Finding Trading Patterns in Stock Market Data

Keith V. Nesbitt
Charles Sturt University

Stephen Barrass
CSIRO

A combined visual and auditory design for helping traders detect patterns and predict stock market direction, tested experimentally, yields clues for future data mining research.

A problem facing many areas of industry is the rapid increase in data and how to deal with it efficiently. In many cases, these large amounts of data are a useful resource, and the problem then becomes one of how to extract meaning from that data. *Data mining* is the process of exploring abstract data in the search for valuable and unexpected patterns.

We can consider the available tools for data mining in two broad categories: automated intelligent tools and human perceptual tools, as Figure 1 (next page) shows. *Automated intelligent tools* implement well-defined strategies for finding rules or patterns in data. These tools exploit a computer's capability to perform error-free, repetitive tasks and to efficiently process large amounts of data without human intervention. With *human perceptual tools*, on the other hand, the focus is on keeping the human in the process by displaying data to users and letting them search for patterns. These tools take advantage of the human capability to perform subtle pattern matching tasks.

In the past, human perceptual tools often provided users with a 2D visual display to explore. However, new user interface technologies developed in virtual environments let researchers extend visual displays to 3D, and also incorporate sound and haptic (touch) feedback. Unfortunately, designing such multisensory displays is complex; although some researchers have formalized design frameworks for both visual[1] and auditory display.[2,3] Despite such frameworks, this field is still largely an embryonic research area and fraught with many difficulties.

This article describes our design and evaluation of a multisensory human perceptual tool for the real-world task domain of stock market trading. The tool is complementary in that it displays different information to different senses—our design incorporates both a 3D visual and a 2D sound display. The results of evaluating the tool in a formal experiment are complex.

The data mined in this case study is bid-and-ask data—also called depth-of-market data—from the Australian Stock Exchange. Stock market traders typically analyze this data in real time. It captures offers made by potential buyers (bids) and sellers (asks) of a particular stock. Our visual-auditory display is the bid-ask-landscape, which we developed over many iterations with the close collaboration of an expert in the stock market domain. From this domain's perspective, our project's principal goal was to develop a tool to help traders uncover new trading patterns in depth-of-market data.

In this article, we not only describe the design of the bid-ask-landscape but also report on a formal evaluation of this visual-auditory display. We tested nonexperts on their ability to use the tool to predict the future direction of stock prices. The experiment's null hypothesis was that nonexperts couldn't predict the direction of the stock price using this tool. Surprisingly, the null hypothesis was proved false, leading to the possibility that useful patterns were detected in the display.

Design framework

We developed the work described here using a multisensory design framework called the MS-Taxonomy.[4] Our basic motivation for designing a multisensory display for data mining is to widen the computer–human bandwidth. By computer–human bandwidth, we mean the amount of information—displayed by the computer—that users can perceive through their senses. We can achieve this widening by mapping different data attributes to the different senses. We base our measure of performance with the bid-ask-landscape tool on the characterization of multisensory displays that McGee et al. describe:[5]

- *Conflicting.* Contradictory information is displayed to each sense. Performance is worse in the multisensory display.
- *Redundant.* The same information is displayed to each sense. Performance with both the single-sensory and multisensory displays is the same, but there might be a reduction in workload or an increase in confidence.

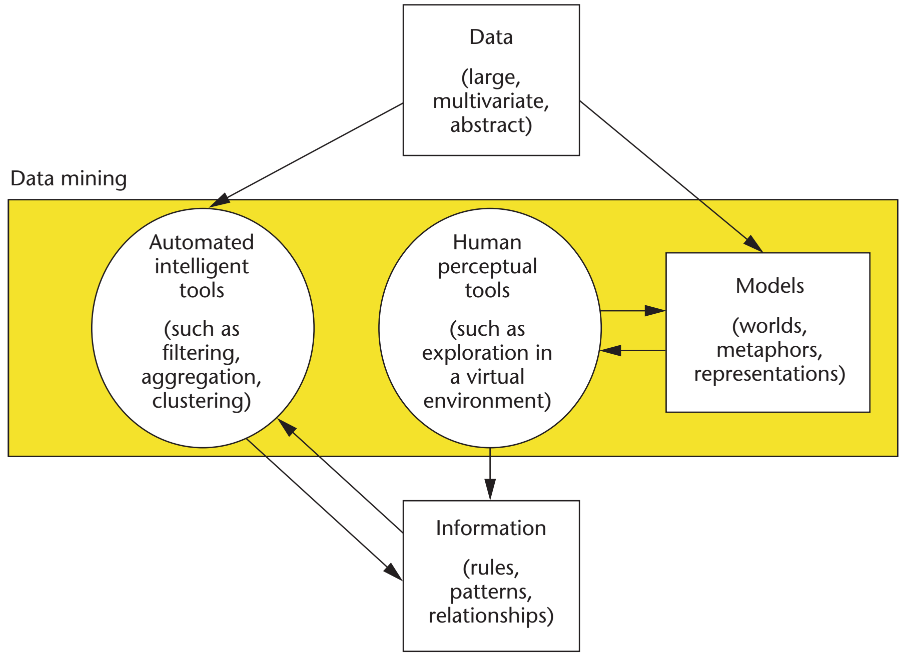

1 The two categories of data mining and the relationship between them.

- *Complementary.* Different information is displayed to each sense. Performance with the multisensory display is superior to the separate single-sensory displays.

Stock market

To better explain the motivation for our application and to provide a background to our experimental design, we offer the following overview.

Technical analysis

Stock market data contains many attributes, far more than traders can readily comprehend. Traders nonetheless attempt to determine relationships between the data attributes that can lead to profitable trading of financial instruments. Traders apply two types of complementary analysis to trade on the stock market: technical and fundamental.[6]

In *fundamental analysis,* traders study 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. Traders consider these factors against more global concerns such as the general economic trend. Traders have traditionally used fundamental analysis to trade the market.

