IST-2003-511592 STP

MICOLE

Multimodal collaboration environment for inclusion of visually impaired children

Specific targeted research project

Information society technologies

Deliverable D7: Report on the Use of External Memory Aids to Improve Navigation Within Data

Due date of deliverable: 31.01.2007

Actual submission date: 30.01.2007

Start date of project: 1.9.2004 Duration: 36 months

Name of the partner responsible for the deliverable: University of Glasgow

Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)		
Dissemination Level		
PU	Public	**X**
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Summary

This report describes the work done by four different partners (UGLAS, METZ, ULUND and, KTH) in Work Package 2, Senses and Accessibility, and in particular task 2.2 Multimodal control and Navigation, subtask 2.2.2 Developing External Memory Aids. The key aims of research in this area were to look at how we could support people exploring information with the use of *external memory*. This reduces working memory loads and allows the making and comparison and re-finding of points of interest much simpler. This is currently very difficult for visually-impaired users and imposes a high working memory load.

We investigated external memory aids in several different modalities. We looked at force-feedback beacons that could move users to a point of interest, and studied the types of forces that were appropriate. We also looked at force playback. This would allow a user to be dragged back to a point of interest later on, or dragged around a shape that another user has marked. Results showed that multimodal force and audio feedback were the most effective at guiding users.

In addition to the force aspects of haptic feedback, we also investigated tactile displays for presenting external memory aids. Here we looked at how the different parameters of cutaneous perception could be used to encode information into a tactile cue. The results showed that blink rate and shape were most effective, with pattern size causing problems for users.

The final area was the use of multiple forms of feedback for memory aids in an application setting. This was a haptic line drawing application that used speech, gesture/sound feedback and haptic force beacons to support users drawing simple shapes and then finding and navigating around the shapes.

From all of this work we have distilled a set of design guidelines for others to use (WP5) and fed requirements into the MICOLE architecture in WP4.

Table of Contents

1. Introduction

Understanding and manipulating information using visualisations such as graphs, tables, bar charts and 3- dimensional (3D) plots is a very common task for sighted people. The skills needed are learned early in school and then used throughout life, for example, in analysing information, creating presentations to show to others, or for managing home finances. The basic skills needed for creating and manipulating graphs are necessary for all parts of education and employment. Blind people have very restricted access to information presented in these visual ways. It is currently very hard for them to create, manipulate and communicate visualisations such as graphs and tables. As Wies *et al.* state “Inaccessibility of instructional materials, media, and technologies used in science, engineering, and mathematics education severely restricts the ability of students with little or no sight to excel in these disciplines. Curricular barriers deny the world access to this pool of potential talent, and limit individuals’ freedom to pursue technical careers” (Wies *et al.*, 2001). Traditional methods of presenting visualisations to blind and visually impaired people include Braille diagrams, heat-raised paper, screen readers and screen magnifiers. There are several drawbacks inherent with these methods in that they are either unable to respond quickly to dynamic changes in data (hard copies need to be produced of Braille and heat raised diagrams which is often slow and difficult for a blind person without sighted assistance), they are inherently serial in nature and therefore highly memory intensive (e.g. a screen reader reading values from a graph or a Braille table), use an abridged form of the data (pre-recorded descriptions of graphs delivered via spoken word or Braille versions of tables), or are simply inaccessible to potential users (only 26% of visually impaired UK university students read Braille, and screen magnifiers are useless to those with no residual vision).

Previous research has shown that in the absence of visual information, users are able to perceive and interpret multimodal (haptic and audio) representations of common graph types such as line graphs, bar charts and pie charts. Experimental results showed that a multimodal representation of line graphs was significantly more accurate than a raised paper based representation; however, exploration times were significantly slower (Yu and Brewster, 2003). This increase in time was attributed to the point interaction nature of the PHANTOM. Limiting the user to a single point of contact precludes the use of Exploratory Procedures (Lederman and Klatzky, 1987) such as enclosure and contour following that are important for perceiving size and shape of objects efficiently (an essential action for comprehending the data in graphs). The lack of spatially distributed cutaneous information on the finger tip means that the users instead have to integrate a series of temporally varying cues as they traverse the graph. Exploration is therefore slow and highly memory intensive as little context can be provided through a single point of stimulation. These problems are further exacerbated when dealing with large data sets or data exhibiting a high dimensionality.

A fundamental problem faced by blind people when interacting with visualisations (or any complex information) is that there is no easy way to mark points of interest or to access *external memory* (Zhang and Norman, 1994); a sighted user might mark a graph with a pen to indicate an interesting point to return to later, or write something in the margin as a reminder. Such external memory is a very powerful tool for sighted people and can significantly reduce working memory requirements. This is not possible for blind people and means that they may easily get lost in the data, overloaded, and makes it hard for them to mark interesting points in the data. This slows down interaction, increases workload and means that it is more likely that mistakes will be made. As Stevens suggests, providing access to an external memory aid will give very substantial benefits to blind users (Stevens, 1996). This document describes the design and evaluation of external memory aids for blind and visually impaired users accessing complex visualisations.

1.1. Objectives of the Work Package

In WP2 we carry out basic research through empirical experiments and prototypes to find out how to use different senses in user interfaces for visually impaired children. One main objective is to explore the idea of cross-modal equivalence and multi-sensory perception through a series of empirical studies that contrast different representations of information across different modalities (task 2.1). This will allow us to come up with the best methods of presenting information given different disabilities of our users and different technologies that they may have at their disposal.

The second objective is to design, develop and evaluate a range of navigation and control techniques to allow users to explore, navigate and share data, visualisations (such as graphs, tables and charts) and mathematical formulae (Task 2.2). These will use combinations of modalities such as sound and touch to enable seamless interaction with data and with other users. By careful design we will ensure that our tools are usable by blind people themselves and that they can be supported by their carers and teachers. This deliverable report (D7) focuses on the work in this task, deliverable D5 focuses on work in Task 2.1.

A final, but significant objective, is the construction of design guidelines that inform the creation of multimodal presentation and navigation tools. These guidelines will be used extensively in WP4. There will be a continual effort to feed the results gained in this work package into WP4 throughout, but dedicated time will be assigned to formalising the insights gained from this process towards the end of the work package.

For this deliverable we concentrate on work done in subtask 2.2.2 *Developing External Memory Aids.* Here we have been developing tools to support the fundamental problem of *external memory* to mark, find and compare data points. Results from Task 2.2.1 (Basic Navigation and Interaction) set the foundations by giving us a range of tools for users to navigate around our interfaces successfully. Prototypes have been built to assess the usefulness of external memory aids and we have looked at aids presented using force, tactile and audio. Full experimental evaluations have been undertaken with our users to ensure that the tools we produced were usable and solved the intended navigation problems (M5).

We have fed results into work package 4 on software architectures and work package 5 on evaluation and design recommendations.

1.2. Progress

During the year we had three meetings, Lund, November 2005 to discuss the final prototypes for our year one work, Stockholm in February 2006 where we planned the work to be done after the first set of deliverables were submitted, and then Glasgow in August where we checked progress towards work on the objectives. We also ran the HAID workshop in Glasgow after our MICOLE meeting and many partners presented their work to a larger, international audience (www.dcs.gla.ac.uk/~mcgookdk/multivis/workshop.html).

