Automated and Perceptual Data Mining of Stock Market Data

Keith V. Nesbitt
School of IT, Charles Sturt University
Bathurst, NSW, Australia
knesbitt@csu.edu.au

Abstract
Data mining large abstract data sets for useful patterns is an attractive proposition. It may, for example, reveal useful trading rules in stock market data. There are two basic types of data mining, Automated data mining tools use algorithms that allow the computer to search the data. Perceptual data mining tools present the data to the user's senses (vision, hearing, touch) in a way that the user can search for useful patterns. The two methods are not disjoint, as rules discovered with the user's perception can then be automated. This paper describes a case study, where a visual-auditory interface was used to uncover patterns in stock market data. Results from a formal evaluation are reported. The paper also includes a discussion on how to incorporate these results into an automated tool using an agent framework.

INTRODUCTION

A problem facing many areas of industry is the rapid increase in the amount of data and how to deal with it. However, these large amounts of data could also be considered a resource. The term *data mining* has been used to describe the diverse methods used to explore abstract data in the search for valuable and unexpected patterns [8].

The available tools for *data mining* can be considered in two broad categories, *Automated Intelligent Tools* and *Human Perceptual Tools* (see figure 1). *Automated Intelligent Tools* implement well-defined strategies for finding rules or patterns in data. These systems take advantage of a computer's capability to perform error-free, repetitive tasks and to process large amounts of data efficiently. *Human Perceptual Tools*, on the other hand, display the data to the user and allow the user to search for patterns. These systems take advantage of the human capability to perform subtle pattern matching tasks.

In fact the two types of data mining are closely related. New rules discovered by users of *Human Perceptual Tools* can be automated. These new patterns thus then form the basis of a new *Automated Intelligent Tool* (see figure 1). The primary motivation of this paper is to illustrate this close relationship with a real world case study.

The are a number of different approaches to developing *Automated Intelligent Tools* [17]. These include supervised methods, such as, Neural Networks and Linear Regression Models. They also include unsupervised methods, such as, K-means and self-organising maps. Because the idea of a Human Perceptual Tool is less familiar this paper reviews in more detail work in this area. In particular the review concentrates on the design of multi-sensory displays for virtual environments.

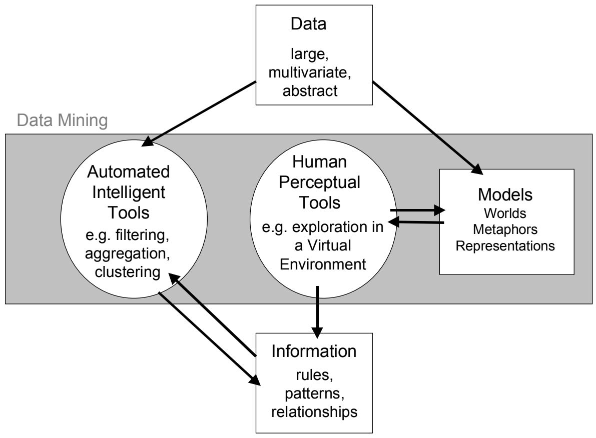

Figure 1. The two categories of data mining and the relationship between them.

The case study described in this paper is from the domain of stock market trading. The abstract data in this case uses *bid* and *ask* data from the Australian Stock Exchange. This data is also described as depth of market data. It is gathered in real time and captures offers made by potential buyers (bids) and sellers (asks) of a particular stock. The design of a auditory-visual display for predicting the direction of the stocks price is described. The results from a formal evaluation are detailed. The null hypothesis of the experiment was that "non-experts could not predict the direction of the stock price using this tool". Surprisingly the null hypothesis was proved false.

This motivates the possibility that useful patterns have been detected in the display. The reasoning used by each user to make decisions was also captured during the experiment. These "heuristics" are sometimes in conflict or competition. Even single users perceived conflicting sensory information when using the multi-sensory displays. One way to handle such complexity is to use an agent model. Such models have proved successful for modelling complex domains [7] and could provide a useful framework for implementing an automated tool. The preliminary design for such a tool is discussed. The paper concludes with an outline of further work.

PERCEPTUAL DATA MINING

It can be expected that *Human Perceptual Tools* are particularly useful where:

- Unpredictable exceptions may occur in the data.
- Heuristics are required to filter subtle variations.
- The target is unknown or cannot be precisely formalised by rules.
- The problem requires intuitive knowledge that is hard to formalise, such as, past experience.