Technical analysis is "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."[6] In technical analysis, traders ignore the underlying factors that determine price and assume that the price of a financial instrument already quantifies these underlying factors. Technical analysis relies on patterns found directly in the stock data. Because this work relies on the user's finding patterns directly in the data, it is based on technical analysis. Many traditional analysts don't support the assumptions made by technical analysts, because it ignores the underlying market factors on which stock prices are based and so is thought to be less reliable. Technical analysis lends itself to both visual and auditory displays of stock market data, as we explain.

Visual displays

Although some parts of the investment community dismiss technical analysis, it's still widespread and often used in conjunction with fundamental analysis to assist in timing the market entry and exit points for trading stocks most profitably. Technical analysis is also known as charting because it frequently involves visual analysis of 2D charts constructed over time to show variations in price, volume, or other derived indicators such as price momentum. For example, daily charts might show price bars that record the opening, closing, maximum, and minimum price for a period of trading. A simple graph of closing price over time can be used for longer term analysis and might often be augmented by a volume histogram.

The practice of visual charting is also quite old if we consider that the Japanese technique of candlestick charts is an early form of technical analysis and which originated sometime in the 17th century for use with futures trading (see Figure 2).

Traders who apply the charting approach make trading decisions from patterns that they observe in the data—for example, they might make a buy decision after observing two consecutive peaks in price (see Figure 2 and Figure 3). By recognizing new patterns through charting, traders could gain an edge over other traders.

Of course, analysts have long used computers to chart large amounts of stock market data and to produce (in real time) charts from live market data feeds. Researchers have developed novel visualizations to enable the perception and analysis of more information

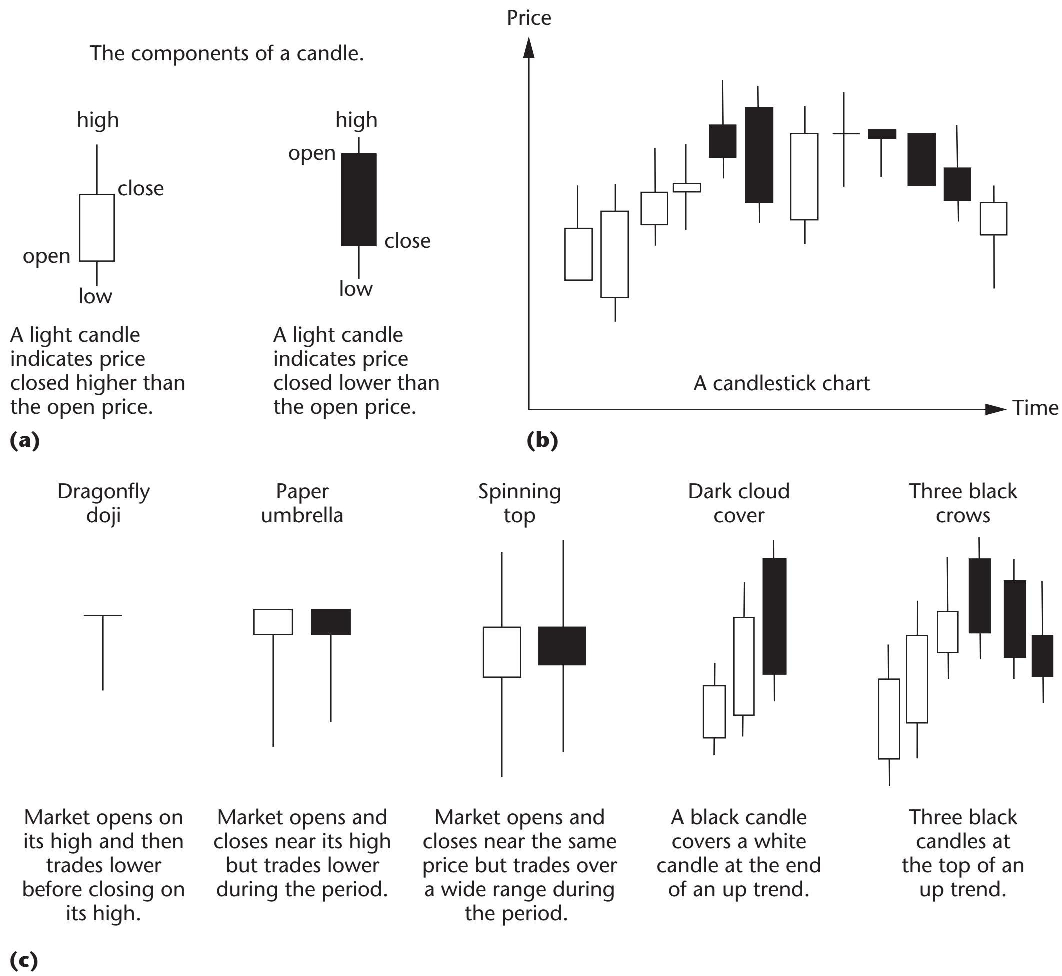

2 **(a) Components of a candlestick chart, (b) a candlestick chart, and (c) some typical patterns used in candlestick charting.**

