on three-dimensional visual models, there are also a growing number of applications where auditory and force displays are being used to help in data interpretation.

With multi-sensory interfaces we can potentially perceive and assimilate multivariate information more effectively. The hope is that mapping different attributes of the data to different senses, such as the visual, auditory and haptic (touch) domains will allow large data sets to be better understood. However, multi-sensory interpretation is a very complex field and involves understanding the physiological capabilities of each sense and the perceptual issues of individual and combined sensory interactions. Associated work is looking at extending these stock market models described here to provide multi-sensory feedback [7].

To assist in designing more intuitive multi-sensory interactions a classification of natural metaphors has been developed [8]. Associated with this classification are guidelines for integrating these metaphors in a way that best supports the human perceptual capability [9]. The harder question still remains, that is, to experimentally prove that this approach results in 'better' models for human-based data-mining.

## 7 Acknowledgements

The ideas in this work have resulted from close collaboration with Bernard Orenstein of Agents Incorporated in Sydney. Bernard has provided both support and invaluable expertise in the field of technical analysis. The integration of these models into the 'Workbench' was made possible with the assistance and support of Martin Göbel and Bernd Fröhlich from the Virtual Environment group. This group is part of the Institute for Media Communication (IMK) at the German National Research Center for Information Technology (GMD) [4].

## References

1. Durlach, N.I., Mavor, A.S. Virtual Reality. Scientific and Technological Challenges. National Academy Press, Washington, DC. 1996.
2. Stuart, R. The Design of Virtual Environments. McGraw-Hill, New York. 1996. ISBN 0-07-063299-5
3. Krüger, W., Bohn, C., Fröhlich, B., et al. The Responsive Workbench: A Virtual Environment. IEEE Computer, pp.42-48, July, 1995.
4. Internet web site: Institute for Media Communication, GMD, German National Research Center for Information Technology. <http://viswiz.gmd.de>
5. Technical Analysis : Course Notes from Securities Institute of Australia course in Technical Analysis. (E114), 1999. <http://www.securities.edu.au>
6. Specialised Techniques in Technical Analysis : Course Notes from Securities Institute of Australia course in technical Analysis (E171), 1999. <http://www.securities.edu.au>
7. Nesbitt, K. V. and Orenstein B.J. Multisensory Metaphors and Virtual Environments applied to Technical Analysis of Financial Markets. Proceedings of the Advanced Investment Technology, 1999. pp 195-205. ISBN: 0733100171.
8. Nesbitt, K. V. A Classification of Multi-sensory Metaphors for Understanding Abstract Data in a Virtual Environment. Proceedings of IV 2000, London. 2000.
9. Nesbitt, K. V. Designing Multi-sensory Models for Finding Patterns in Stock Market Data. Proceedings of International Conference on Multimodal Interfaces, Beijing. 2000.