

# Visualisation tools for blind people using multiple modalities

## Abstract

*Purpose:* There are many problems when blind people need to access visualisations such as graphs and tables. Current speech or raised-paper technology does not provide a good solution. Our approach is to use non-speech sounds and haptics to allow a richer and more flexible form of access to graphs and tables.

*Method:* Two experiments are reported that test out designs for both sound and haptic graph solutions. In the audio case we compare a standard speech interface to one with non-speech sounds added. The haptic experiment compares two different graph designs to see which was the most effective.

*Results:* Our results for the sound graphs showed a significant decrease in subjective workload, reduced time taken to complete tasks and reduced errors as compared to a standard speech interface. For the haptic graphs we again show reductions in workload and some of the problems that can occur when using such graphs.

*Conclusions:* Using non-speech sound and haptics can significantly improve interaction with visualisations such as graphs. This multimodal approach makes the most of the senses our users have to provide access to information in more flexible ways.

## Introduction

One of the main deprivations caused by blindness is the problem of access to information. Visualisation is a fundamental method for understanding information but at the present time it is almost completely unavailable to blind people. It is concerned with “exploring data and information in such a way as to gain understanding and insight into the data” [5]. Visualisation occurs in many aspects of life from reading a newspaper (where a graph might depict changes in the Pound against the Dollar) to statistical packages used by many students to view data

from experiments in the form of tables, graphs or 3D plots. This is almost always done visually so it is hard for a blind person to get at the information in an easy, usable way. Being able to perceive and understand complex information is particularly important for students at school and university so this is where we are focusing our research.

Current techniques for displaying information non-visually rely mainly on synthetic speech and Braille. Users hear a line of digits read out or, if they read Braille, feel a row of digits. Consider a sighted person reading a matrix of numbers. He/she would immediately be able to make certain inferences about the data. For example, there may be larger numbers at the bottom right or top left. A blind person would be unable to capitalise on these patterns in the data. He/she would just hear rows of numbers spoken, one after another. This is very slow. Properties of human short-term memory also mean that listeners are unable to hold in mind enough information to make any non-trivial observations - they become overloaded [1]. In addition, there is no way to get a quick overview or 'glance' at the data as a sighted person could [14]; it has to be listened to serially. Things become even worse when graphs or complex 3D plots are used because there are almost no techniques for presenting these in a non-visual way.

Our approach to solve the problem of visualisation is to make use of the senses that our users do have, namely hearing and touch. Using multiple sensory modalities will allow a richer and more flexible way to get at the information non-visually. We are focusing on the use of non-speech sounds to complement the synthesised speech already available and the use of haptic or (force-feedback) devices to allow users to feel virtual objects [13, 15]. There has been little research into how these techniques can be used in visualisation systems for blind people. As Edwards *et al.* [7] say "Currently the greatest obstacle to the exploitation of the variety of communications channels now available is our lack of understanding of how to use them".

## NON-SPEECH SOUNDS

Synthesised speech has been used in screen-reading software for blind people for many years. However, there has been little use of structured non-speech sounds that can provide access to information in very different ways [4]. *Sonification* [9] (literally *visualisation in sound*) has successfully been used to present, amongst others, census, geographical and environmental data. Non-speech sound can provide alternative ways to understand data. For example, such sounds are much better at providing access to continuous data than speech is. Little of the work in this area has specifically investigated sonification for blind users. One of the main exceptions that is the work on *Soundgraphs* by Mansur *et al.* [11]. The *Soundgraphs* system allows the presentation of line graphs

in sound. Time is mapped to the x-axis and pitch to the y-axis. The shape of the graph can then be heard as a rising or falling note playing over time. Mansur's initial results were very promising with users able to identify types of curves along with maximum/minimum points very easily. It also allows listeners to get an overview by listening to all of the data very quickly; again something that speech is not good at. We have taken this early work and are developing it further. In this paper we investigate whether or not such sounds can be combined with speech in a more realistic application to allow users to browse complex graphs in real time and whether or not this gives any usability benefits.

## HAPTIC INTERACTION

Haptic (or force-feedback) devices allow users to feel virtual objects [12]. Minsky *et al.* (in [3]) describe the technology thus: "Force display technology works by using mechanical actuators to apply forces to the user. By simulating the physics of the user's virtual world, we can compute these forces in real-time, and then send them to the actuators so that the user feels them". The haptic device used in our research is the PHANToM from Sensable Technologies [12] (see Figure 1). This is a very high resolution, six degrees-of-freedom device in which the user holds the end of a motor-controlled, jointed arm. It provides a programmable sense of touch that allows users to feel textures and shapes of virtual objects, modulate and deform objects with a very high degree of realism. For our research we are allowing users to feel graphs and tables as if they were really present in front of them.