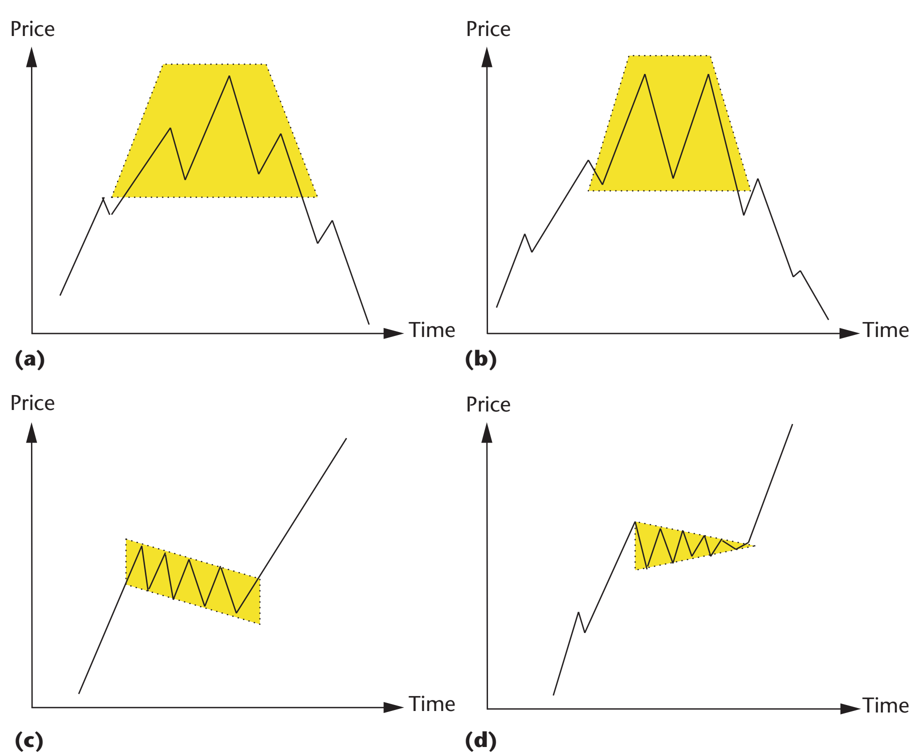

3 **A graph of closing price over time is often used by technical analysts. (a–b) The head-and-shoulders and double-top shaped shadings signal the reversal of an up trend; (c–d) the flag and pennant shaped shadings signal the continuation of an up trend.**

more quickly, such as treemaps for visualizing stock portfolios[7] and Wright's visualization of live data feeds of liquidity information.[8] Wright's work is in fact one of the few reported attempts to visualize depth-of-market data. One drawback is that Wright's display does not retain historical information and so might be less suitable for detecting temporal patterns. A simple way to incorporate time into a 3D model is to generate a group of bar charts showing bids and asks at discrete time steps, which creates a type of 3D order book. However, we generated the 3D visual model for our case study using a more complex mapping to support a focus-and-context display.[1] The most critical information for the trader is displayed in detail around the current trading price. This forms the visual focus of the display. Less relevant information provides context and is displayed to the user's peripheral vision.

Auditory displays

Researchers developing sound displays have also studied stock market data. These displays usually provide a mapping from the data to sound parameters such as pitch and volume. By monitoring the sound, the user monitors changes in the market. In 1990, Frysinger developed his audification technique for creating a sound display by directly playing back stock market price data as a sound waveform.[9] However, the sounds were difficult to interpret. Because stock market doesn't follow physical-acoustic laws, perhaps audification doesn't produce natural sounds that users can understand based on their everyday listening experience.

In contrast to audification, sonification indicates a mapping between the data attributes and the qualities of sound to display information.[3] Researchers have taken a number of approaches to developing sonifications. The simplest approaches describe mappings to sound attributes such as loudness (the perceived property that relates to the sound's amplitude), pitch (the perceived property that relates to the sound's frequency) and timbre (the perceived "quality" of a sound that distinguishes it from other sounds of the same pitch and volume).[3] The act of interpreting the properties of a sound source is often called *musical listening*, as opposed to William Gaver's *everyday listening*, where listeners interpret sounds in terms of the events that cause them.[3] For example, in the SonicFinder application, a monitoring sound—in the form of pouring water—was displayed when a user copied a computer file. The pitch of the pouring sound would gradually change to indicate the amount of copying that had been completed.[3]

Researchers have applied a range of sonification techniques to stock market data. In one experiment, for example, Brewster and Murray mapped stock market price data for a single share to pitch to investigate if sounds provide a viable alternative for trading shares on mobile computing devices with limited screen space.[10] Test subjects' performance was evaluated in a task where the aim was to maximize profit by buying and selling shares. Subjects could use either the sonification or a line graph to monitor the share price over time. Results showed no difference in performance between the two modes, but subjects reported a significant decrease in workload with sonification, which let them monitor the price while using their devices' visual display to buy and sell shares.

To investigate the orthogonality of the perceived sound qualities of pitch and loudness, Neuhoff et al. sonified both share price and volume for a single stock.[11] They mapped price to a tone's frequency and share volume to the sound amplitude. Subjects were tested on their ability to judge the share price and volume from the resulting sound. The interaction of loudness and pitch in sound perception is well known, but this experiment was the first to measure the effect in a sonification of abstract data. As expected, the results revealed perceptual interactions, in which asymmetries could potentially distort the display's interpretation. This highlights a problem with designing auditory displays, especially when using multivariate mappings.

One advantage of using sound displays is the ability to simultaneously monitor separate streams of sound.[3] Ben-Tal et al. explored this facet of sound perception by developing a simultaneous display of share price and volume for two stocks.[12] To ensure the separation of sound streams, they mapped the two stocks to perceptually distinct vowel-like sounds. They mapped the day's closing price to the number of sound bursts and the volume of trade to the bursts' duration. Next, they mapped the data for each day to a second of sound for periods of up to a year. They observed that it was possible to categorize high-volume, high-price trading days as loud, dense sounds, and low-volume, low-price days as pulsed rhythmic sounds.

Despite no compelling evidence for using sound displays to detect trading patterns, the intuition is that sound forms an excellent medium for monitoring temporal patterns that might occur. Because the domain expert—an experienced trader from an Australian trading firm—who collaborated in our case study has an interest in temporal patterns, we developed an auditory display. Although researchers had applied traditional share price data in a number of previous sonifications, they hadn't previously attempted the sonification of depth-of-market data. Our domain expert felt that the novelty of such a display provided a good opportunity for uncovering new types of patterns.

