Experimenting with Haptic Attributes for Display of Abstract Data

Keith V. Nesbitt
Department of Computer Science and Software Engineering,
University of Sydney, Sydney, NSW. 2006. Australia
knesbitt@cs.newcastle.edu.au

Abstract

Information perceptualization has been defined to include the three domains of Information Visualization, Information Sonification and Information Tactilisation [1]. While Information Visualization is a developing field and Sonification is an emerging field of study, very little work has been done to experiment with haptics for the display of abstract information. This paper describes some pilot studies designed to test the feasibility of using three different attributes of touch to display nominal or ordinal data. These attributes are surface hardness, surface roughness and object inertia. The results are applied to the display of stock market data on the Haptic Workbench.

1. Introduction

Information Visualisation is a developing field of study that aims to display abstract information to the human visual sense. The emphasis is on allowing perceptual cues in the visualisation to convey information about structure, patterns or rules in the data [1]. Many examples of information visualisation have been developed over the years in a wide range of fields including the display of stock market data [2], network maps of telecommunication systems [3] and for understanding complex software systems [4].

By contrast Information Sonification is only a newly emerging field and while examples of using sound for information display have been used in some traditional fields such as sonar tracking in submarines and as alarm systems in aircraft the attempt to map abstract data attributes to the attributes of sound is a relatively new endeavor [5]. Still a number of interesting applications have been developed which include a simple sound enhanced interface for the Macintosh Finder called SonicFinder [6], use of sound to help understand and debug parallel computer programs [7], and auditory versions of 2D statistical scatter plots [8].

Information Tactilisation, as the display of abstract information to the sense of touch has been termed [1] has the potential to also emerge as a new field of study. Some examples exist, force feedback has been used to display flow fields in blast furnace data [9], to assist with interpretation of seismic exploration data [10] and to understand soil properties [11]. Despite these examples, as yet only a little work has been done in this area. The availability of suitable hardware to display feedback to the sense of touch has been one drawback. Recently the accessibility of the PHANToM™ force feedback device [12] has meant it is now possible to provide the user with forces that act on the proprioceptive sensors, hence displaying to the user some attributes of touch.

Section2. Haptic Workbench and the Reachin API describes the Virtual Environment hardware as well as the platform used to develop and perform the studies. *Section 4. Pilot Studies* describes the rationale, methodology and results of each study. *Section 5. Feeling Stock Market Data* introduces the application of the pilot studies to some stock market models. Future work is discussed in S*ection 6. Conclusion*.

2. Haptic Workbench and the Reachin API

The Haptic Workbench (figure 1) is a single-user, multi-sensory interaction platform**.** It provides support for auditory, haptic and stereographic visualization [13]. It was developed at CSIRO's Division of Mathematical and Information Sciences, Canberra and incorporates the PHANToM™ force feedback device [12] from SensAble Technology Inc. [14] with the 'Virtual Workbench' [15] stereographic display platform, developed at the Institute of Systems Science of the National University of Singapore.

The Reachin API produced by ReachIn Technologies, Sweden [16] (and previously known as Magma) was developed to provide an integrated haptic and graphic rendering API. It is an object oriented C++ API and is based on the familiar hierarchical scene graph concept where the scene graph contains nodes that describe a 3D scene. The API is in fact an extension of the VRML Scene graph model with additional haptic properties. For example as part of the "Appearance" node a "SimpleSurface" can be added which makes the surface of an object "feelable" by implementing a surface repulsion algorithm [17]. The Reachin API greatly reduces the complexity of implementing haptic objects and is fully extensible to implement new force models. The experiments described in this paper use some of the simple models provided by the Reachin API, the “SimpleSurface” and the “RoughSurface” as well as one extended model of forces called the “SpringSlotDynamic”. This node builds on the Reachin API node called a “SlotDynamic” and was implemented at the CSIRO Virtual Environment Lab [18].