[Insert figure 1 about here]

Petrie and colleagues have done key initial work in the use of haptic devices for blind people [6]. They have shown that, for example, blind people are more discriminating than sighted in their assessment of roughness and texture when presented via a haptic device. They has also shown that the perception of larger objects is easier than smaller. Haptic devices allow users to interact directly with their data – to get an overview of a graph users could just run their fingers along it. This has many advantages over raised paper graphs that are used by blind people [10]. For example, our system is dynamic so that we can render a haptic scene in real-time, rather than having to wait for a raised paper graph to be printed. Our scenes can be fully three-dimensional rather than just raised lines. Users can also change the graphs themselves, e.g. by changing the value of X in the graph and seeing what effect it has, just as a sighted person might do with pen and paper. In this paper we investigate how we might use haptic devices to present visualisations such as line graphs.

## **Our research**

In this paper we present two experiments to investigate the use of non-speech audio and haptic graphs. Our design approach is based on three steps. Our initial prototypes were developed with expert input from a local blind organisation (Glasgow and West of Scotland Society for the Blind). This allowed us to prototype our initial designs to ensure they would be effective with our target user group. We then tested these on sighted people (who could not see any screen display) and then finally a full experimental evaluation was carried out with blind people. We used sighted people as we have a ready supply of them and our pool of blind participants was much more limited (we had to travel long distances to get large enough groups of blind users to test). Using this approach meant that we could collect data from sighted people and refine our designs before we tested with blind people, thus ensuring we had the best interface designs we could when we used blind participants.

The rest of this paper presents two experiments: one on sound graphs and one on haptic graphs. At this stage the two technologies have not been integrated. We wanted to learn more about how each worked alone before we combine them later on in our research.

## **Sound graphs**

We were interested in finding out whether there is any value in using sound graphs [11] for improving the accessibility of graphs and tables containing numerical information for visually impaired people. So far, there is little evidence to show convincingly whether this strategy can work for real world data being investigated in real-time by a user. By observing participants tackle data comprehensions tasks we hope to identify the strong and weak points of our design.

The application uses the MIDI (Musical Instrument Digital Interface) synthesiser of a PC soundcard to generate non-speech sounds and the Microsoft speech synthesiser to generate speech. Participants have access to the auditory messages through headphones and navigate the data using a numeric keypad.

Data values are mapped directly to sounds; the value on the y-axis of a graph is mapped directly to the MIDI note number, the X-axis is mapped to time. This means the listener can hear the graph play with rising and falling notes representing the data. For this experiment we restricted the range of y-values to 127 (the MIDI note range). The volume of the sounds is kept constant. The mapping also involves panning the sound sources. The

sounds are localised along a line joining the left and right ears and positioned according to the position of their relevant cell. For example, the first value might be heard on the left, the last value on the right and the intermediate values are heard in between the left and right ears.

The keyboard interface to our system is based around the standard numeric keypad. This is similar to the way standard screen readers such as JAWS (one of the most popular screen readers, from Henter-Joyce, Inc.) are controlled. It allows users to browse graphs in speech (with each value read in turn) or by pitch (listening to the different values being played back as a sound graph), or a combination of both (listening to values until an interesting one is reached and then its value can be spoken). These different modes are chosen by keys on the keypad.

## THE EXPERIMENT

An experiment was conducted to investigate if the addition of sound graphs to a speech-based interface would improve the use of graphs by blind users. The purpose of this experiment was to evaluate how successfully participants tackle data comprehension tasks under two auditory conditions namely (1) speech only ('Speech') and (2) speech and pitch ('Pitch'). Our main hypotheses were:

H.1 Overall subjective workload would be reduced under the pitch condition. The sound graphs would make it easier for participants to understand and use the data so would generate less workload.

H.2 Time taken to complete the tasks would be reduced in the pitch condition. The sound graphs would make answering the questions faster as they provide a better way to get at the data.

H.3 Number of correct answers to the data comprehension tasks would be increased. The sound graphs would make it easier for participants to make comparisons and understand the data.

An initial test of our system on 16 sighted subjects allowed us to ensure the design was effective ready for testing on blind people.