Depth-of-market data

Many charting techniques let researchers look for patterns across periods of trading. Although we can generalize these periods to short time frames, they traditionally cover one-day periods. The market's shorter-term players, such as day traders, might need to make minute-by-minute decisions from live feeds of stock data. In this case, each transaction can be charted so that the trader can exploit variations that occur every minute.

Besides price information, day traders might take advantage of data provided by the market's depth. *Depth of market* refers to the number of buyers and sellers currently trying to trade a financial instrument. A buyer might make a bid to purchase a specified volume of shares while at the same time a seller might ask a price for some specified volume. The balance of

Table 1. Depth-of-market data.

Buyers (bids)		Trade	Sellers (asks)	
Volume	Price	12.03	Price	Volume
14,533	12.03	1	12.04	42,450
28,850	12.02	2	12.06	20,540
23,000	12.02	3	12.07	8,261
2,121	11.99	4	12.09	35,000
41,000	11.98	5	12.10	120,515
17,000	11.97	6	12.11	574

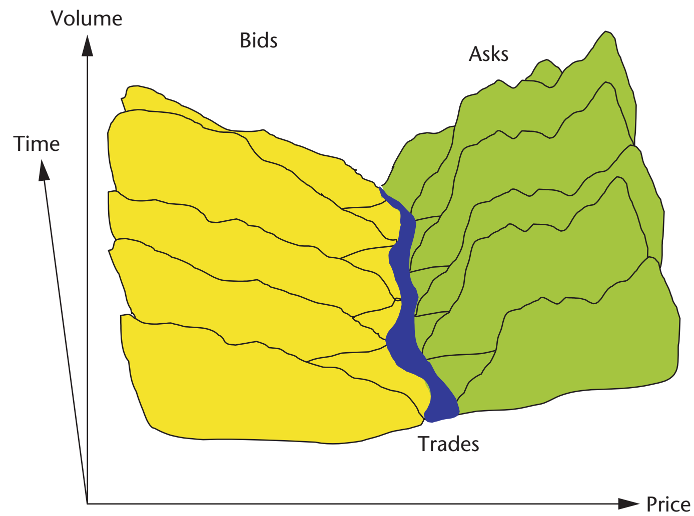

4 **Conceptual model of the depth-of-market visualization.**

bids and asks determines the current market state. The difference between the highest bid and lowest ask is the spread.

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 might sell when the volume of bids around the last trade price outweighs the volume of asks; or sell if asks outweigh bids.

Traditionally, trading software displays depth-of-market data in a table (see Table 1), which updates every 30 seconds or so. The top of the table displays the highest bid and lowest ask. A wider spread usually indicates a lower likelihood that a trade will occur. The table shows the price of the last trade 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: Bid-ask-landscape

We designed the visual display for the bid-ask-landscape based on discussions with our domain expert. The aim was to produce a display that better showed temporal patterns in depth-of-market data. Our display incorporates a natural landscape metaphor, that is, it displays depth-of-market data as hills and valleys that changed over time (see Figure 4). The changes to the landscape would ideally provide temporal signals that would help a trader predict market direction.

Three-dimensional surfaces frequently help researchers represent relationships between three variables, such as price, volume, and time. However, rather than using a single surface to display the data, the bid-ask-landscape subdivides the space, with two "surfaces" to represent the data. In this case, buyers (bids) are represented by one surface on the left in Figure 4 and the sellers (asks) are represented by another surface on the right. The buyers' bids naturally fall at a lower price, while the sellers' asks reflect a higher price. However, the surfaces are not independent but change in response to actions made by the other. The display captures the changing tension between the two market forces, and anticipates that a valley will form at the display's center as buyers and sellers exchange bids and asks to make a trade. A third surface represents these trades (see Figure 4).

In the figure, bids are yellow and asks are green. The highest bid and lowest ask are next to each other and close to the center. Less important data spreads to the periphery. The trades are shown as a blue river that also tracks the market spread.

Price is the key component that connects the buyers and sellers, which we placed on the horizontal axis. The volume of all bids and asks at that price point determines the height of the hills and valleys. The landscape is slightly complicated—bids accumulate as they move away from the maximum price, and asks accumulate as they move away from the minimum price ask. This can be understood by thinking of a buyer who is prepared to bid at $10. That same buyer is also prepared to bid at anything less than $10. The situation is similar for a seller who asks $12; that seller is also presumably happy to sell at anything greater than $12.

After we initially evaluated the display, we made a modification to incorporate a fish-eye distortion around the center of the bids and asks, as Figure 5 (next page) shows. By distorting the space within this region—stretching it along the *x*-axis—more space is dedicated to bids and asks close to the current area of trading. This new design emphasizes the use of detailed foveal-vision close to the landscape's center, which contains the most important data. The bids and asks closest to the last trade are likely to have the most impact on the short-term market direction. Less detailed peripheral vision lets traders monitor large changes on the data periphery.

The display's final axis is for time. Hence the 3D orthogonal space is defined by the quantitative attributes of price, volume, and time, as in Figure 4. Note that as new data updates the display over time, the visualization takes on the form of an evolving landscape resembling a valley between two hills with a river flowing through it. We propose that this metaphor may help users interpret the visualization from familiar natural properties such as cliffs' steepness or hills' height.