During the 1990s, the accent for Human Perceptual Tools was on designing visual displays of data. This approach is called *Visual Data Mining* [17]. A number of example applications have been described and the field of Information Visualisation [4] has emerged. There are also some good examples using sound display [9] for finding patterns. Haptic (touch) displays are still relatively uncommon although they have been employed in novel ways to investigate force fields [3] and fluid flow fields [14].

To implement a Human Perceptual Tool typically requires a Virtual Environment. There have been a number of descriptions of Virtual Environments that define their properties [6]. These properties include immersion in a three-dimensional, synthesised world and multi-sensory interaction within that world [6]. One goal of Virtual Environments is to widen the bandwidth between human and computer. With multi-sensory interfaces the user can potentially perceive and assimilate multi-attributed information more effectively. By mapping different attributes of the data to different senses, such as the visual, auditory and haptic sense, it may be possible to better understand large data sets.

Displaying more data to the user makes perceptual data mining an enticing approach to searching large data sets. Designing a display using multiple displays is particularly difficult as sensory interactions can occur [15]. Previously the design space itself was not well enumerated. However, recent attempts have been made to better categorise the design space, gather guidelines and also provide a process for designing these displays [15].

TECHNICAL ANALYSIS

Stock market data contains many attributes, far more than traders can readily comprehend. Yet traders attempt to determine relationships between the data attributes that can lead to profitable trading of financial instruments. Predicting the market has some obvious rewards. However, the complex dynamics and interactions that drive the stock market are very difficult to model. Although a number of agent based systems have been developed to predict aspects of the market none of these have been proved successful in a all types of market conditions [1][10][11] [12].

Market traders make use of two types of complementary analysis. They are known as *technical analysis* and *fundamental analysis* [16].

Fundamental analysis studies the underlying factors that determine the price of a financial instrument. For example, factors such as, a company's profit, market sector, or potential growth can influence the share price. These factors can be considered against more global considerations such as the general economic trend. Fundamental analysis is the more traditional form of analysis used to trade the market.

Technical analysis is defined as "the study of behaviour 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" [16]. Unlike fundamental analysis, technical analysis ignores the underlying factors that determine price and makes the assumption that the price of a financial instrument already quantifies all these underlying factors. Technical analysis is based on patterns that can be found directly in the data. This case study is based on Technical Analysis.

CASE STUDY – DEPTH OF MARKET TRADING

This case study focuses on the search for short term trading rules based on *depth of market data*. Short-term players of the market, such as day traders, may wish to make minute by minute decisions from live feeds of stock data. The *depth of market* refers to the number of buyers and sellers currently trying to trade a financial instrument. A buyer may make a *bid* to purchase a specified volume of shares while at the same time a seller may *ask* a price for some specified volume. The balance of bids and asks determines the state of the current market. The difference between the highest bid and lowest ask is known as the *spread*.

Table 1. Traditional table used to display depth of market data.

Buyers		Trade 12.03	*Sellers*	
Volume	Price		Price	Volume
40,230	12.03	1	12.04	20,000
1,000	12.02	2	12.06	68
34,000	12.02	3	12.07	450
500	11.99	4	12.09	12,000
2300	11.98	5	12.10	6,000
10	11.97	6	12.11	10,200

The task of buying and selling shares to make a profit on short-term variations in market prices is called *discretionary trading*. The emphasis is on making a small profit many times during the period of trading. A trader may decide to sell when the volume of bids around the last trade price outweighs the volume of asks, or conversely sell if asks outweigh bids.

Traditionally the depth of market data is displayed in a table (see table 1), which updates every 30 seconds or so. The bids and asks are sorted so that highest bid and lowest ask are shown at the top. The difference between the highest bid and lowest ask indicates the spread. A wider spread usually indicates a lower likelihood that a trade will occur. The price of the last trade is shown for comparison with the current spread. Other more peripheral information includes the volume of stock in bid and ask quotes, and the context provided by lower bids and higher asks.

CASE STUDY – DESIGN

The visual part of the multi-sensory display was designed to provide extra context and historical information that is not easy to interpret from a conventional tabular display. The data is represented by a series of surfaces formed by the volume of bids, asks and trades at each price (see Figure 2). Price is on the X axis, volume on the Y axis and updates in time are appended on the Z axis. Bids are coloured yellow and asks are coloured green to separate them. The highest bid and lowest ask are next to each other close to the centre of the display while less important data spreads to the periphery. The trades are shown as a red river that also tracks the market spread (the difference between highest bid and lowest ask).