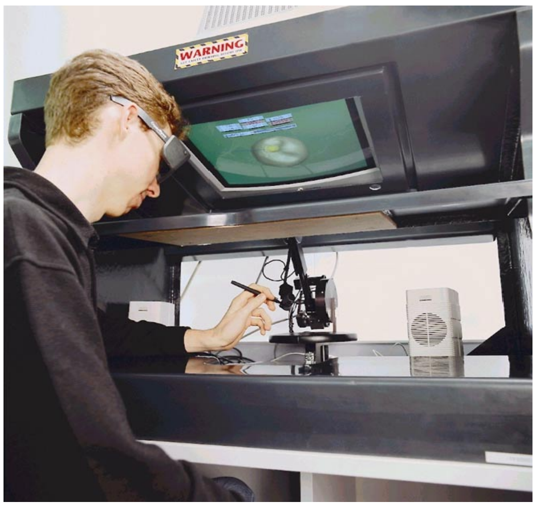

Figure 1. The Haptic Workbench at CSIRO

3. Pilot Studies

Three different touch parameters were studied for the mapping of abstract data. These were surface roughness/smoothness, surface hardness/softness and object inertia. Each of these three properties is described in more detail below.

Initially for each of these properties ten cubes were rendered. Each of the cubes was designed to have a noticeable change in the property of interest. It was hoped that users would be able to rank these ten cubes based on a property like smoothness/roughness. For the purpose of this preliminary study this would indicate the usefulness of each property for displaying ordinal data with ten classes.

An original attempt at producing a linear step change in the parameter values did not allow the user to sufficiently differentiate between some cubes and also produced large changes between others. After some experimentation with values, this was largely a trial and error process, it was possible to produce cubes that provided a noticeable difference in the desired property. It was noted that some of the parameters changed in the model required a doubling of the parameter value at each step to produce a noticeable difference. This revealed power related function (a doubling at each step) should not be surprising given the common occurrence of such relations in perceptual theory.

Once the ten cubes were implemented it was shown to just three users before it became obvious from feedback that the users had a good deal of trouble disambiguating between many of the cubes. While it was a simple task to tell the roughest cube from the smoothest cube and also to rank some cubes in between, it was not possible to differentiate sufficiently between the cubes to rank all ten. It was decided to change the implementation and reduce the number of rendered cubes from ten to five. The aim at this stage of the study was just to determine at least some usefulness of these three haptic properties for the display of ordinal data.

Surface smoothness/roughness was implemented using the Reachin API "RoughSurface" node. This is a haptic surface description that can be used to describe an appearance of a VRML shape. In this case the shape has the simple geometry of a cube. Most parameters were held constant and only the "mean" value was changed. The roughness is actually implemented by modeling the mean friction on the surface of the object. In this experiment the "mean" value varies from 0.1 for a very smooth surface through to 0.9 for a very rough surface. Five cubes in total are produced, with "mean" values of 0.1, 0.3, 0.5, 0.7 and 0.9. The cubes are arranged randomly in a circular pattern to avoid any ordering in the spatial configuration. The material of the cube surfaces was kept as simple as possible, a simple white without texture. In this experiment the user was asked to rank the cubes based on the perceived roughness.

Surface hardness/softness was implemented using the Reachin API "SimpleSurface" node. A “SimpleSurface” is a haptic surface description that can be used to describe the “appearance” of a VRML shape. Again in this case the shape has the simple geometry of a cube. All default parameters were held constant and only the "stiffness" value was changed. This "stiffness" value varies from 30 for a very soft surface through to 480 for a very hard surface. Five cubes in total are produced, with "stiffness" values of 30, 60, 120, 240 and 480. Note how the value of this parameter needs to be doubled each time to produce a perceptual difference in the surface's hardness. The cubes are arranged randomly in a circular pattern to avoid any ordering in the spatial configuration. Again the material of the cube surfaces was kept to a simple white colour without texture. In this experiment the user was asked to rank the cubes based on the perceived hardness of each object.