Bid-ask-landscape auditory display

Whereas the bid-ask-landscape's visual display shows the accumulated volume of bids and asks over time, the auditory display shows each individual bid and ask as they occur and provides information about relations

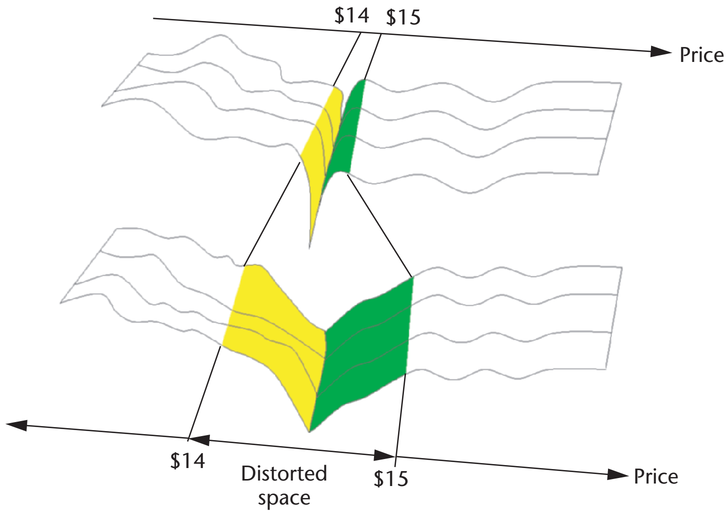

5 Space is distorted around the central valley where trades are occurring. This design affords more space for displaying the detail information most likely to impact trading prices.

Table 2. Display mapping used in the auditory display. The letters and numbers under "Pitch" refer to the seven major notes from the fifth octave in the key of C (C5, D5, E5, F5, G5, A5, and B5).

	Category	
Position (degrees)	**Pitch (musical notes)**	**Price (cents from last trade)**
–90	C5	≤ –4
–60	D5	–2, –3
–30	E5	–1
0	F5	0
30	G5	+1
60	A5	+2, +3
90	B5	≥+4

Table 3. Categories of price importance derived from the price of each bid and ask.

Bid or Ask offer of Last Trade (cents)	Price Importance
–4 or more	–low
–2, –3	–medium
–1	–high
0	0 high
1	+high
2, 3	+medium
4 or more	+low

between bids and asks such as the current spread. The auditory display also provides more detailed information about the market's current activity—that is, the number of bids and asks occurring, as well as the volume of each bid and ask. We combined displays to help users determine the direction of the next trade price, either up or down, from the last trade. For this reason, the auditory and visual display feature different types of information.

The auditory display has two levels—schema and perceptual.[2] At the schema level, we adopted a marketplace metaphor, suggested from the task scenario analysis we performed with the help of our domain expert. In this marketplace, vendors shout the price of produce and shoppers reply with offers or agree to trade, as Table 2 depicts. Listeners can interpret the direction of the next trade from this familiar experience.

At the perceptual level, we map information onto the perceptually scaled auditory variables of pitch and loudness. The design also addresses issues of perceptual grouping and segregation in the overall auditory scene. For example, the use of distinct timbres for displaying the bids and asks helps the listener group the sounds of buyers and distinguish these from the sounds of sellers.

We designed the sound display to enable quick, confident, and accurate answers to the global question, What is the direction of the next trade? It should also enable answers to intermediate questions about the 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? The individual events in the auditory display should allow answers to questions such as, Are there any bids? or What is the volume of the most recent ask?

We implemented the schema level with samples of a male voice saying the word "buy" and a female voice saying the word "sell." We put the male voice at a pitch of E5 (that is, the note E in the fifth octave in the key of C) and the female voice at G5. We edited the samples to a duration of 0.25 second. The difference in samples' timbre, vowel formant, and pitch lets users hear both "buy" and "sell" distinctively when played simultaneously.

The data we want to display consists of the price and volume of bids, asks, and trades. We mapped the price data for each bid and ask into seven ordered information categories of price importance, as Table 3 shows. Each category represents the importance of the bid or ask relative to the last trade. As with the visual display, this mapping emphasizes bids and asks close to the price of the last trade. The volume data, considered less important than price, is mapped to three ordered levels of volume importance on a logarithmic scale. Shares less than or equal to 10,000 have low volume importance; shares less than or equal to 100,000 have medium importance, and shares equal to or greater than 100,000 have high importance. A combined mapping from price and volume importance to overall importance places each offer into one of seven ordered categories of importance, as Table 4 shows. Finally, we use each offer's overall importance to define the final auditory parameters in displaying that offer, as Table 5 shows.

In Table 5 we mapped the seven ordered information categories to seven ordered pitch categories, C5 through B5, within the same octave. This mapping allows the perception of order without overlapping har-

6 **The bid-ask-landscape was implemented and evaluated on the Barco Baron at CSIRO's Virtual Environment Laboratory in Canberra. (Photos courtesy of Matt Adcock, CSIRO, Mathematical and Information Science, Canberra.)**

Table 4. Overall importance categories derived from price and volume importance.

Price Importance	High Volume Importance	Medium Volume Importance	Low Volume Importance
–low	–low	–low	–low
–medium	–medium	–low	–low
–high	–high	–medium	–low
0 high	0 high	0 medium	0 low
+high	+high	+medium	+low
+medium	+medium	+low	+low
+low	+low	+low	+low

monic or octave pitches. However, pitch doesn't allow for the perception of the central zero (fourth information category). We addressed the zero by a redundant mapping to seven categorical spatial locations in the stereo display, ordered from left to right with the zero in the middle. The listener can identify the zero elements at the display's spatially absolute center. As price diverges from the last trade, the sound moves further away from the central position. Mapping the information categories to both pitch and space increases the perceptual segregation and perceived order of the categories. The three ordered categories of share volume are mapped to three equal steps in loudness.

This mapping reflects that volume close to the last trade price can strongly influence the direction of the next trade, but that share volume isn't important in the price periphery. Note that no trade information is explicitly displayed. Our early design iterations displayed each trade as it occurred; however, our domain expert on this project found that this complicated the display without providing useful information.