Note that as new data is used to update the display the visualisation takes on the form of an evolving 'landscape' that looks like a valley between two hills with a river flowing through it. It is proposed that this 'ecological' metaphor may help users interpret the visualisation from familiar natural properties such as 'steepness' of cliffs or 'height' of the hills.

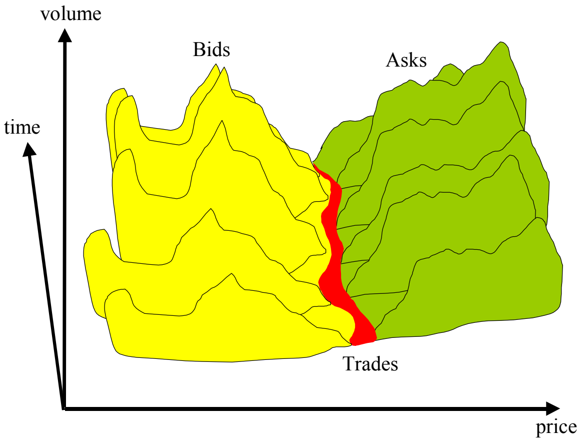

Figure 2. A conceptual model of the depth of market visualisation.

The sound display design method has two levels - a schema level and a perceptual level [2]. At the schema level a 'market place' metaphor is adopted. In this market place sellers shout the price of produce and buyers reply with offers or agree to trade. It is proposed that listeners using the display may interpret 'direction' of the next trade from this familiar experience. At the perceptual level we map information from the data attributes onto perceptually scaled auditory variables. This mapping considers issues of perceptual grouping and segregation in the auditory scene. While the final display is quite simple to interpret the full mapping is very complex and has been described elsewhere [15].

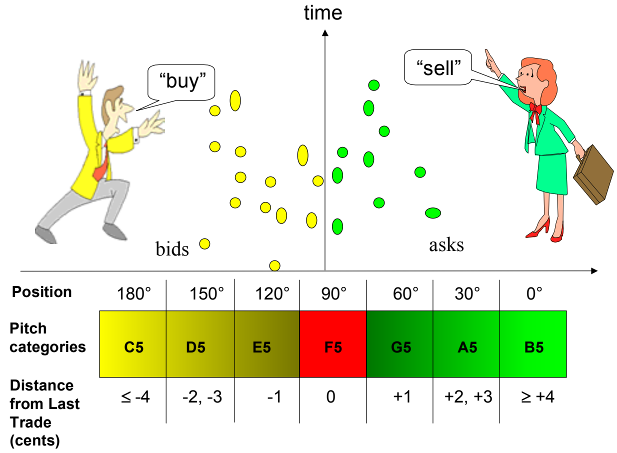

Figure 3. A conceptual model of the depth of market sound display.

Bids and asks are heard as they are made and come from the left or right as though from a crowd (see figure 3). If the bid or ask is lower than the last trade it is heard to the left, if it is the same then it comes from the centre and if it is higher it comes from the right. A flurry of bids to the right could indicate demand to buy at a higher price than the last trade and could indicate upward movement in trade price. A pattern of bids to the left mixed with asks to the right might indicate equilibrium in the market.

The resulting display should enable quick, confident and accurate answers to the global question "what is the direction of the next trade?" - either up or down. It should also enable answers to intermediate questions about relations between data elements such as "how wide is the current spread between bids and asks?" and "where is the current activity relative to the last trade?". At the local level it should allow answers to questions about individual elements such as "are there any bids?" or "what is the volume of the most recent ask?"

CASE STUDY – EVALUATION

The experiment was designed to answer the questions:

- Can people use the visual (V), auditory (A) and combined multi-sensory (M) displays to predict the direction of the next trade from depth of market data?
- What differences are there in performance with the visual, auditory and multi-sensory displays?
- Do people find consistent patterns in the data?
- How do people make decisions from these displays?

The multi-sensory depth of market display was implemented and evaluated on a BARCO stereo Projection Table, (see figure 4), in the CSIRO Virtual Environment Lab, Canberra. The null hypothesis for the evaluation was that subjects cannot predict the direction of the next trade from these displays. The alternative hypothesis that subjects can predict market direction depends critically on the technical trading hypothesis that all the information needed is contained in the data. Since it relies on this premise the data was recorded from real trading data for two shares over a day of trading on the Australian Stock Exchange. The data were divided into 6 subsets, three for training and three for evaluation. These subsets were randomly allocated to visual training, visual evaluation, auditory training, auditory evaluation, multi-sensory training and multi-sensory evaluation for each Subject. The Order of presentation of the different Modes of display (Visual, Auditory, Multi-sensory) was also randomised.