## ***EXPERIMENTAL DESIGN AND PROCEDURE***

Sixteen participants were used in a repeated measures design. Participants used both the speech and pitch interfaces in a counterbalanced order. Participants (for this and the experiment on haptic graphs described below) were recruited from two schools for blind children and adults in the UK (RNIB College, Peterborough

and Queen Margaret College, Hereford). The age of the participants ranged from 23 to 57, with 7 women and 9 men. They had diverse visual impairments but none had any residual vision for tasks relevant to our experiment.

In order to get a full range of quantitative and qualitative results time, error rates and subjective workload measures were used. Time and error rate reductions would show quantitative improvements and workload differences would show qualitative differences. This gives a balanced view of the usability of a system [2]. Hart and Wickens ([8], p 258) define workload “...as the effort invested by the human operator into task performance; workload arises from the interaction between a particular and task and the performer”. Workload is broken into six different factors: Mental demand, physical demand, time pressure, effort expended, performance level achieved and frustration experienced. The NASA-Task Load Index (TLX) [8] is a set of six rating scales and was used for estimating these subjective workload factors in both of the experiments described here.

Two sets of questions were used for the data comprehension tasks: Student statistics and London crime statistics. The data comprehension tasks were constructed in the same style as exercises used to test the basic graph or table analytic skills of late primary or early secondary school students in the UK (see Figure 1). For example, here are the questions asked of the London Crime Statistics data:

1. State the year(s) in which the highest number of murder cases was/were reported.
2. State the year(s) in which the highest number of robbery cases was/were reported.
3. Which type(s) of crime had a consistently high number (~>50) of cases reported?
4. Which type(s) of crime show(s) a consistent increasing trend?
5. State the year(s) in which the numbers of hate crime cases was/were lowest.
6. Which type(s) of crime show(s) a consistent decreasing trend?

[Insert Figure 2 about here]

Participants had training on each of the interfaces (speech or pitch) before they used the interface. During the training phase the design of each of the systems was described to the participants and they were then allowed to experiment with a different data set (describing the gross national products of a country over a number of years). The experimental conditions were then administered and the participants were then asked to work through a series of questions such as those above. At the end of each condition they were given a NASA TLX workload test (see Table 1).

[Insert table 1 about here]

## ***RESULTS AND DISCUSSION***

Figure 3 illustrates the average values for the various workload categories (out of 20). These data were analysed with within-groups T-tests. Overall subjective workload (the last category in Figure 3) was significantly reduced in the pitch condition ( $T_{15} = 4.28$ ,  $p < 0.01$ ). A more detailed analysis showed a significant decrease in all of the workload categories for the pitch condition: mental ( $T_{15} = 4.46$ ,  $p < 0.01$ ), physical ( $T_{15} = 2.15$ ,  $p < 0.05$ ), temporal ( $T_{15} = 4.00$ ,  $p < 0.01$ ), effort ( $T_{15} = 2.15$ ,  $p < 0.05$ ), performance ( $T_{15} = 3.70$ ,  $p < 0.01$ ) and frustration ( $T_{15} = 2.67$ ,  $p < 0.01$ ). Our first hypothesis is thus confirmed.

[Insert figure 3 about here]

Participants took on average 133 seconds to answer a question in the speech condition (standard deviation = 55) and 44 seconds in the pitch (standard deviation = 14), this difference was statistically significant ( $T_{15} = 6.2$ ,  $p < 0.01$ ). This confirms hypothesis H2. On average participants got 4.6 correct answers in the speech condition and 5.3 in the pitch. This difference was statistically significant ( $T_{15} = 2.2$ ,  $p < 0.05$ ). This confirms hypothesis H3.

The results from the experiment showed that the use of non-speech sounds had a very significant effect on all aspects of the use of the system. They significantly reduced the workload and allowed users to answer more questions correctly more quickly. Using non-speech sounds along with speech in a graph browsing interface will provide significant usability and accessibility benefits for users.

### **Haptic graphs**

Our haptic work is currently about investigating the usefulness of haptic devices for presenting various forms of graphs to blind people. Using force feedback devices has several advantages over conventional methods such as: simpler and more rapid graph creation, flexible modification, and presentation of dynamic, time-varying data. The key question is: how should the graphs be designed so that they are usable. The following experiment is an attempt to evaluate some initial designs.