2. External Memory Aids using Force Feedback

Force feedback potentially has an important role in providing external memory aids for visually impaired computer users. A force feedback device like the PHANTOM (from SensAble Technologies) can physically interact with the user and affect their movements in an environment. This could be used for example to direct users towards areas of interest in the computer environment, or to trace paths on a map. However, when using technology that has such an active role in directing the user through an environment it is important that we develop techniques that are safe, useful, usable and improve performance. Here we describe studies working towards this goal. We examine techniques to both guide users to discrete points of interest in the environments and as well as guide them along continuous paths. Here we further describe two systems using a combination of force and audio feedback to provide external memory aids. The force feedback allows us to physically guide the user through the environment while the audio feedback provides additional context information that can be processed by the user in parallel to this guidance.

2.1. Beacons – Guiding the User to Points of Interest

2.1.1. Introduction

With one point haptic interaction in a non-visual setting, it is easy to miss objects or get lost in haptic space (Colwell *et al,* 1998). Some navigational tools have been suggested, such as “magnets”, “crosses” (allowing the user to feel if he or she is aligned with an object) or a “ball” (to feel things from a distance) (Sjöström, 1999). The attractive force in particular has been used and found to be helpful in many circumstances (e.g. Wall *et al.* (2002) and Langdon *et al.* (2002) and is included as a standard tool in the current OpenHaptics software from SensAble). For graph exploration, Roberts *et al.* (2002) and more recently Pokluda and Sochor (2005) presented different versions of guided tours, while Wall and Brewster (2004) tested the use of external memory aids, so called “beacons”, which the users could place on a surface and which then could be activated to drag the user back to this particular location. Text labels have been used extensively to help users obtain an overview of maps Wood *et al.* (2003) or traffic environments, for example Magnusson and Rassmus-Gröhn (2005).

Other suggested ways to help the user with navigation/learning are automatic guiding constraints, referred to as “fixtures”, which have been used for tele-operation, shared control tasks, tracking and training, often in a medical context Prada and Payandeh, 2005), or to have the user cancel forces generated by the haptic device Saga *et al.* (2005).

In our previous work on navigational tools (Magnusson and Rassmus-Gröhn, 2005, and Magnusson, Danielsson and Rassmus-Gröhn, 2006) we had confirmed the results obtained by Wall & Brewster [7] which was that attractive forces can be useful for helping users to locate targets when using the PHANToM. We had also noted that the type of attractive force used could influence the results and we decided to perform a small study to compare different types of attractive forces. The results of this study are presented below.

During this work we have had discussions with Steve Wall & Johan Kildal at UGLAS, and we are grateful for their comments & suggestions.

2.1.2. Details

The current study is motivated by the fact that to make effective attractive force beacons for a PHANToM environment, one needs to know more about how different users are able to work with different types of attractive forces. Thus we decided to test forces that increased towards the target, forces that were kept constant over distance and forces that increased as you move away from the target. Our previous results indicated that an 1/r force probably was too strong at close distances, while results from Wall et al indicated that the linear force produced too strong forces at longer distances. Because of this we decided to keep these two types as “endpoints” of our scale and add other forces in between. The six different radial dependencies we used were: constant force, tanh(r), 1/r, 1/sqrt(r), sqrt(r) and r. To avoid vibrations on the beacon position, the forces that did not tend to zero at small distances had a short linear part for very small distances (inside a radius of 0.006m). This linear part was attached so that the force function was continuous throughout the whole space (although the derivative would be discontinuous at the breakpoint). The forces were adjusted by hand to feel roughly the same at medium distance (0.05 m). The radial dependence of the forces is illustrated in Figure 1. The forces were always directed along $\hat{e}_r$ towards the beacon.

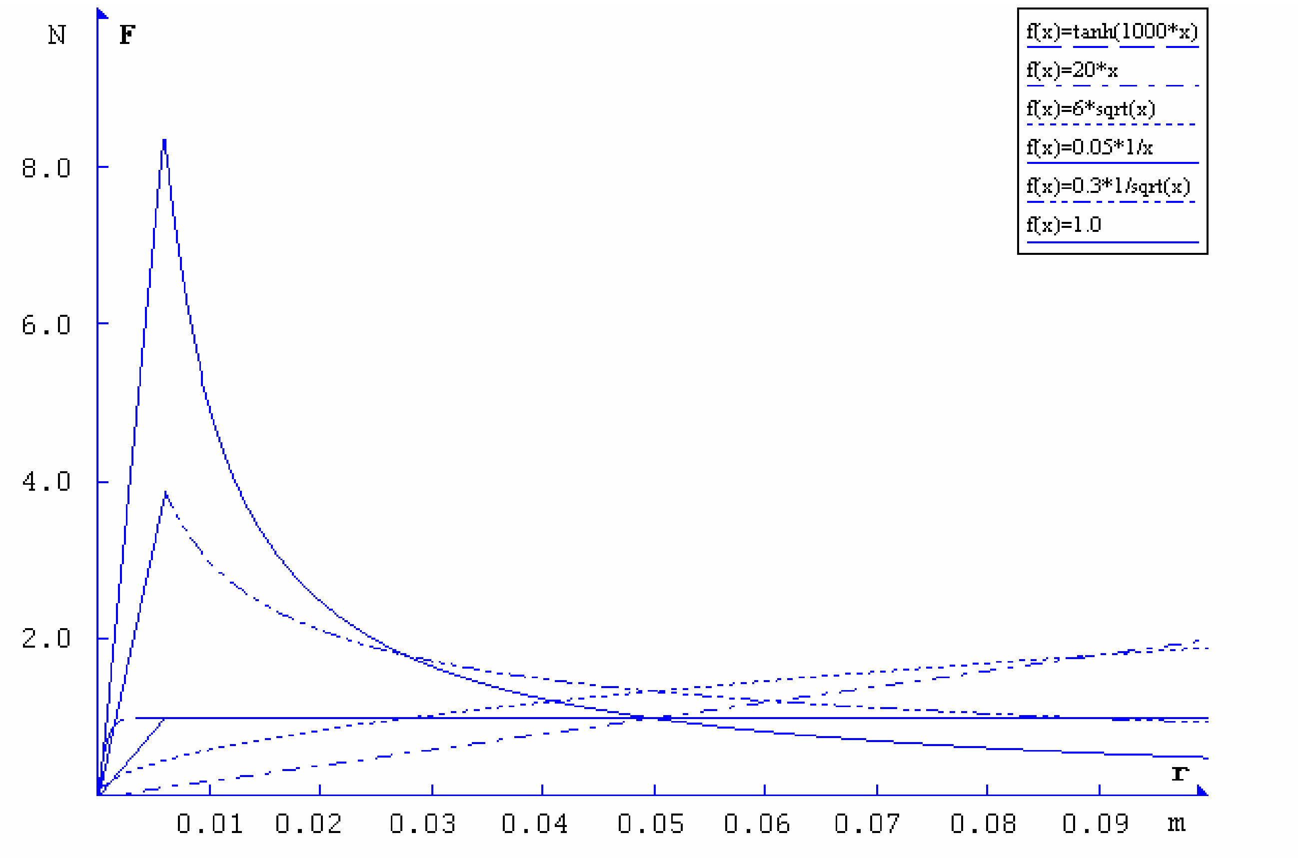

Figure 1. The radial dependence of the different forces used. The tanh and constant forces differ only at very small distances.

Description of the interface and any relevant equipment

For this test the PHANToM 1.0 premium model was used, since it has more precise force rendering. It should be noted that the strength of the forces needs to be adjusted if another PHANToM model is used. The test environment was the Haptic drawing program developed within MICOLE.