Figure 4. The Barco Projection Table at the CSIRO VE Lab being used during the evaluation.

Each Subject carried out the experiment one at a time with the researcher. At the start the Subject was given a written introduction to trading with depth of market stock data, and allowed to ask questions. They then carried out a training session followed by an evaluation session for each Mode. In the training session the Subjects were shown a display of historical data which was paused at 10 random points where they were told the direction of the next trade - up or down. In the evaluation the Subjects were asked to predict the direction of next trade at 10 random pause points as either up or down. The up and down movements were between 1 and 7 cents, with 80% of the decision points involving only 1 or 2 cent changes. After each evaluation the Subjects were asked for comments.

The Subjects comprised 13 males and 2 females between the ages of 20 and 42. Only one subject had any familiarity with depth of market data and none had traded on the stock market. For each Subject 10 predictions were recorded for each of the 3 Modes (V, A, M). At the end of the evaluation subjects were interviewed and asked how they made decisions using each display. Each Subject typically took 45 minutes to complete the training and evaluation.

CASE STUDY – RESULTS

The analysis was separated into 'total' predictions and predictions of price increase ('up') and predictions of a price decrease ('down'). Analysing totals correct out of 10 by regression analysis showed no significant effect for variation in Subject, Mode or Order. (see table 2). Next the analysis considered proportions of correct predictions for trades that went up in price separately from trades that went down using generalised linear models [5]. In the up direction there no significant effects, for Subject Mode or Order. However, in the down direction some Subjects performing significantly better than others (P=0.014), and there was also a significant variation with Mode (P=0.029). However the Order was still not significant.

Table 2. P values for significance (P<0.05) of main variables in experiment.

Direction	*Subject*	*Mode*	*Order*
all (up+down)	0.884	0.397	0.953
up	0.981	0.812	0.793
down	**0.014**	**0.029**	0.653

Table 3. Experimental analysis for prediction groups (all, up, down).

G	M	#	T	%	PC	s.e.	P
all	V	92	150	61.3	0.613	0.047	**0.000**
	A	105	150	70.0	0.070	0.047	**0.000**
	M	105	150	70.0	0.070	0.047	**0.000**
up	V	50	85	58.8	0.583	0.077	0.128
	A	46	78	59.0	0.595	0.079	0.141
	M	42	64	65.6	0.658	0.085	**0.017**
down	V	43	64	67.2	0.649	0.050	**0.008**
	A	59	72	81.9	0.831	0.036	**0.000**
	M	63	86	73.3	0.734	0.046	**0.000**

Table shows: prediction group(G), Mode (M), number of correct predictions (#), total number of predictions (T), percentage correct prediction (%), estimate of mean proportions correct (PC) and standard error (s.e.) for one binomial trial, and the exact two-tailed binomial of calculated probability (P).

Next we compared how well the Subjects were actually predicting the direction of the next trade, using the regression model to estimate mean proportions correct (PC) and standard error (s.e.) for one binomial trial, and using an exact two-tailed binomial to calculate probability (P) of the result occurring by chance. This analysis shows Subjects predict the direction of the next trade at levels significantly above chance in all three Modes, with A (70%), M (70%) and V (61.3%) (see table 3).

Subjects were best at predicting down trades from A (83.1%) followed by down trades from M (73.4%), up trades from M (65.8%), and down trades from V (64.9%). The prediction of up trades from V (58.8%) and A (59.0%) separately was not significantly above chance. The effectiveness of M in the up direction where the individual V and A are not so effective indicates that complementary information from both Modes is needed to predict up trades. Subjects predict down trades much better than up trades, at levels so far above chance that the overall analysis shows significant prediction for all 3 Modes despite the up results. The prediction of down trades was best with the A (81.9%), followed by M (73.3%), and then V (67.2% correct). These results suggest that A is the best source of information for predicting down trades, and that the combination of A with V reduces performance in the M display.

CASE STUDY – TOWARDS AN AGENT MODEL

The results from this evaluation are quite complex and further discussion of the issues raised about perceptual displays and predicting stock market direction are discussed elsewhere [13]. The results are encouraging but a great deal of caution is required. Many stock market prediction systems are highly dependent on the current market and the particular stock being predicted. Hence it would be desirable to automate the predicting process so a broader array of conditions could be analysed.