Our system displays all bids and asks sequentially as they occur in time. The overall effect of listening to the display is that bids and asks are heard in real time as they are made. The bids and asks come from the left and right, respectively, as though from a crowd. If the bid or ask is lower than the last trade, the user hears it to the left. If the offer is the same as the last trade price, the user hears it from the center; and if it's higher, 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 market equilibrium.

The general level of market activity can also be discerned from the sound display. A highly active market sounds like a continuous hubbub. A mid-range of activity—around 20 events per time step—sounds intermittent and overlapping. When market activity is low, users can hear the individual events. Silence, of course, indicates a lack of activity.

Implementation and evaluation

We implemented the 3D visual display on a Barco Baron Stereo Projection Table, at the CSIRO Virtual Environment Laboratory in Canberra, Australia, as Figure 6 shows. The Barco screen connects to an SGI Onyx2 computer with synchronized shutter glasses and a Polyhemus head tracker. Sennheiser HD540 headphones provide the sound. We built the visualization on the Avango Virtual Reality framework, developed by the IMK-VE group at the Fraunhofer Institute for Media Communication. The User Datagram Protocol sends data updates from the visualization to the sonification. We built the sonification with the Avango sound server and Max synthesis system.

The visual display is updated every 3 seconds. Each

Table 5. Mapping from overall importance to pitch, spatial position, and loudness.

Overall Importance	Pitch Category	Spatial category or pan(degrees)	Loudness (range)
–low	C5	–90	low
–medium	D5	–60	low, medium
–high	E5	–30	low, medium, high
0 high	F5	0	low, medium, high
+high	G5	30	low, medium, high
+medium	A5	60	low, medium
+low	B5	90	low

Table 6. Results of tests for significant variation ($p < 0.05$) of the main variables. The two darkened cells below indicate significant results.

Direction	Subject	Mode	Order
All	0.884	0.397	0.953
Up	0.981	0.812	0.793
Down	0.014	0.029	0.653

visual time step represents 30 seconds of real-time data recorded from the stock market. Every bid and ask event triggers an auditory sound. We maintain the bid-and-ask event sequence; however we compressed time by a factor of 10. This ensures that test subjects hear all bids and asks occurring within the 30 seconds of real-time data in the 3 seconds before the next update of the visual model.

Experimental design

Our null hypothesis was that subjects couldn't predict the direction of the next trade from the visual and auditory displays. The alternative hypothesis—that subjects can predict trade direction—depends critically on the technical trading hypothesis that the data contains the needed information. Because we based our design on this premise, we used data recorded from real trading data for two shares during the opening and closing hour of a trading day on the Australian Stock Exchange. We divided the data into six subsets, three for training and three for evaluation: visual training, visual evaluation, auditory training, auditory evaluation, multisensory training, and multisensory evaluation. We randomly allocated each test subject one of these subsets. We also randomized the presentation order of the different display modes (visual, auditory, and multisensory).

Each subject carried out the experiment individually with a researcher present to record responses. At the start, we gave subjects a written introduction to trading with depth-of-market stock data, and allowed them to ask questions. Next, they participated in a training session followed by an evaluation session for each of the three modes. The training time was intentionally minimized because we'd developed the display to work perceptually and thus be intuitive. In the training session, we showed the subjects a display of historical data that was paused at 10 random points. At each point, they were told the direction of the next trade price—up or down. This mimicked the actual testing where the subjects were asked at 10 random locations to predict whether the next trading price was above (up) or below (down) the last trade. The up and down movements were between 1 and 7 cents, with 80 percent of the decision points involving only 1- or 2-cent changes. After each evaluation, we asked subjects for comments about how they used the display to make decisions.

There were 13 male and 2 female subjects 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. We recorded 10 predictions for each subject in each of the three modes; thus collecting 150 data points for each display mode. The experiment, including training and testing, typically took 45 minutes for each subject.

Results and analysis

We originally analyzed all of the 450 predictions made in the experiment as a single set of results. However, because some patterns in the display might be selective for either up or down movements of the trade price, we next performed a separate analysis considering the 227 predictions made at up locations and the 223 predictions made at down locations.

Overall, the analysis aimed to elicit answers to the following questions:

- Did results show significant variations across subject, mode of display, or for the order in which each mode was used?
- Can people use the visual, auditory, or multisensory displays to predict the direction of the next trade from depth-of-market data?
- Does the multisensory display function in a complementary, redundant, or conflicting manner?
- What differences are there in performance with the visual, auditory, and multisensory displays?
- Do people find consistent patterns in the data?
- How do people interpret and make decisions from these displays?

Three possible causes of variation existed in the results:

- mode (auditory, visual, multisensory),
- order (for example, auditory, followed by multisensory, followed by visual), and
- subject.

Subjects were tested on the three displays in random order. It could be expected that the ordering of display types in the experiment might influence the results. Subjects might, for example, undergo a training effect.

We analyzed the total correct predictions out of 10 by regression analysis and, as Table 6 indicates, found no significant effect for variation in subject, mode, or order. This lack of effect led us to analyze, using generalized linear models, proportions of correct predictions for trades that went up in price separately from trades that went down. For predicting trades in the up direction, we again found no significant effects. In the down direction the order still wasn't significant. However, we found that some subjects performed significantly better than others ($p = 0.014$) when predicting down movements, and that there was a significant variation in performance with mode ($p = 0.029$).

Next, we analyzed the results to see how well the subjects actually predicted the direction of the next trade. If subjects guessed their prediction randomly, then the expected result would be a binomial distribution with a probability of 0.5. To determine how well the subjects were performing, we compared the results to exact binomial probabilities that these results (or more extreme ones) would be obtained by chance. This corresponds to calculating the size of the two tails of the binomial

Table 7. Experimental analysis for all decisions and for decisions made at up and down movements of trade price. The probability of the result occurring by chance is shown. A $p < 0.05$ is considered significant. The darkened cells below indicate significant results.