Description of external memory aids

This study concerns the design of force beacons. Since the force design will affect user experiences and actions, we have focused on this aspect of the force beacons in this test.

Description of user interactions/Screenshots/Experiment methodology

The test consisted of 5 tasks. In the first the user was asked to rate how well they liked the different forces as they were being held to a point in space by the force. In the second task the user was guided to three different points quite far apart by use of the forces and the user was asked to rate the forces for this type of task. The user initiated beacon changes himself/herself. The task three was the same as the task 2 apart from the fact that the test leader initiated the change of beacon. In task four the user was guided between two nearby points by the use of the forces and rated the forces also for this case. The user initiated a change of beacon himself/herself. In the fifth and final task the user was asked to close his/her eyes and to draw a circle starting from the beacon point and then using the beacon force to close the circle, and then to rate the forces for this type of task.

The order the forces was presented was the same for all tasks for one user, but the order was changed between users to avoid learning effects. Figure 2 below shows the points used for the five different tasks.

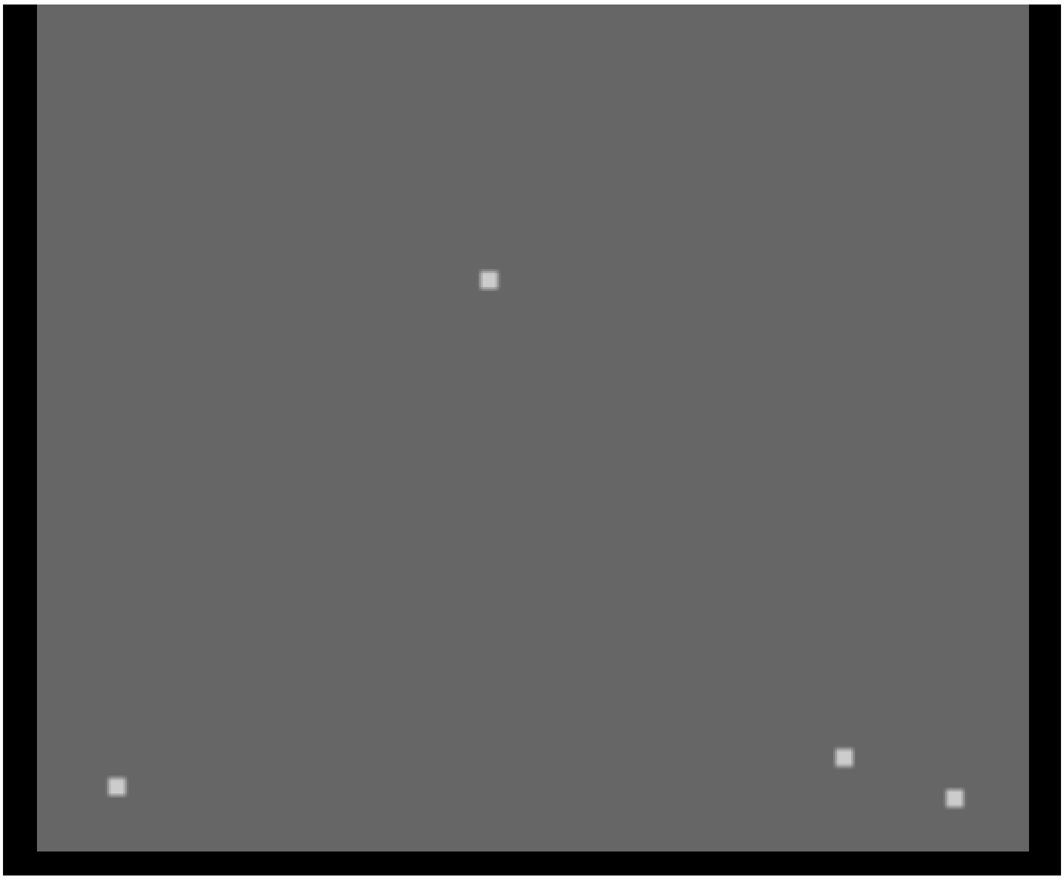

Figure 2. The points used for the test. The top point was used for tasks 1 and 5 while the three points closest to the borders were used for tasks 2 and 3. The two points in the corner was used for task 4.

Participants

Fourteen participants between the ages of 10 and 73 did this test. Based on the assumption that this kind of basic interaction will provide reasonably similar results for blind and sighted participants, due to the limited availability of blind test persons we did this test with sighted users.

Hypotheses

The research question here is basically to find out which of these types of force designs that work well for attractive force beacons, but also to gain some general insights on the design of such forces.

Evaluation results

The main result of the test was the user ratings. These are shown in Figure 3.

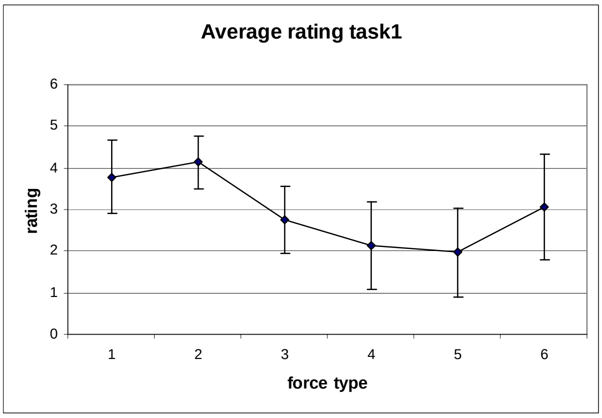

1.

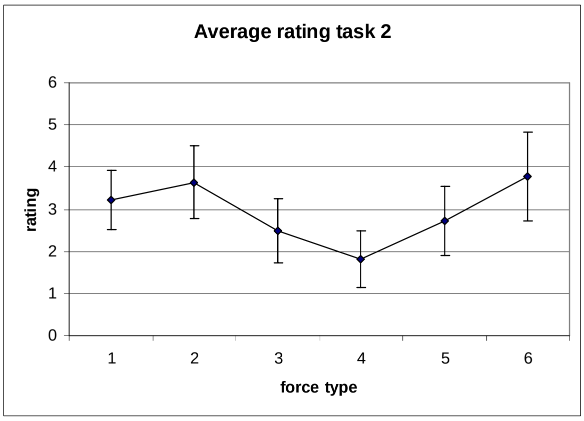

1.

1.

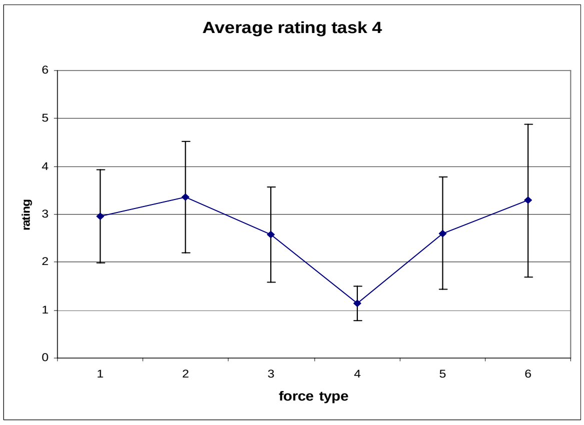

2.

3.