The first question that needs to be addressed is: Do people find consistent patterns in the data? The frequency for the proportion of correct responses at each decision point was analysed (see figure 5). While there were too few predictions made at each decision point to statistically validate the presence of patterns. However, indications are promising, with 10 points consistently predicted and 1 point consistently mis-predicted.

The second question that is addressed is: How do people make predictions? The comments from each Subject were recorded after the evaluation in each Mode. In the Visual display nine Subjects commented that they made decisions based on ‘size, height, slope and steepness of cliffs’, ‘how close the peaks were to the centre’, and ‘bending and trends in the river or valley’. Three said they could not understand how to make decisions. One commented that it was not clear whether the ‘forces shown by hills’ were pushing or pulling against each other. Overall three Subjects said they preferred the Visual display.

In the Auditory display the Subjects made decisions from ‘frequency of calls’, ‘closeness to the centre’, and ‘loudness’. Six said they found it easy to understand the Auditory display. One commented that the words ‘buy’ and ‘sell’ could be interpreted as commands rather than labels which could lead to a prediction in the opposite direction. One commented that information about the last trade price was missing from the Auditory display.

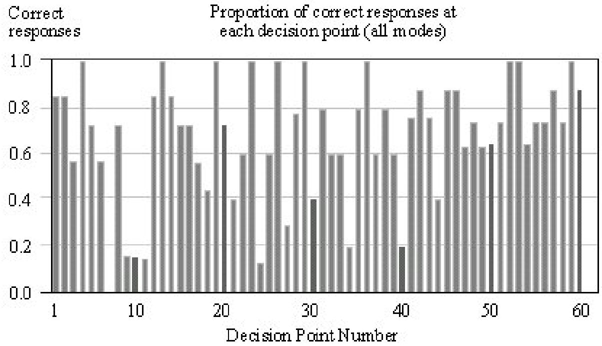

Figure 5. Proportion of correct predictions at the 60 decision points in the data.

In the Multi-sensory display Subjects commented that the combined display contained more information than the Visual or Auditory alone, and that the Visual display provided ‘context, history, past and general trends’, while the Auditory display provided ‘most recent trends’, ‘focus’, ‘eagerness’ and ‘presence’. Four mentioned that if the Visual display was not clear they used the Auditory display to make the decision. Four commented that Visual and Auditory displays were sometimes in conflict, and 3 of these relied on the Auditory while the other relied on the Visual in this situation. One found the conflicts in Multi-sensory display made it more ambiguous and preferred the Auditory alone. Another found the Auditory display distracting and preferred the Visual over the Multi-sensory display.

What is interesting about the comments recorded from Subjects is the different heuristics used to predict trade direction and also their level of confidence. In effect the Subject’s approach to interpreting the display provides a way of generating rules to automate the data mining process. The Subjects are in fact intelligent agents and we have tried to capture their beliefs, desires and intentions. What is interesting is that Subjects do not have an agreed approach and even find there own senses in disagreement. Sometimes the information from the Auditory and Visual displays conflict and at other times they collaborate. Some Subjects are strongly biased to the visual patterns while others favour the auditory information. This complex system indeed suggests that an Agent framework might be best suited to automate the data-mining.

A simple hierarchical model is proposed (see figure 6). At the lowest level there are two main families of agents. One family predicts on the basis of ‘Auditory’ rules. This in effect incorporates shorter-term fluctuations in the data. The second family uses ‘Visual’ rules that capture longer-term trends in data. Agents may collaborate or disagree within these families.

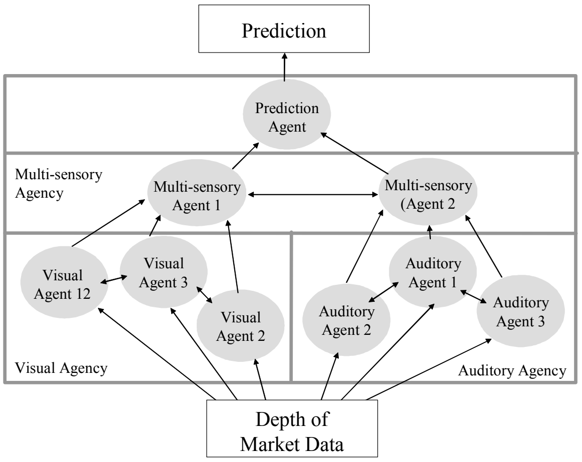

Figure 6. A proposed agent model for automating the prediction process.