Object inertia was implemented using a "SpringSlotDynamic". This is a class derived from the Reachin API "SlotDynamic" node. Objects constrained by this node's behaviour can be moved from their starting position by a force. Objects can only be moved in a straight line and will spring back to their starting position when the pushing force is removed. This is actually modeled by using a spring behind the object. By assigning different "mass" values to the "SpringSlotDynamic" it is possible to change the amount of force required from the user to push the objects. The user can push down on a cube, but feels a spring force pushing in the opposite direction. This corresponds to a sense of inertia related to the mass of the object.

As in previous examples this inertia concept was applied to VRML shape that had the simple geometry of a cube. The surface hardness or stiffness was maintained constant at 100 and only the mass" value was changed. This "mass" value varies from 2 for a very easy to move object through to 32 for a very difficult to move object. Five cubes in total are produced, with "mass" values of 2, 4, 8, 16 and 32. Once again note how the value of this parameter needs to be doubled each time to produce a perceptual difference in the objects inertia. The cubes are arranged randomly in a straight line across the screen allowing the user room to push down on each cube and feel its inertia. As in previous experiments the material of the cube surfaces was kept to a simple white colour without texture. In this experiment the user was asked to rank the cubes based on the perceived ease of movement of each object.

A group of ten users who had prior experience with the haptic workbench and so had used the PHANToM™ previously completed the experiment. While some users had considerable experience with using this interaction environment others had much less experience but were considered familiar enough so that no training time was required. There was no time constraint for the trials but each user typically spent about 2-3 minutes on each of the three tasks. The users were motivated to produce a good ranking by explaining that the experiment was designed to test their proficiency with using haptic feedback.

The results from the experiments were very conclusive. Only one user in ten made an error when ranking the objects using the property of surface roughness. This was mistake was due to uncertainity between the fourth and fifth roughest object. No users made errors ranking the objects by surface hardness. Again, one user in ten (a different user) made a single error ranking the objects by inertia. This error was between the ranking of the third and fourth heaviest objects. Despite the high success rate a number of users reported that the tests were not easy and they felt only reasonably confident about their results. Most users ranked the last task, that of determining object inertia while moving the different cubes as the hardest to complete. Though one user made an error the majority of users felt that the surface roughness task was the easiest.

A number of interesting observations and comments were collected. The procedure most users followed was to differentiate between the two extremes, for example, the smoothest and the roughest cubes. When these two were ranked they would iterate between the remaining cubes, comparing and contrasting the property of interest. Most users also commented that they adopted different strategies for using the PHANToM™ depending on the property of interest. For example, determining the roughness required a rubbing action across the surface, with some users alternating between using a gentle grip on the stylus for light rubbing and then a firm grip for applying more forceful rubbing. To determine surface hardness some users found that not only pushing but also tapping on the surface helped to distinguish the value of the property.

Some users commented that in the last experiment, were the cubes were pushed against a spring, that some visual cues were also present. The speed with which objects sprung back to their original position when the user's force was removed was dependent on the mass assigned to the spring. This visual feedback was a reasonably subtle cue that had not been noticed when implementing the experiment. However, three of the ten users reported using this to help in the ranking.

While the results from these studies did not serve as a formal experiment they did provide a positive indication that haptic properties could be used to display categorical data. The next stage of the study was to apply these haptic properties to the real domain area of Technical Analysis.

4. Feeling Stock Market Data

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" [19]. Analysts and traders attempt to make profitable trades by determining relationships within the data. 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.

Previous work has resulted in the development of some 3D visual models of price, volume and other market indicators to try and assist a Technical Analyst find patterns in stock market data [20]. Another focus of previous work has been the study of multi-sensory displays for increasing the amount of data that can be displayed using different sensory channels [21]. As a direct impetus of these two pieces of work we have extended two visual stock market models to incorporate feedback of the haptic properties studied above.