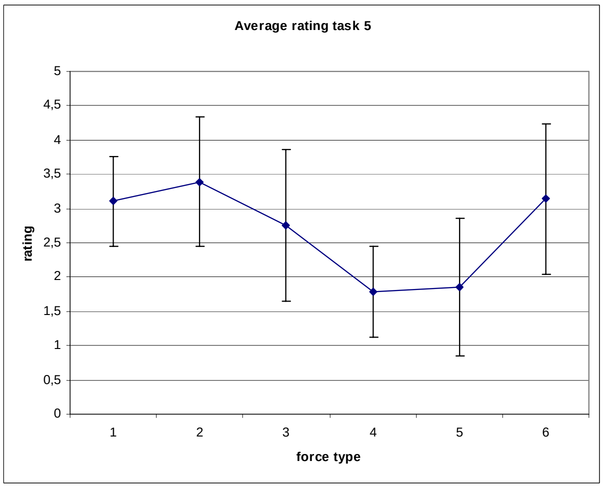

1.

Figure 3. User ratings on a scale 1-5 (1 bad, 5 good) for the different force types: 1. constant force, 2. tanh(r) type force, 3. sqrt(r) type force, 4. linear force (r type force) , 5. 1/r type force, 6. 1/sqrt(r) type force.

Added to this was the observation that the forces that increased with distance (r & sqrt(r)) did not always manage to actually guide the user exactly to the point they were supposed to reach.

To add further illustration of the way the different forces affect a drawing Figure 4 shows screen shots of an experienced PHANToM user trying to draw a circle from a point (using the force as a reference to close the circle) and trying to draw a triangle (using the force to drag the PHANToM between the points). The drawings were done without visual feedback.

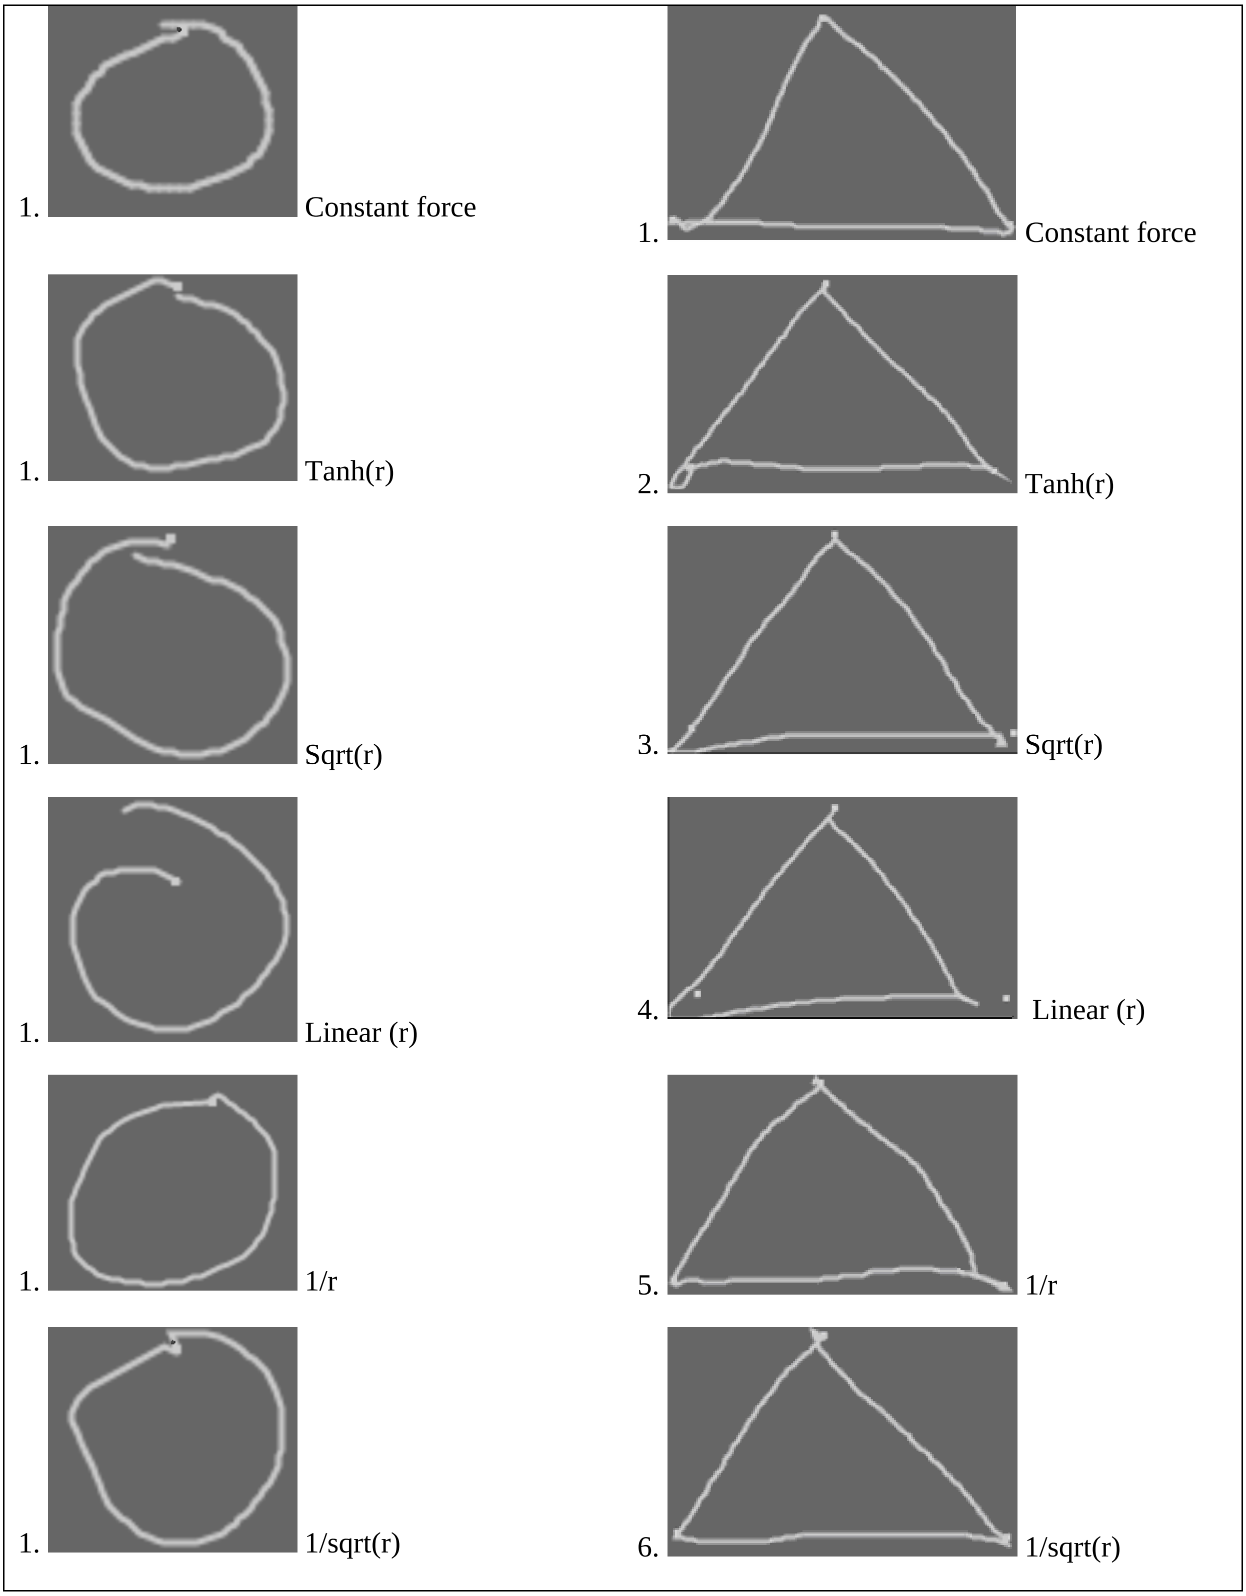

Figure 4. Drawing results for the different types of forces. For the triangles for sqrt(r) and r the user hit the walls restricting the drawing area (that is why it looks as if some of the image is outside the picture). For the triangle the top point was the starting point.