## DESIGN OF HAPTIC LINE GRAPHS

Our line graphs are constructed by using geometric objects and polygon-based models (based on the GHOST software tools from SensAble Technologies, Inc.). The axes and the background of the line graphs are formed based on primitive geometric cylinder objects. The data lines are modelled by using polygons. An illustration of a line graph model is given in [Insert figure 4 about here]

. In the figure, the PHANToM-controlled pointer moves in a three-dimensional space in which all virtual objects are located on the background. The lines are raised from the surface of the background and form V-shaped channels. By moving the pointer along the channel, users can trace the path of the line thus a mental image of the line can be constructed. Multiple lines can be constructed on the same graph. Users only need to move the pointer out from the top of the channel in order to locate other lines. To handle multiple lines on a graph, friction keys, which change the frictional property of the line surface, are used as a discriminative feature. Therefore, by judging the degree of friction, users can distinguish different lines.

[Insert figure 4 about here]

The main reason for making data lines into a V-shape channel was to retain the pointer on the line effectively. In our initial trials [16] we found that users had problems with keeping the pointer on raised objects, such as cylinders. We hence concluded that a concave shape could effectively solve this problem ([Insert figure 5 about here]

shows two example graphs). There is also a problem at line intersections as users sometimes did not notice the intersections, or alternatively failed to continue along the same line but crossed on to a different one without noticing. We hoped that the concave design would avoid this.

## PILOT STUDY

### *PILOT STUDY DESIGN*

An exploratory pilot study was conducted to validate the design of the haptic line graphs. Fifteen sighted subjects participated, all were undergraduate Computing Science students at the University of Glasgow. Two sets of six graphs were designed: each graph in the first set contained up to three separated lines, whilst those in the second contained two intersected lines. Two sample graphs are given in [Insert figure 5 about here]

. Participants had four minutes to explore each graph. Within the time limit, they needed to perform a set of tasks that included:

- Counting the number of lines in the scene;
- Identifying lines by their friction key;
- Counting the number of bends in the line;
- Counting the number of intersections, if any.

[Insert figure 5 about here]

Moreover, participants needed to remember the relative position of the lines and their location so that they could sketch the graphs after the exploration. This set of tasks tested the accuracy of the information conveyed through the interface. In addition to the above we collected measures similar to the sound graphs experiment above:

- Task completion time;
- NASA Task Load Index workload data;
- A log file containing cursor positions.

The log file was a record of all of the pointer positions (x-y coordinates) each participant moved through during each graph exploration. This allowed us to see how the participants had explored the graphs. The format of the experiment was similar to that in Table 1.

## ***RESULTS FROM THE PILOT STUDY***

[Insert table 2 about here]

summarises the results obtained. Participants scored highly in all of the main aspects of line recognition. Participants in the non-intersected group did well at identifying the number of lines present with 86.67% correct (in the intersected condition there were always two lines present so we expected this score to be high). [Insert table 2 about here]

also shows that a high percentage of correct answers was obtained from the identification of frictional properties. There was no significant difference between the conditions in terms of total time taken to complete

the tasks. Overall subjective workload was significantly higher in the intersected group (non-intersected=10.6 and intersected=12.21,  $T_{14}=3.75$ ,  $p=0.002$ ).

[Insert table 2 about here]

## ***DISCUSSION***

The findings of the pilot study showed that participants' performance was generally very good with many of the main features recognised. The significant difference in the workload scores reveals that participants had to work harder in order to achieve these results in the intersected condition. Among the six categories of workload mental demand and effort received highest ratings at 13.9 and 11.3. This may indicate that sighted people are not used to perceiving information by touch alone. In summary, results obtained in the pilot study are promising as participants can use the haptic interface to get a high percentage of correct answers.

Several problems with the haptic line graph design were identified, however. The friction keys were effective at marking the different lines but they could hinder movement on a line, especially when the line had sharp bends and strong friction. This combination could mislead the participants into misjudging a sharp bend as the end of the line, giving an incomplete image of the graph. [Insert figure 6 about here]

shows an example of this kind of problem. [Insert figure 6 about here]

a shows the graph that was explored. In [Insert figure 6 about here]

b the cursor position log indicates the trace of a participant's movements (darker lines show areas explored the most). The pointer trace follows the path of the darker, higher friction line but stops at the sharp corner half way along. In the participant's sketch ([Insert figure 6 about here]

c), an incomplete picture of the graph is drawn for line 2. This suggests that higher levels of friction can make a participant think that the end of a line has been reached at a sharp corner, rather than at the real line end. In order to overcome this problem, additional cues are required to assist the user.

[Insert figure 6 about here]

## EXPERIMENT

Based on the results and observations from the pilot study, several additional features were planned to improve the haptic interface. However, only two of them were implemented in time for the evaluation experiment due to time constraints. They were: instant line identity and line endpoint indication. The lines were numbered and whenever the pointer touched a line, the PHANToM generated a sequence of clicks (a gentle movement in the z axis) corresponding to the line number. The line end point indication was a continuous vibration that only occurred when the pointer reached the end of a line. It was designed to reduce the confusion between sharp bends and line end points.