Direction	Mode*	Correct Predictions	Total Predictions	Correct Predictions (%)	Standard Error	Probability of Results
All	V	94	150	62.6	0.047	0.0024
	A	105	150	70.0	0.047	0.0000
	M	105	150	70.0	0.047	0.0000
Up	V	50	85	58.8	0.077	0.1284
	A	46	78	59.0	0.079	0.1405
	M	42	64	65.6	0.085	0.0169
Down	V	44	65	67.7	0.050	0.0060
	A	59	72	81.9	0.036	0.0000
	M	63	86	73.3	0.040	0.0000

*V = visual; A = auditory; M = multisensory.

distribution, to perform a two-tailed test. Analyzing the combined results showed us that subjects could predict the direction of the next trade at levels significantly above a chance guess in all three modes. As Table 7 indicates, there is a 70 percent above chance of a correct prediction with multisensory display mode and with auditory and 62.6 percent above with visual.

In the next phase of analysis, we calculated the significance of the results for decisions made at up and down locations (see Table 7). Once again, we compared results against chance using a two-tailed binomial distribution. The analysis shows that for up movements in trade price, subjects don't predict the direction of the next trade at levels significantly above chance using the visual (58.8 percent) or auditory (59.0 percent) display. However, subjects perform significantly better (65.6 percent) than chance when using the multisensory display. For down movements in trade price, subjects perform significantly better than chance in the auditory (81.9 percent), combined (73.3 percent), and visual (67.7 percent) modes.

Given that our goal was to design a complementary multisensory display, analyzing the results for subject performance with the different modes wasn't straightforward. When subjects predicted upward movements, the combined display was the only mode where performance was significant. This indicates that the multisensory display was complementary. However, subjects were able to predict down trades 81.9 percent of the time from the auditory display and only 73.3 percent of the time using the multisensory display. This indicates that the multisensory display presents subjects with conflicting information when used for predicting down movements. When all responses are considered, we find that performance with the auditory and multisensory displays are similar—subjects predicted market direction 70 percent of the time with both. This indicates that the multisensory display is redundant for the prediction task.

Returning to our original motivation of data mining, we'd ideally validate the presence of patterns, identified by users, in the data. One subject commented that specific points in the display enabled the subject to make a decision with certainty; at other points, the subject wasn't so sure what to decide. This suggests that, at some places in the data, clear-cut patterns exist.

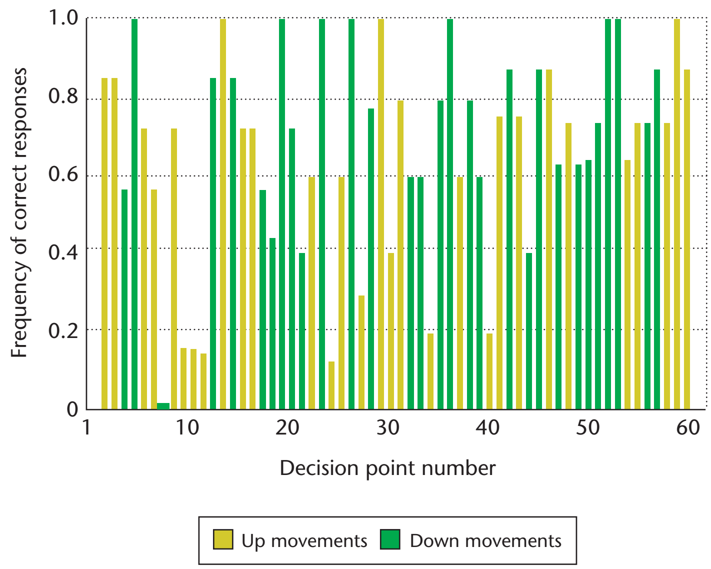

7 **Proportion of correct predictions at each of the 60 decision points. The direction of price movement at each decision point is indicated by color.**

In our data set, subjects made decisions about market direction at 60 unique decision points. We analyzed the frequency of correct responses at each of these decision points (see Figure 7). In total, subjects consistently predicted direction at 10 of the 60 decision points, and 7 of these were at places where down movements in price occurred. At one point in the data, subjects were also consistently wrong with their predictions, which might indicate that users misinterpreted a pattern. Although the validation of these patterns requires a more careful follow-up, the patterns might well provide the basis for viable trading rules.

To understand how subjects made decisions, we recorded their comments after their evaluation in each mode. After using the visual display, nine subjects commented that they made decisions based on size, height, slope, and steepness of cliffs; on how close the peaks were to the center; and on bending and trends in the river or valley. Three subjects said they couldn't understand how to make decisions. One commented that it wasn't 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, subjects made decisions from frequency of calls, closeness to the center, and loudness. Six subjects 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.

In the multisensory display, subjects commented that the combined display contained more information than visual or auditory display alone, and that the visual display provided context, history, past, and general trends, while the auditory provided most recent trends, focus, eagerness to trade. Subjects also commented that the auditory display increased their level of presence, in that it made the model feel more real and maintained their attention. Four subjects said that the visual and auditory displays were sometimes in conflict, and three of these resolved this conflict by relying on the auditory information; the other relied on the visual display. One subject felt the conflicts in the multisensory display made it more ambiguous and preferred the auditory display alone. In contrast, another subject found the auditory display distracting and preferred the visual display over the multisensory display.