2.1.3. Discussion

If we summarize the results it is seen that the constant, tanh(r) and 1/sqrt(r) are the ones that are most highly rated by the users. The may be an indication that the 1/sqrt(r) type force is slightly better liked compared to the constant/tanh type force for the tasks where the user is more passively dragged around, in contrast to these where the user moves around under the influence of the force, but user preferences vary and it is hard to say anything definite on this point.

What is clearly seen, however, is that the linear force does badly – particularly at short distance. This could of course have been adjusted to some extent by increasing the slope of this force, but then it would have been too strong at large distances.

It is interesting to note that users are able to handle much larger forces at the end of the movement compared to the strength tolerated initially. Even the 1/r type force (which most users felt snapped too hard to the target) was possible to interact with, while the linear and sqrt type forces generated initially for the longer distances created a lot of problems and many users actually dropped the pen when the force started to pull.

Figure 4 further illustrate some further features of the different forces. We can see that the constant, tanh(r), 1/r and 1/sqrt(r) type forces all are able to take the user to the required points. Looking closely at the triangle for the 1/r type one can see that this force is a bit too weak initially (in the right hand corner it takes some time before the user gets pulled towards the target, which can be seen by the two lines joining). The r and sqrt(r) forces on the other hand do worse on this point – the user ends up outside the points in several cases.

During the test we also noted that users needed different strengths of the forces, which indicate that it should be possible for the user to adjust the strength of the force.

Finally we did not note any real difference between the smooth tanh(r) force and the non-smooth constant force.

The advantages and disadvantages of the different forces are summarized in table 1.

1. Force type	1. Advantages	1. Disadvantages
1. Constant	2. Easy to predict. Takes the user to the targets while not interfering too much with user exploration.	2. Does not snap to the targets which can cause overshooting before the user actually gets to the target.
1. Tanh(r)	3. Same as above	3. Same as above
1. Sqrt(r)	4. Does not interfere so much with user exploration.	4. Does not reliably take the user to the targets. Increasing the strength to improve this makes the force to strong at large distances.
1. Linear	5. Same as above.	5. Same as above. Does not work at all at short distances while being almost too strong at large distances.
1. 1/r	6. Gets the user reliably to the points. Outside the vicinity of the point it does not interfere so much with user exploration.	6. Too weak at large distances. Very hard to pull free from a point.
1. 1/sqrt(r)	7. Gets the user reliably to the points. Outside the vicinity of the point it does not interfere so much with user exploration.	7. Some users thought it was a bit hard to pull free from a point.

Table 1. Advantages and disadvantages of the different force types.

2.1.4. Conclusions

To summarize the above, we can say that for the type of tasks studied, users like forces which do not interfere too much with exploration and that, depending on the task, some short distance snap-to-point behaviour is useful.

We also note that the users are able to handle significantly larger forces at the end of a movement compared to in the beginning.

Thus we recommend constant type forces, possibly with some 1/sqrt(r) like snapping behaviour in the vicinity of a point to make sure you actually reach it. This snapping has to be weighed against the possible disadvantage of interfering with user exploration close to the points. We also recommend that the user should be able to adjust the strength of this type of force.

2.1.5. Guidelines

We recommend constant type forces, possibly with some 1/sqrt(r) like snapping behaviour in the vicinity of a point to make sure you actually reach it. This snapping has to be weighed against the possible disadvantage of interfering with user exploration close to the points.

On a more general level we can say that the initial force must be possible to feel, but it cannot be too strong. For short distances some increase in the force is acceptable – but this increase should not be too strong. These considerations point towards the type of force recommended above.

Users need different strengths of the forces, so it should be possible to adjust the strength of the force individually.

Finally it should be pointed out that to avoid vibration the force needs to be continuous in towards r=0, but that it is ok to use a force curve that is not smooth everywhere (at least for the tasks in the present test).

Future developments

As a result of the above tests we will implement the recommended type force with adjustable strength in the drawing program.

2.2. Audio-haptic line drawing application

2.2.1. Introduction

The purpose of this application is to allow visually impaired users to create and access graphical images. The application is and will be developed in close collaboration with a user reference group of five blind/low vision school children. The objective of the prototype is twofold. During the early development stages, it has foremost been used as a research vehicle to investigate user interaction techniques and do basic research on navigation strategies and help tools. In the later phase, during end-evaluations the prototype will be tailor-made for use in schoolwork and the final application should be possible to use in different school subjects.

Getting access to 2D graphics is still a large problem for users that are severely visually impaired. Using a haptic display in combination with audio feedback is one way to enable access. There are many issues to address, e.g. how to provide an overview, to what extent users are able to interpret a combination of lines or line segments into a complex image, how to design the lines to get appropriate haptic feedback, what hardware to use etc. General guidelines to create and develop haptic applications and models are collected in (Sjöström, 2002). Applications making practical use of non-spoken audio and force-feedback haptics for visually impaired people are e.g. applications supporting mathematical display (Yu & Brewster, 2002), (Yu et al., 2003) & (Bussell, 2006) , games (e.g. (Iglesias et al., 2004), (Magnusson et al., 2002), and (Magnusson & Rassmus-Gröhn, 2005)) and audio-haptic maps (e.g. (Iglesias et al., 2004), and (Magnusson & Rassmus-Gröhn, 2005)).

There are few tools that enable blind people to create computer graphics. As described in (Kennedy, 1993) and (*Art Beyond Sight Art Education for the Blind*, 2003), there are indeed people who are blind who have an interest in hand drawing. In (Kamel, 2003), a CAD application is presented that enables users to create drawings with the help of audio and keyboard. This is accomplished by a structured approach of dividing a drawing into small parts and to enable the user to draw small segments of a drawing. In (Hansson, 2003), a study on a haptic drawing and painting program is presented, that is a pre-study to the presented work.

2.2.2. How This Fits Into the Deliverable

The work on the audio haptic line drawing has investigated external memory aids in the shape of speech information, gesture-sound feedback and haptic force beacons. The speech feedback provide information which help users identify parts of the drawing, while the gesture-sound feedback has been used both as a navigational tool (continuous feedback) and as more iconic sounds as responses to specific user actions. The design of the forces to be used in haptic beacons was studied in a separate study. These force beacons are used both "statically" as a means to get back to specific points, and dynamically to allow a mouse user to guide a PHANToM user.

2.2.3. Interface and Equipment

The prototype makes it possible to make black & white relief drawings. The Reachin 4.1 software is used to control the haptic device, which can be either a PHANToM OMNI or a PHANToM Premium. The sound effect playback is based on the FMod API, and the Microsoft SAPI 5.1 text-to-speech API is used to make the application talk. A mouse can be used simultaneously to the PHANToM to enable collaborative use, although the application can be used by a single user also.