### ***EXPERIMENT DESIGN***

The final evaluation of our design was conducted on 15 blind people (recruited from the same sources as the sound graphs experiment above). The graphs from the pilot study were used again but the four simplest graphs were taken out to shorten the length of the experiment. Graphs were arranged in two groups with each having two intersected and two non-intersected graphs. The difference between two conditions in this case was that one group had the enhanced line features whilst the other did not so that we could assess the effectiveness of our graphs in general with blind participants, and our new line identification techniques in particular to see if they made exploration easier. The experiment again followed the design in Table 1. Experiment procedure and measurements were similar to those in the pilot study except that participants were not asked to sketch the graph after explorations.

### ***RESULTS AND DISCUSSION***

[Insert table 3 about here] summarises the results obtained. Again, participants generally did well at identifying the number of lines present and their frictional properties. There were no significant differences in the mean scores between the two conditions. However, there was a large reduction in the variance of the scores in the identification of the number of lines in the enhanced condition. An F-test showed this to be significant ( $F_{14}=3.09$ ,  $p=0.021$ ).

Similarly to the pilot study, mental demand and effort received the highest overall ratings for the subjective workload measurements. Participants rated mental demand at 12.33 for the standard graphs and 10.2 for enhanced ones and this difference was statistically significant ( $T_{14}=2.35$ ,  $p=0.034$ ). In the effort category 11.46 and 10.73 were obtained for the standard and enhanced graphs respectively (with no significant difference).

[Insert table 3 about here]

The scores obtained show that blind participants were able to use the haptic graphs well. This is important as this group reflects the end users of our system so we now know that haptic graphs are appropriate for them. But the participants achieved the same mean scores on the enhanced graphs as the standard ones, which meant that the newly implemented features had little effect. However, the usefulness of the enhancements was reflected in other areas such as a significant reduction in the variance in the number of lines identified, subjective mental demand and users' informal feedback after the experiment. The reduction in variance for the line identification suggests that in general more users were doing better with the line identification technique and it deserves further investigation.

Although no direct comparison can be made because the pilot study and the experiment were different, the scores were roughly the same between the blind and sighted people. This was despite our subject groups being very different. The sighted participants were all computing science students who had very good mathematics skills, whereas there was a wide variation in age, educational background and familiarity to graphs with the blind participants.

## **General discussion**

The results from our experiments with audio and haptic graphs on blind people show that we can present such data very effectively. The sound graphs allowed a far better interaction with a graph than speech alone – all aspects of our usability tests (workload, time to complete tasks and error rates) were improved with the combination of speech and pitch.

Our haptic graphs are in an earlier stage of development than our sound ones so the experiments here were more exploratory and the results less clear-cut. However, the usefulness of the haptic graphs has been confirmed. The blind participants could use the haptic interface successfully and perceive correct information about the graphs. Enhanced navigation features can minimise the differences between individual's capabilities and reduce mental workload. Some problems still exist in the current design, such as the combinational effect of sharp corners and strong friction, and path changes at intersections. Further improvement is needed to solve these problems.

Reductions in mental demand occurred in both experiments and these are very important in this kind of interface. We want our designs to be as easy to use as possible so that users can concentrate their mental

resources on solving the mathematical problems they are working on, and not on using the interface. Our haptic and audio interfaces have succeeded in this regard.

In both cases the way of presenting graphs is new and users had not come across them before. This is likely to mean that users will continue to improve beyond the 1 – 1.5 hours of each of our experiments. A longitudinal study would be important to do in this area as many of our eventual users might use the system for many years as, for example, they go through school and then on to university.

The next stage in our work is to combine these two presentation methods. This will allow users to trace the shape of the graph with a finger whilst hearing pitch change as they move over the graph. A click with the pointer will then speak the value of the current cursor position. This system is in development and will be tested on a range of visually impaired people to ensure its effectiveness when it is finished.

## **Conclusions**

A problem for blind school and university students is access to visualisations such as graphs, tables and data plots. These are very common but can be very hard for people without sight to use. Our research has demonstrated two new methods by which we can overcome this problem. The first use non-speech sounds and the second haptic feedback. Experimental evaluations of these new interfaces have shown that they are effective methods of presenting such visual material in an alternative form. Designers of systems for visually impaired people can now use these new techniques knowing that there is strong evidence to show that they will significantly improve their products.