At a higher level the results of the 'Auditory' and 'Visual' agents are processed by an agency of 'multi-sensory' agents. These agents incorporate longer-term and shorter-term information. In fact it is quite normal to incorporate longer term trading information into shorter-term trading rules [16]. For example, a bull market is a long-term up trend. In such a market it is expected that short-term trends are more likely to go up then down. It is likely that many agents with different policies will generate an array of predictions with different levels of confidence. A top-level agent, the 'Prediction' agent must resolve any conflicts and arrive at a final prediction.

FURTHER WORK

The Agent model needs to be implemented and tested. Work is under way to formalise the 'Visual' and 'Auditory' patterns that have been detected in the perceptual tool. The flexibility of an agent system makes the task of adding and removing 'rule' agents quite simple.

An automated tool will allow better testing of the results from the perceptual data-mining tool. However, the perceptual tool has already yielded promising results and more work is also ongoing to test the multi-sensory display in 'live' trading situations.

What this case study shows is that the two types of data-mining, perceptual and automated can in fact work together to yield cooperative benefits to the individual in search for patterns in abstract data.

ACKNOWLEDGMENTS

This work was carried out in collaboration with staff at the CSIRO Virtual Environment Laboratory in Canberra. I would like to thank Stephen Barrass and Warren Muller who provided major contributions to this work. Stephen's expertise is in the area of sound displays and he was instrumental in the design of the sound displayed used in this work. Warren provided invaluable help with experimental design and analysis.

REFERENCES

[1] Arthur, W. B. Holland, B. LeBaron, R., Tayler, P. Asset pricing under endogenous expectations in an artificial stock market. In W. B. Arthur, S. Durlauf, and D. Lane, editors, The Economy as an Evolving Complex System II, pages 15-44. Addison-Wesley, Reading, MA, (1997).

[2] Barrass, S. "Auditory Information Design". Ph.D. Thesis, Australian National University, Canberra, Australia (1997).

[3] Brooks, F. P., Ouh-Young, J. M. et al. "Project GROPE– Haptic Displays for Scientific Visualization." *Computer Graphics* 24(4): 177-185. (1990)

[4] Card, S. K., J. D. Mackinlay, et al., Eds. *Information Visualization. Readings in Information Visualization*. San Francisco, California, Morgan Kaufmann Publishers, Inc. (1999).

[5] Dobson, A. J. "An Introduction to Generalized Linear Models", Chapman and Hall, London, (1990).

[6] Durlach, N. I. and A. S. Mavor, Eds. *Virtual Reality: Scientific and Technological Challenges*. Washington, D.C., National Academy Press. (1995)

[7] Gallimore, R J., Jennings, et. al. "Cooperating Agents for 3D Scientific Data Interpretation"., IEEE Transactions on Systems, Man and Cybernetics, (1996).

[8] Groth, R. *Data mining: a hands-on approach for business professionals*. Upper Saddle River, N.J., Prentice Hall. (1998).

[9] Kramer, G. *An Introduction to Auditory Display. Auditory Display:Sonification, Audification and Auditory Interfaces*. G. Kramer, Addison-Wesley Publishing Company. (1994).

[10] Lettau, M. Explaining the facts with adaptive agents: The case of mutual fund flows. Journal of Economic Dynamics and Control, 21:1117-1148. (1997).

[11] T. Lux. The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions. *Journal of Economic Behavior and Organization*, 33:143-165. (1998).

[12] Lux, T. and Marchesi, M. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397:493-500, 1999.

[13] Nesbitt, K. and Barrass, S. (2002) "Evaluation of a Multimodal Sonification and Visualisation of Depth of Market Stock Data", Proceedings of the International Conference of Auditory Display. Kyoto, Japan, pp 233-238. (2002).

[14] Nesbitt, K. V., Gallimore, R. et al. "Using Force Feedback for Multi-sensory Display". *2nd Australasian User Interface Conference AUIC 2001*, Gold Coast, Queensland, Australia, IEEE Computer Society. (2001)

[15] Nesbitt, K. V., "Designing Multi-sensory Displays for Abstract Data", PhD Thesis, University of Sydney, Sydney, Australia (2003).

[16] Nicholson, C. *Technical Analysis*. Course Notes E114. Sydney, Australia, Securities Institute and Australian Technical Analysts Association. (1999).

[17] Soukup, T. *Visual data mining: techniques and tools for data visualization and mining*. New York, John Wiley & Sons. (2002).