The application consists of a room with a virtual paper sheet, which both users can draw a relief on. The PHANToM user can also feel the drawn relief (or an imported pre-drawn relief). The virtual paper can be inscribed in a limiting box, however, in some tests (described below) the limiting box has been removed to provide a larger work area. When the PHANToM pen is in touch with the virtual paper the user draws on it while pressing the PHANToM switch. The mouse user draws while pressing the left mouse button. The haptic image is produced as positive or negative relief depending on which alternative is selected. The relief height (depth) is 4 mm. The drawing can be seen on the screen as a grayscale image – a positive relief is seen as black, and a negative relief is seen as white. The paper color is grey.

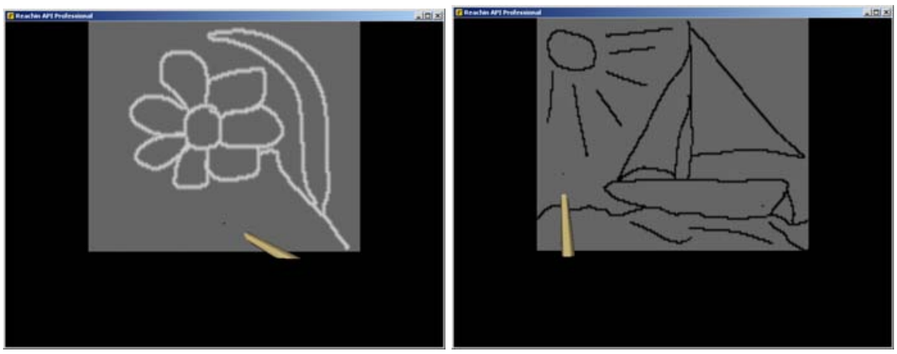

Figure 5. Screenshots of the drawing program in negative relief to the left and in positive relief to the right.

While drawing, the PHANToM user can feel the other lines that have been drawn previously and thus connect parts of lines and feel intersections. Furthermore, each line (or single object) is attached with a spoken number and text tag. This number and text tag is spoken by the application each time a user selects an object by touching it with the PHANToM pen or hovers over it with the mouse cursor.

There are a number of tools that can be of help to the users while drawing. The objects can be manipulated in the following ways:

- Moving
- Resizing
- Copying
- Pasting
- Deleting
- Changing text tags
- Converting to outline shapes: rectangles/squares, ellipses/circles, straight lines

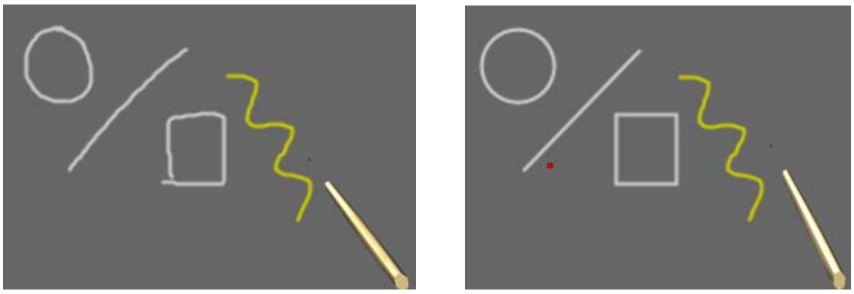

Figure 6. Screenshots showing the outline shape tool. The user has drawn freehand shapes and converted them into a perfect circle, a rectangle and a straight line. The curve to the right is selected, which is indicated by yellow colour.

All manipulation tools (except tagging, so far) are fitted with a feedback sound, and auditory icon that is supposed to help the user understand the progress of the manipulation. The auditory icon is designed to resemble a real world manipulation of similar nature. E.g. the copy function sound effect is a camera click.

Furthermore, the mouse user can guide the PHANToM user by a pulling force that drags the PHANToM pen tip to the mouse cursor position. This also enables the mouse user to point and reference to objects. The PHANToM user can also place force beacons specific points in the environment, to enable fast and accurate return to the points.

The PHANToM user can also choose to turn on a sound field to aid localization. When the cursor moves in the virtual room, the pitch of a position tone is changed, brighter upwards, and mellower downwards. The mode information is conveyed by the volume and timbre of the tone. In free space, a pure sine wave is used. When the user is in contact with the virtual drawing paper (not pressing the PHANToM switch) the volume is louder. And when the user is drawing (and thus pressing the PHANToM switch) the tone is changed to a saw-tooth wave. Also, to differentiate the walls of the limiting box from the virtual paper a contact sound is played when the user hits a wall with the PHANToM pen.

A png import function has also been implemented. The files imported must be grayscale and a multiple of 256*256 pixels – and be exactly the same size as the first imported png (currently a 256*256 pixel size image). A complete grayscale is actually translated into different relief heights, which makes it possible to import any grayscale image and get some haptic information from it. Images not adapted to haptic/tactile reading for blind users are very hard to understand, however, the grayscale can also be used e.g. to smooth out relief lines.

A drawing can be saved as a png file. This, however, does not make it possible to save the drawings as consisting of separate objects that can be manipulated, and it also does not save text tags on drawn objects.

2.2.4. User evaluations

The application has been gradually developed since 2005. It has been continuously evaluated by a reference group of 5 school children, aged 10 to 16. Two (2) of the participants are blind from birth, and three (3) participants have various forms of low vision. All of them read Braille and are integrated in normal schools in southern Sweden.

In March 2006, the application was tested in a formal pilot test with 11 sighted adults (aged 25 – 72). The users made the test without visual feedback from the screen.

Qualitative evaluations with reference group with children

The reference group has (to date) used the drawing program at 5 different group meetings. Not all of the reference group participants have been present at every meeting. Design work has been iterative and the users have been presented with new features and changes in the prototypes at every meeting. All evaluations have been qualitative. The first two evaluations were of an informal nature, with few and loosely formulated test tasks. Instead an open discussion took place in which children and their parents or other close relations and the researchers discussed topics triggered by the prototypes tested. The third evaluation also incorporated some formal test tasks. During these tests, drawing has been tested with and without audio feedback, with positive and negative relief and with program interaction by virtual haptic buttons and keyboard commands. We have used the PHANToM Omni and Premium models both. The test tasks for the 3 first occasions are summarized below:

- Draw and feel lines in negative and positive relief (haptics).
- Use vertical or horizontal work area.
- Change the relief using virtual button or keyboard button (haptics).
- Draw and feel an image (haptics + audio).
- Draw a specified shape – a rectangle and an Arabic number (haptics + audio).
- Explore and recognize Arabic numbers (haptics + audio).
- Explore and recognize two simple geometric shapes (haptics + audio)

At the fourth meeting, the shape drawing tools were used. All three users were guided to complete a specific task. First, the users were helped to learn the manipulation tools, and the commands that went with them. As a final test, the users were asked to use the tools to draw a specified shape - a house – with a rectangular body, two round windows, a rectangular door between the windows and a pointy roof. Since the task quite informal and not specifically directed towards handling of keyboard commands, the users were helped to remember commands when needed. Three of the reference group members participated in this test.

The fifth meeting introduced the collaborative version of the application. The accompanying persons (parent or friend) used the mouse and the visually impaired pupils used the PHANToM. The screen was faced to the users, so that the seeing person could get the visual info. The force beacons were also introduced, but no special task was designed to make them useful, since the tasks were collaborative. Aside from some general experimentation with the application, the users were asked to solve 2 tasks of school work nature. Both tasks were focused on exploring, marking and text-tagging in prepared drawings. One task was to explore a positive relief of a closed figure consisting of angles and connecting straight lines, and to identify, mark and tag the right angles of the figure. The other task was to explore an outline map of the Skåne region, and mark it with 3 points of interest, and tag them with appropriate info. 4 users participated in this meeting and worked with these test tasks. Two of the test sessions were recorded on audio tape and one with a video camera.