## **Acknowledgments**

The authors would like to thank GWSSB, the Royal Blind College in Hereford and the RNIB College in Peterborough for their participation in the experiments. This research work is funded by EPSRC Grant GR/M44866, ONCE (Spain) and Virtual Presence Ltd.

## **References**

1. Baddeley, A. *Human Memory: Theory and Practice*. Lawrence Erlbaum Associates, London, 1990.
2. Bevan, N. and Macleod, M. Usability measurement in context. *Behaviour and Information technology* 13, 1 & 2 (1994), 132-145.
3. Blattner, M. and Dannenberg, R.B., Eds. *Multimedia Interface Design*. ACM Press, Addison-Wesley, New York, 1992.

4. Brewster, S.A., Wright, P.C. and Edwards, A.D.N. A detailed investigation into the effectiveness of earcons. In *Proceedings of ICAD'92* (Santa Fe Institute, Santa Fe) Addison-Wesley, 1992, pp. 471-498.
5. Brodlie, K., Carpenter, L., Earnshaw, R., Gallop, J., Hubbard, R., Mumford, A., Osland, C. and Quarendon, P., Eds. *Scientific visualization: techniques and applications*. Springer-Verlag, Berlin, 1992.
6. Colwell, C., Petrie, H., Kornbrot, D., Harwick, A. and Furner, S. Haptic virtual reality for blind computer users. In *Proceedings of ACM ASSETS'98* (Marina del Rey, CA.) ACM Press, 1998, pp. 92-99.
7. Edwards, A.D.N., Pitt, I.J., Brewster, S.A. and Stevens, R.D. Multiple modalities in adapted interfaces. In *Extra-Ordinary Human-Computer Interaction*, Edwards, A.D.N. (Ed.), Cambridge University Press, Cambridge, UK, 1995, 221-243.
8. Hart, S.G. and Wickens, C. Workload assessment and prediction. In *MANPRINT, an approach to systems integration*, Booher, H.R. (Ed.), Van Nostrand Reinhold, New York, 1990, 257-296.
9. Kramer, G., Eds. *Auditory Display*. Addison-Wesley, Reading, MA, 1994.
10. Kurtz, M. Rendering drawings for interactive haptic perception. In *Proceedings of ACM CHI'97* (Atlanta, GA) ACM Press, Addison-Wesley, 1997, pp. 423-430.
11. Mansur, D.L., Blattner, M. and Joy, K. Sound-Graphs: A numerical data analysis method for the blind. *Journal of Medical Systems* 9 (1985), 163-174.
12. Massie, T. and Salisbury, K. The PHANToM Haptic Interface: A Device for Probing Virtual Objects. In *Proceedings of the ASME winter annual meeting, symposium on haptic interfaces for virtual environments and teleoperator systems* (Chicago, IL), 1994.
13. Ramstein, C. and Hayward, V. The Pantograph: A Large Workspace Haptic Device for Multi-Modal Human-Computer Interaction. In *Proceedings of ACM CHI'94* (Boston, MA) ACM Press, Addison-Wesley, 1994, pp. 57-58.
14. Stevens, R.D., Wright, P.C., Edwards, A.D.N. and Brewster, S.A. An audio glance at syntactic structure based on spoken form. In *Proceedings of ICCHP'96: Interdisciplinary aspects on computers for helping people with special needs* (Linz, Austria) R. Oldenbourg, Vienna, 1996, pp. 627-635.
15. Vince, J. *Virtual reality systems*. Addison-Wesley, Wokingham, UK, 1995.

16. Yu, W., Ramloll, R. and Brewster S.A. (2001) Haptic graphs for blind computer users. Brewster, S.A. and Murray-Smith, R. (Eds.). *Haptic Human-Computer Interaction*, pp 41-51, Berlin, Springer LNCS, Vol. 2058, 2001.

## Figure and table labels

Figure 1: A PHANToM device with overlaid arrows showing all possible movements (from SensAble Technologies, Inc.).

Figure 2: The London crime statistics data.

Table 1 Experiment schedule.

Figure 3 Results showing the TLX workload scores (standard error bars are shown).

Figure 4: Haptic line graph model.

Figure 5: Example experiment graphs. (a) Non-intersected graph, (b) Intersected graph.

Table 2: Results from the sighted participants (all figures in %).

Figure 6: Sketch analysis. (a) Reference graph (lighter, lower line has low friction and darker line has higher friction, (b) cursor log, (c) a participant's sketch.

Table 3: Results from the blind participants (all figures in %).