Conclusion

The aim of developing a multisensory display for data mining is to increase the computer–human bandwidth and thus provide the user with complementary information through different senses, rather than redundant or conflicting information through multiple senses. In our experiment, subjects' performance with the auditory and multisensory displays was similar, indicating that the combined display was redundant for finding patterns in depth-of-market data. However, analyzing the results for prediction of up and down trades showed that the situation wasn't so straightforward. The subjects could predict down trades from the auditory display more than 80 percent of the time, and three scored 100 percent. The lower prediction for downward movements using the multisensory display indicates that the visual display might have introduced noise or conflict when combined with the auditory display. Yet, in the upward direction, the combined multisensory display was the only mode where performance was significant, indicating that the combination of auditory and visual displays was complementary for predicting upward movements.

This evaluation suggests that useful trading patterns occur in both the auditory and visual displays. However, we need to perform much more work to isolate these patterns and determine their usefulness for real-world trading across a wider range of data. For example, future work will integrate a real-time data feed of depth-of-market data into the display. This will let the domain expert (stock market trader) as the principal user evaluate the display over a longer time period with a wider cross-section of data. Further evaluation is also required to compare trader performance on traditional table displays with the new multisensory display. Rather than an experimental evaluation, it's also possible to extract the heuristics used by subjects and incorporate these into an intelligent tool for automatically finding similar patterns in historical data. This will accommodate a more automated approach to testing the patterns found by users of the bid-ask-landscape.

Despite the display's success to date, we've identified several issues that must be addressed before we can qualify the display as an effective human perceptual tool for finding useful trading patterns. We need to address the following issues:

- Why do subjects predict down trades so well from the auditory display? Perhaps the auditory display is biased toward down decisions, or the display highlights a feature. The nature of the market might also bias results in the down direction.
- Why do some subjects perform significantly better with the auditory display yet there's little difference with the visual or multisensory displays?
- Why is the multisensory display complementary for up trades but redundant for down trades?
- What are the significant patterns in the visual and sound displays? Indications are that the display more easily and correctly predicts some decision points than others.
- How well do the predictions from these displays compare with a table? How do predictions from experts compare with nonexperts? Is the display still useful under real-world trading conditions? These questions could be explored by designing an experiment to measure the trading performance of experts using a real-time display.

The domain expert we are working with would like us to further extend the technical analysis approach to study temporal patterns in live depth-of-market data. ■

Acknowledgments

We thank stock market domain expert, Bernard Orenstein of Object Trading Pty Ltd., who was fundamental to the design and evolution of the bid-ask-landscape. We also thank Martin Göbel, Bernd Fröhlich, John Plate, and other members of the IMK-VE group of the Fraunhofer Institute for Media Communication for their assistance with the display's development. We thank Warren Muller from CSIRO Mathematical and Information Sciences, Canberra, Australia, for his assistance in designing and analyzing the results of the experiment.

References

1. S.K. Card et al., *Readings in Information Visualization: Using Vision to Think*, Morgan Kaufmann, 1999.
2. S. Barrass, "Auditory Information Design," doctoral dissertation, Dept. of Computer Science, Australian Nat'l Univ., 1997.
3. G. Kramer, "An Introduction to Auditory Display," *Auditory Display: Sonification, Audification and Auditory Interfaces*,

vol. XVIII, G. Kramer, ed., Addison-Wesley, 1994, pp. 1-78.
4. K.V. Nesbitt, *Designing Multi-Sensory Displays for Abstract Data*, doctoral dissertation, Dept. of Computer Science, School of Information Technologies, Univ. of Sydney, Australia, 2003.
5. M.R. McGee, P.D. Gray, and S.A. Brewster, "The Effective Combination of Haptic and Auditory Textural Information," *Proc. 1st Workshop on Haptic Human–Computer Interaction,* Springer, 2001, pp. 33-38.
6. C. Nicholson, "Technical Analysis," *Course Notes E114*, Securities Inst. and Australian Tech. Analysts Assoc., 1999.
7. W. Jungmeister and D. Turo, *Adapting Treemaps to Stock Portfolio Visualisation*, tech. report, CS-TR-2996 , Dept. of Computer Science, Univ. of Maryland, 1992.
8. W. Wright, "Information Animation Applications in the Capital Markets," *Proc. IEEE Information Visualization Symp.*, IEEE CS Press, 1995.
9. S.P. Frysinger, "Applied Research in Auditory Data Representation," *Extracting Meaning from Complex Data: Proc. SPIE/SPSE Symp. Electronic Imaging,* vol. 1259, E. Farrell, ed., SPIE, 1990, pp. 130-139.
10. S.A. Brewster and R. Murray, "Presenting Dynamic Information on Mobile Computers," *Personal Technologies*, vol. 4, no. 2, 2000, pp. 209-212.
11. J.G. Neuhoff, G. Kramer, and J. Wayand, "Pitch and Loudness Interact in Auditory Displays: Can The Data Get Lost In The Map?" *J. Experimental Psychology Applications,* vol. 8, no. 1, Mar. 2002, pp. 17-25.
12. O. Ben-Tal, M. Daniels, and J. Berger, *De Natura Sonoris: Sonification of Complex Data*, World Scientific and Eng. Soc., 2001.

Keith Nesbitt *is a lecturer in computer science at Charles Sturt University in Bathurst, Australia, where he is also establishing a 3D Multisensory Games Laboratory. His research interests include designing multisensory information displays, virtual environments, and user interface technology. Nesbitt has a BMath and a MComp from the University of Newcastle, and a PhD in computer science from the University of Sydney.*

Stephen Barrass *leads advanced audio interfaces research at CSIRO Australia. His research interests include interaction design, design patterns, perceptual interfaces, interactive sonification, brain–computer interaction, and sports biofeedback. Barrass has a PhD in computer science from the Australian National University in Canberra, Australia. He is a board member of the International Community for Auditory Display (http://www.icad.org).*

Readers may contact Keith Nesbitt at the School of Information Technology, Charles Sturt Univ., Panorama Ave., Bathurst, 2795, Australia; knesbitt@csu.edu.au.