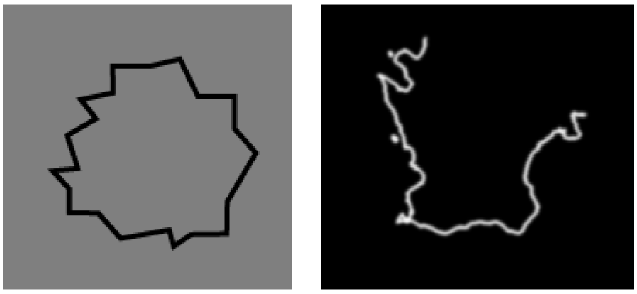

Figure 7. Screenshots showing tasks for collaboration, the left figure for the angle recognition task and the right figure for the map task.

Formal pilot study with sighted adults

A formal pilot user test was conducted on the prototype version to study image recognition tasks concerning geometrical figures, curves and road signs (Figure 8) and to find out how different images, sound/no sound and positive/negative relief influenced the recognition performance. The functions of the prototype version included sound field mapping to pitch and pan, vertical drawing paper, limiting box and png file save. The PHANToM Premium was used for this test, and a pair of headphones was used when sound field feedback was available.

Figure 8. Sample recognition task material. Two geometrical figures and two curves (one open and one closed) as well as two road sign examples.

Eleven (11) users were presented with different recognition tasks and were asked to describe or reproduce the experienced relief using verbal description (for geometrical figures), the drawing program itself (for curves) or pencil and paper (for road signs). The images, sound and relief parameters were varied in four different test cases to overcome the learning effect bias in the test.

The users were asked to rate the difficulty of the recognition tasks on a scale from 1 to 5, where 1 means least difficult, and 5 means most difficult. For each task the time to examine figures was measured. When the user considered himself/herself ready with the examination the time was stopped and the user was free to use as much time as he/she needed to reproduce the task figures. For the road sign recognition, the user was also asked to point out the test task road signs among a collection of road signs with similar features.

2.2.5. Results

Results from qualitative evaluations

All users seem to enjoy the program, and have found it easy both to draw and feel the lines. For line following negative relief seems to be preferred, although one user expressed a liking for positive relief. In general both types of relief seem to be wanted. Our first test was done using the PHANToM Omni model, while later tests due to practical reasons were done using a PHANToM Premium. This caused problems since the Premium pen is less easy to hold (particularly for blind children who are not as used to holding a pen as sighted children are). Furthermore the tiny switch on the Premium was harder to manipulate than the buttons on the Omni – particularly for children with more problems with their fine motor skills. Both vertical and horizontal work areas have been tested. The horizontal work area puts less strain on the arm, and allows for external physical constraints (e.g. the table) to stop the user from pushing through the haptic surface too much. The vertical work area on the other hand seems to generate more distinct haptic feedback – uses express that shape differences may be easier to feel with this orientation.

There was no clear preference for keyboard buttons over virtual buttons. The advantage of the keyboard is that it can be accessed without moving the PHANToM, but our users have to remember a lot of keyboard commands already, and thus keyboard use may lead to an unwanted increase of the memory workload.

For the tests with both haptics and audio, all three test users were able to use the application as intended, and the different task results for the users seem to match personal differences in fine motor skills and the ability to remember images. Two of the three users could draw a square and a number, two could recognize the Arabic numbers and all three could recognize the simple geometric shapes (circle/triangle/square). During the test session some general observations were also made. It seemed as if some of the users were helped by the difference in sound character to know which program mode they were in. This helped especially one user who previously had had big problems releasing the switch to feel the painting. The sounds, however artificial, did not disturb the users very much, although one musically trained user found them disturbing. That same user also indicated the similarity of the sound with the aiming sound used for target shooting for blind users. Another user expressed great enjoyment with the sound and spent quite much time playing with it.

In the test on the drawing tools, no sound field audio was present, although the previous test indicated that some of the users were helped to understand the program mode. Partly, this was because the sound field could interfere with the auditory icons for the manipulation feedback. All three users completed the test with no apparent use problems other than remembering commands. The tools worked as expected, although a performance bug was discovered when using an older computer during the test. The users drew one house each. Two houses ended up as intended, one got a square roof and the door in the wrong place. The 2 users that managed to solve the task as intended both commented on the problem to find the corner of the house and connect the roof to it.

Figure 9. The house drawings from the three participants in the user trials. The two houses on the left look as intended from the instructions given. The leftmost house also has a chimney on top of the roof. The rightmost has placed the door on the left side instead of the middle and the roof is a rectangle. However, that same user has used the limiting box in a creative way to place the house on the "ground". The coloured rectangles in the drawings indicate the object currently selected.

All participants also commented on the audio feedback for the manipulation tools. They found the application noticeably more interesting than in previous meetings and said "Great fun!" and "Cool!" when they heard the sounds playing. At the particular test session, no tools for the creating of straight lines was implemented, but it was asked for and therefore realized shortly afterwards. More tools were also asked for, e.g. tools for making triangles or stars.

The results of the fifth test with the collaborative environment are very diverse. Although all pairs of users can be said to have succeeded in solving the tasks at hand, marking and also using the text input to label points of interest in the map and the right angles in the angle task.

Figure 10. Two pupils collaborating in the map marking task.

In most pairs, the visually impaired user was the one that marked out the points on the map and the angles in the figure. The PHANToM is a lot easier to draw with than a mouse, but there is also the issue of speed – it is simply much faster to obtain an overview of the scene visually compared to using the PHANToM. Thus, the sighted participant guided the other user in all pairs – sometimes verbally, and sometimes with the dragging force. Especially one user, who, despite continuous test sessions with the PHANToM has not learned to use it efficiently, was guided with the force by a parent. This actually made the child understand better how to move with the PHANToM in order to get good feedback. In the pairs with children and parents together (3 of the pairs), the sighted parent was the one who lead the work and prompted the child to do things. One pair was a child (12 yrs) and a friend of the same age, and the lead was not so clearly taken by the sighted pupil. However, it seemed that the particular visually impaired pupil relied very much on the friend, and sat passive with the PHANToM in hand waiting for things to happen. Therefore, the test leader had to prompt the visually impaired pupil to work actively.

Results from formal pilot study

The formal test took approximately one hour per user to complete. 10 out of 11 test persons were able to complete the test, although a few test tasks had to be cancelled due to time restraints.

In general, the geometry recognition tasks were found to be the easiest. The time to complete the examination of the geometrical figures was shorter than the time to examine curves and road signs (table 1).

Image	Mean estimated difficulty (1 to 5)	Mean examination time (seconds)
Geometrical figures	2	35
Curves	3	100
Road signs	4	191

Table 1. Estimated difficulty (1 is least difficult) and examination time for different images.

The geometrical figures recognition task also indicated that negative relief is to be preferred over positive relief, both subjectively and by time measure. However, there appears to be no significant difference between the presence and the absence of the sound field. Results are presented in more detail in (Rassmus-Gröhn et al., 2006).