Figures and tables

Figure 1

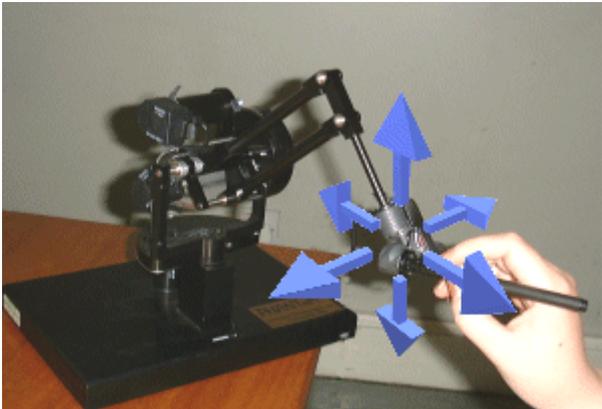


figure 2

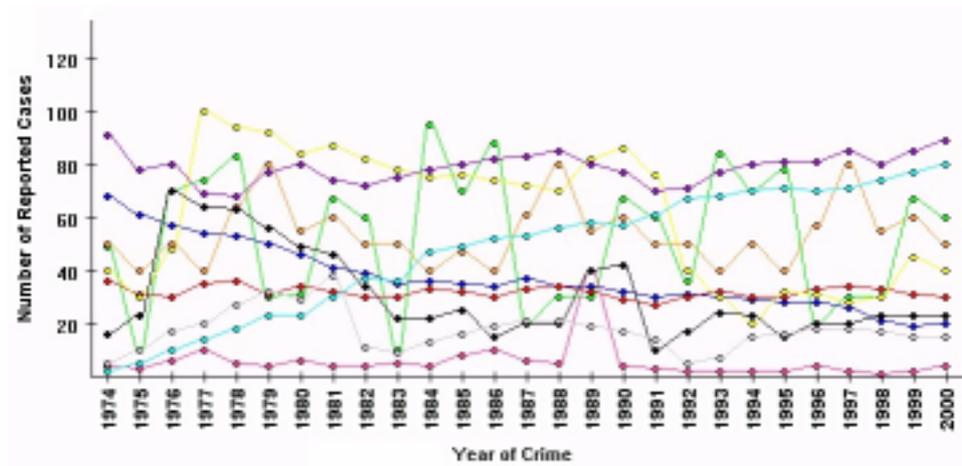


table 1

| Group (4 people each) | Demonstration and Training | Session I                                    | NASA TLX Administration | Demonstration and Training | Session II                                   | NASA TLX Administration | Debriefing |
|-----------------------|----------------------------|--|-------------------------|----------------------------|--|-------------------------|------------|
| A.1                   |                            | <i>(Speech)</i><br><i>Crime Statistics</i>   |                         |                            | <i>(Pitch)</i><br><i>Student Statistics</i>  |                         |            |
| A.2                   |                            | <i>(Speech)</i><br><i>Student Statistics</i> |                         |                            | <i>(Pitch)</i><br><i>Crime Statistics</i>    |                         |            |
| B.1                   |                            | <i>(Pitch)</i><br><i>Crime Statistics</i>    |                         |                            | <i>(Speech)</i><br><i>Student Statistics</i> |                         |            |
| B.2                   |                            | <i>(Pitch)</i><br><i>Student Statistics</i>  |                         |                            | <i>(Speech)</i><br><i>Crime Statistics</i>   |                         |            |