The simplest model is a 3D price chart [20]. This model enhances a normal time series of price bars (figure 2) with a variable such as volume in the third dimension. Volume is often used to confirm price signals such as a trading range breakout. 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 (figure 3). Another indicator that is often compared against price is momentum. Momentum is an indication of how fast the price is moving and in what direction. A high value of momentum indicates a volatile market in either direction while a zero value of momentum indicates that price is not moving. The momentum can be calculated over various time periods.

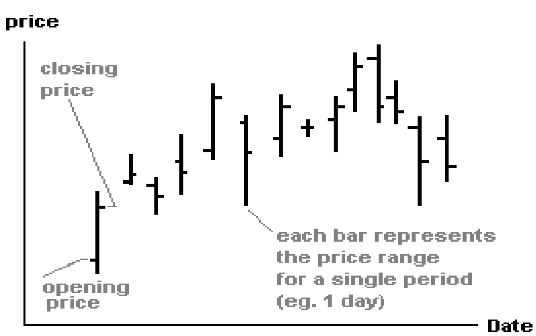

Figure 2. A traditional daily bar chart.

This 3D price chart was enhanced so that a daily momentum calculation was displayed using haptic inertia. Hence a price bar could be moved in the direction of the momentum dependent on its momentum. So if there was zero momentum if was not possible to move the price. If there was large momentum in the positive direction it was possible to easily move the price bar upwards. The smaller the momentum the harder the bar was to move. Momentum in the negative direction, constrained the bars to be moveable in a downwards direction. Once again the ease of movement was correlated with the momentum. High momentum price bars were easy to move and low momentum bars required greater force from the user.

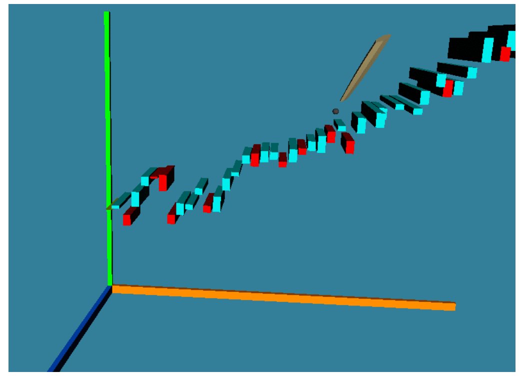

Figure 3. A view of the 3D bar chart.

A more complex 3D visual model is the moving average landscape [20]. In some technical analysis tasks it is useful to smooth out fluctuations that occur in price at each time step. 'Moving averages' [19] 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. To assist with the analysis of signals from multiple moving average curves a surface of moving averages was constructed (figure 4). This surface is constructed by joining together a number of strips. Each strip represents a different moving average curve from 1 day to 30 days. These moving average strips are joined to create a continuous surface. Overlaid on this surface are price bars that extend across the surface (figure 5). The position of the price bars above and below the surface may contain interesting information for the analyst [19].

The volume of trades is also a helpful attribute to display to the analyst. This surface was augmented to display volume with both the haptic property of hardness/softness and the haptic property of roughness/smoothness. In both cases the same information was displayed, namely volume of trades scaled between the low and high values used for each of the haptic properties in the experiments. Volume of trades can give the analysts and indication of how serious a period of trading is. For example, during a low volume of trading, price tends to fall but this is not indicative of real price action. Volume is also important for confirming signals in the price curve. The metaphor adopted for the display of the softness/hardness property was that low volumes of trade feet soft, while high volumes felt hard. This was meant to indicate that high volumes of trade were "firmer" indications of real price action. In a similar way the smoothness/roughness property was used so that price bars with low volumes of trade feel rough and price bars with high volumes feel hard. This was meant to indicate that rough prices are not very accurate indicators or "rough" indicators of real price.