Figure 3

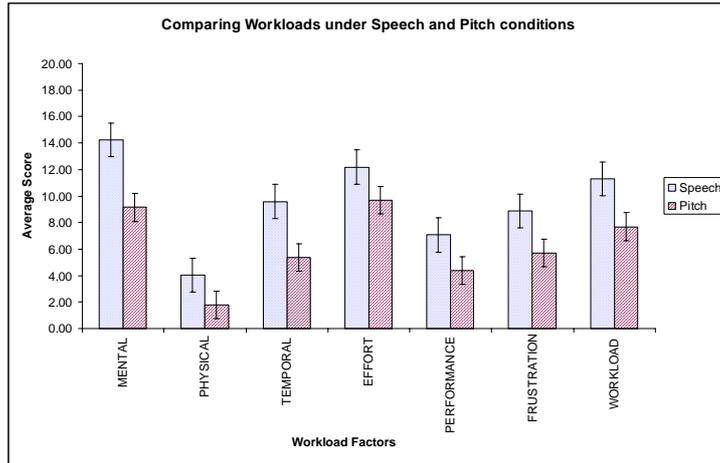


figure 4

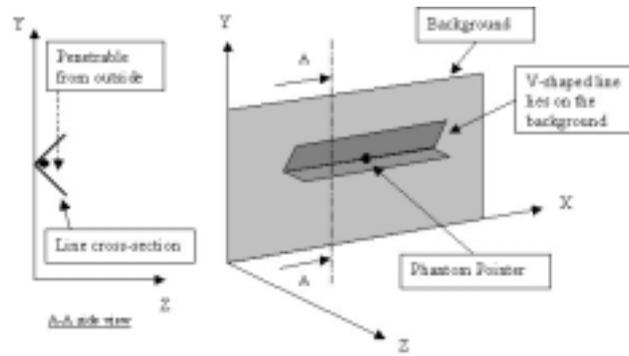
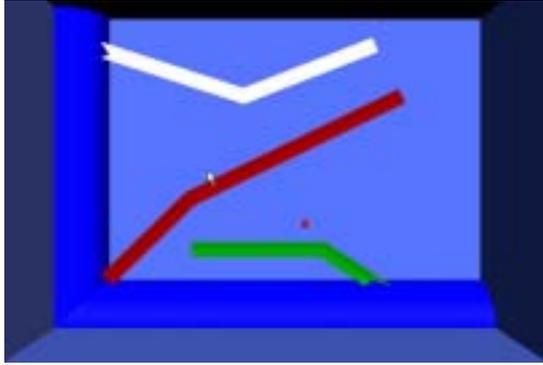
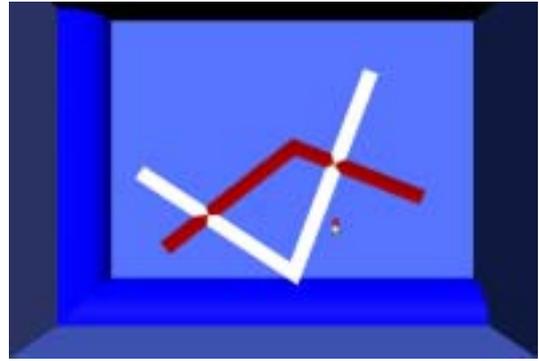


figure 5



(a)

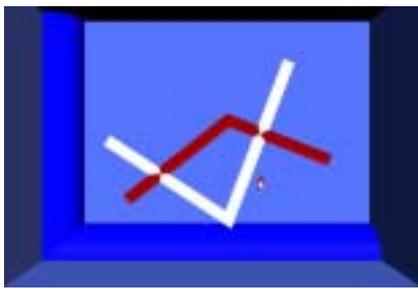


(b)

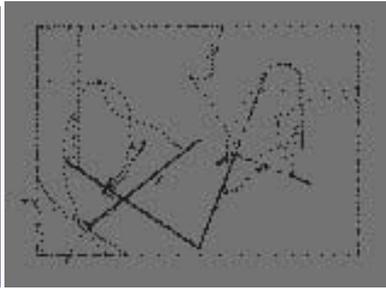
table 2

| Questions            | Non-Intersected |            | Intersected |            |
|----------------------|-----------------|------------|-------------|------------|
|                      | Mean            | Std<br>Dev | Mean        | Std<br>Dev |
| No. of lines         | 86.67           | 15.69      | 96.67       | 9.34       |
| Friction keys        | 94.44           | 6.03       | 98.89       | 4.3        |
| No. of bends         | 76.11           | 16.33      | 88.89       | 15         |
| No. of intersections | N.A.            | N.A.       | 92.22       | 13.9       |

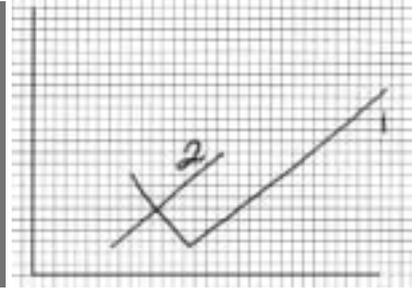
Figure 6



(a)



(b)



(c)

table 3

| Questions            | Standard |            | Enhanced |            |
|----------------------|----------|------------|----------|------------|
|                      | Mean     | Std<br>Dev | Mean     | Std<br>Dev |
| No. of lines         | 86.67    | 28.14      | 88.33    | 16         |
| Friction keys        | 95.56    | 10.11      | 94.81    | 8.26       |
| No. of bends         | 79.26    | 28.44      | 73.33    | 30.23      |
| No. of intersections | 88.33    | 18.58      | 88.33    | 16         |