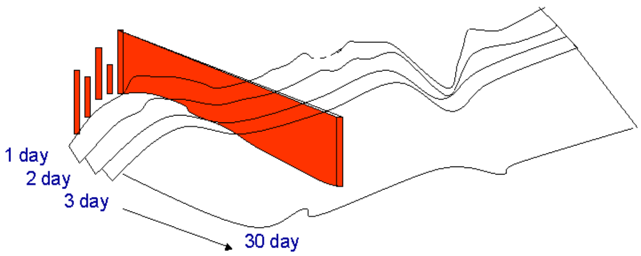

Figure 4. Extending the price bars to intersect with the moving average plane.

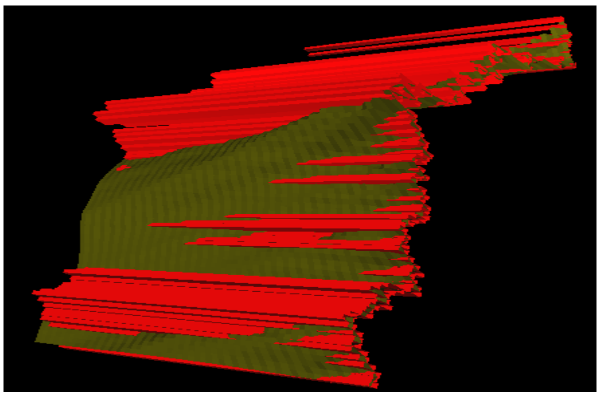

Figure 5. The moving average surface and intersecting price bars.

At this stage of the work we carried out a heuristic evaluation of these applications. The work was still considered too preliminary to justify a more formal experiment. Three users evaluated the system. One of these uses was a expert in technical analysis. All three users reported that the interaction metaphor was intuitive and easy to understand once the context was explained. Other feedback about the "usefulness" of these displays is not conclusive. It is difficult to define experiments were the task is to "find new patterns in the data". A major difficulty with this work is that there are simple interactive ways to present the same data visually. An example is brushing [1] where moving the cursor over the data displays the exact value at that location. The use of visual and indeed auditory feedback of information is quite likely preferable in many instances but the final goal of this work is to extend the human-computer bandwidth by filling up all available channels, visual, auditory and haptic with useful information. Indeed one benefit of haptic data is that it remains hidden until the user wishes to interact with it in some local domain. From the applications described and the informal studies perfomred here it is difficult to make any claims other than the analyst can indeed discern categorical variations in the data based on haptic properties.

5. Conclusion

It has been shown in these experiments that it is possible to use force-feedback from the PHANToM™ to display different haptic attributes to a user. Furthermore it is possible for the user to differentiate and rank these properties. The properties studied in this trial were surface roughness, surface hardness and object inertia. For each property it was shown that it is possible for the user to rank at least five classes of data. Therefore these haptic attributes are suitable for displaying some classes of ordinal data. What is required is more formal psychophysical experiments to determine the number of categories which can be distinguished for the different properties. This could involve determination of the JND for each property.

It has been shown how the ability to feel data attributes can be applied in the analysis of stock market data. There has been no attempt to justify the use of force as better than vision to display data attributes. The end goal of this work is to produce displays that are multi-sensory and we have to imagine that the visual and perhaps auditory domains have been saturated with data before we might employ haptic displays. One advantage of haptics that this work demonstrates is that it does not clutter the visual display. The user is only aware of the haptic properties of objects when they interact directly with them. Once again more formal experiments are required to discern the usefulness of touch in the task of data interpretation.

It is noted that interesting haptic metaphors come to light when we use this sense to display data. For example the stock market price has the concept of movement. Sometimes the price moves easily and other times it is more difficult to move. This can be mapped to object inertia. Roughness of data can be related to the noise in data and mapped to surface roughness. The idea that data is hard or certain and other times soft and so less certain can also be mapped to surface hardness. Despite these possible benefits much more work needs to be done in measuring the benefits of both haptics and multi-sensory displays in some more specific data analysis tasks.

6. 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 'Haptic Workbench' and the experiments themselves were made possible with the assistance and support of staff at the

CSIRO Virtual Environment Lab in Canberra, Australia [18].

7. References

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

[2] W. Wright, W. "Information Animation Applications in the Capital Markets". *Proceedings IEEE Information Visualization 95*, New York. 1995.

[3] R. A. Becker and S. G. Eick (1995). "Visualizing Network Data." *IEEE Transactions on Visualization and Computer Graphics* **1**(1): 16-28. 1995.

[4] N. Nishikawa "ObjectOrrey - 3D Visualization of Class Library Structure". *Proceedings Software Visualization Workshop*, Sydney, 1999.

[5] G. Kramer, G. *An Introduction to Auditory Display. Auditory Display:Sonification, Audification and Auditory Interfaces*. G. Kramer, Addison-Wesley Publishing Company. **XVIII:** 1-78. 1994.

[6] W. W. Gaver. (1986). "Auditory Icons:using sound in computer Interfaces." *Human Computer Interaction* **2**: 167-177. 1986.

[7] J. A. Jackson and J. M. Francioni. "Synchronization of Visual and Aural Parallel Program Performance Data". *Auditory Display:Sonification, Audification and Auditory Interfaces*. G. Kramer, Addison-Wesley Publishing Company. **XVIII:** 291-306. 1994.

[8] J. H. Flowers, D. Buhman, et al. "Data Sonification from the Desktop: Should Sound be part of Standard Data Analysis". *International Conference on Auditory Display*, Palo Alto, California. 1996.

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

[10] K. V. Nesbitt, B. J. Orenstein, R. J. Gallimore, R.J. and J. P. McLaughlin. "The Haptic Workbench applied to petroleum 3D Seismic Interpretation." *The Second PHANToM User's Group Workshop*, 1997.

[11] D. F. Green and J. K. Salsibury. "Soil Simulation with a PHANToM." *The Third PHANToM User's Group Workshop*, Cambridge, Massachusett, USA, MIT. 1998

[12] J. K. Salisbury and M. A. Srinivasan, "Phantom-Based Haptic Interaction with Virtual Objects". *IEEE Computer Graphics and Applications*, September-October, 1997, Vol. 17, No. 5, pp. 6-10.

[13] D.R. Stevenson, K.A. Smith, J.P. McLaughlin, C.J. Gunn, J.P. Veldkamp and M.J. Dixon, "Haptic Workbench: A Multi-Sensory Virtual Environment", *Stereoscopic Displays and Virtual Reality Systems VI*, J.O. Merrit, M.T. Bolas, S.S. Fisher, Editors. Proc. of SPIE, Vol 3639. pp 356-366. 1999.

[14] Internet web site: SensAble Technology. http://www.sensable.com/

[15] T. Poston and L. Serra, "Dextrous Virtual Work". *Communications of the ACM* 39, No. 5, pp.37-45. 1996.

[16] Internet web site: ReachIn Technologies, Sweden. http://www.reachin.se/

[17] Reachin API Programmer's Guide, Reachin Core Technology, Revision No. 300-501-002-03002, 1998-2001 Reachin AB

[18] Internet web site: CSIRO Virtual Environment Lab http://www.cmis.csiro.au/imvs/Technology/workbench.htm.

[19] Technical Analysis : Course Notes from Securities Institute of Australia course in Technical Analysis. (E114), 1999. http://www.securities.edu.au

[20] K. V. Nesbitt. "Interacting with Stock Market Data in a Virtual Environment" *Joint Eurographics IEEE TCVG Symposium on Visualization*, Ascona, Switzerland, SpringerWein New York. 2001.

[21] K. V. Nesbitt, "Designing Multi-sensory Models for Finding Patterns in Stock Market Data". *Proceedings of International Conference on Multimodal Interfaces*, Beijing. 2000.