Since the task did not incorporate any reproduction of the geometrical figures by drawing, there is no information on how users perceived the figures in exact detail. However, 3 out of 10 users described the pentagon as being “a house side” or “a square with a triangle on top”. One user mistook the octagon (8 sides) for a heptagon (7 sides), another user for a nonagon (9 sides). Two users also mistook the hexagon for a pentagon.

For the curve recognition tasks, there appears to be no major difference between the different sound and relief conditions neither concerning subjective measures nor time measures.

A majority (8) of the ten test users drew more than half of the figures reasonably correct. The most common type of major error is that part of the figure is missing. This occurs particularly when the figure contains separate parts or when there is an intersection inside the figure. But also in the case of a single line, this may occur when the figure contains sharp direction changes such as the curve illustrated in Figure 11.

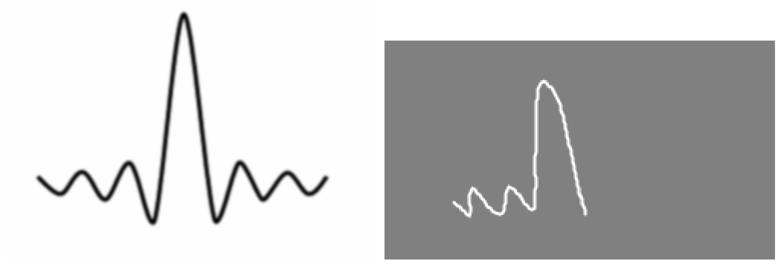

Figure 11. Curve recognition task: the original curve and one example of the curve drawn by a user.

The drawings also illustrate the need for better feedback during drawing, since minor mistakes such with regards to exact/relative positions and shapes were quite common – most users drew the figures from memory (as if drawing in the air) and would easily lose their orientation within the virtual environment.

The road sign recognition test was considered even more difficult. However, despite the obvious problems the users had to examine and to reproduce most of the signs with pencil and paper, the users still pointed out the correct road sign on average 3 out of 4 times when presented with a chart of 24 different road signs.

It is hard to extract any information about benefits of choosing either one of negative or positive relief. Since it was a whole area that was embossed (positive or negative) the scanning of the area was difficult in either mode. Nevertheless, some observations were made that indicate that negative embossment is easier to scan because it more clearly constrains the user to an area.

2.2.6. Discussion

The application has been gradually developed during a long time, letting users repeatedly take part in evaluating newer versions of the program that have incorporated necessary changes, and changes asked for by the users in each new step. It works quite well, although it seems that the youngest members of the reference group would like to see faster progress to keep up enthusiasm. The later prototypes have understandably rendered more spontaneous positive feedback, especially concerning the addition of drawing tools, sound effects and labeling features. The application is quite stable, having undergone gradual efficiency enhancements. Stable and reliable application functionality is especially important when users are visually impaired, and many problems have been discovered during reference group meetings.

Concerning the collaborative use, it seems clear that it is quite hard to design tasks where the users are "equal" - the sighted users were quite dominant in solving tasks in all pairs. Among other things there is a speed difference - vision is often a lot faster than haptic/audio exploration. The test tasks in the collaborative test were considered interesting and it was evident that the parents more clearly saw the use of the system. It has previously been somewhat hard to explain the goal and the usefulness of it.

The formal pilot study has shown that there is a tendency towards the preference for negative relief when it comes to following lines for the simpler line based drawing (e.g. geometrical figures). Also the time measurements indicate that negative relief is easier to use. This approach is also used by e.g. Yu et al. in (Yu et al., 2001). It seems, however, that this effect is less obvious in recognition of more complex line drawings. The study in (Sjöström et al., 2003) also shows that individual preferences for relief vary. The examination time results for the positive relief rely much on the strategy adopted by the user. If a user follows the inside of a closed figure, recognition tends to be found easier and object examination times shorter.

The error made by 3 of the users, who mistook the pentagon with sloped sides for a figure with straight vertical lines is also exemplified by Riedel and Burton in (Riedel & Burton, 2002). The sound field present did not give users any help in determining the slope either, since the stereo panning of the sound has too low resolution.

The sound information was shown not to affect the examination times and recognition accuracy. Pakkanen and Raisamo (Pakkanen & Raisamo, 2005) have previously shown that exact recognition of geometrical objects using a combination of vibrotactile feedback and audio is hard.

Some users also expressed annoyance with the sound, whereas some users enjoyed it despite its artificial sound. One user suggested that the sound feedback should convey information about the placement of the center of the paper rather than the height of the PHANToM pen. Another user suggested that the sound information adjusted with a larger pitch range and better stereo effect might give information about the size of objects or relative shape of similar objects (like a sphere and an ellipse for example).

On two occasions, the absence of sound field feedback did have an impact on a single user's result. Due to technical problems, the contact sound with the walls stopped working after a while, which affected the examination times in the test cases without sound field feedback, since the user mistook the edges along the limiting walls for lines. With the sound field feedback present, the limiting wall contact sound was not as crucial.

2.2.7. Conclusions

- Both positive and negative relief is possible to feel and to work with.
- Negative relief is preferred when working with simple line shapes.
- There are indications that negative relief shortens examination times.
- Both vertical and horizontal work areas can be used.
- Simple shapes can be recognized when they are kept in a specific context.
- The sound feedback can be used to get information about the program mode.
- The sound feedback needs to be improved to provide additional information.
- The PHANToM Premium is hard to use especially for blind users who also don't handle an ordinary pencil very easily.
- The drawing tools implemented are considered useful, but more tools are asked for.
- The sound effects for the drawing tools are considered to add to the experience.
- The collaborative test tasks were interesting, but hard to design in such a way as to make the sighted person and the visually impaired person equal.
- Collaborative drawing with one mouse and one PHANToM is technically possible but more studies on e.g. the application feedback and the collaboration benefits need to be done.

2.3. *Trajectory Playback – Guiding the User along a Path*

2.3.1. Introduction

In deliverable D6, we describe a study testing the use of haptic trajectory playback to learn shapes or gestures. Users were dragged around a trajectory and asked to recreate that trajectory. Here, we describe a second trajectory playback experiment that was designed to build on the findings of the previous study. It was clear from the initial study that the task was particularly difficult for some visually impaired participants. The goal of this follow-up study was to test whether combining the haptic trajectory playback with complementary information from another modality could improve the transmission of the shape information to the user. In this case, audio was combined with the haptic modality.

2.3.2. How this Fits with the Deliverable

Here, we examine non-visual methods to guide users through a trajectory or series of points. This technique can be useful for many reasons, external memory aids being one of these. In some circumstances it is important for a user to not only to navigate to points of interest in the environment, but also to a series of points or along a trajectory in the environment. For example, allowing users to follow a path on a map, or working out the relative positioning of objects by guiding a user through a path between all the objects. This section describes an experiment to evaluate the advantages of a multimodal as opposed to a unimodal solution to this problem.

2.3.3. Background information and relevant literature

Haptic Cueing for Teaching Shape Information

Training users to move through a path non-visually has been examined previously in the literature. Noble and Martin (2006) present a directional cueing study that examines a series of tactile guidance cues designed to guide a blind computer user around a shape. The user interacts with the VTPlayer tactile mouse with a series of eight tactile patterns indicating which direction the user should move in next. Similarly, Crossan and Brewster (2006) used tactile directional cues to guide a user through a complex path represented as a maze. Blind and visually impaired users interacted in a two-handed manner using a PHANTOM device (from SensAble Technologies) in the dominant hand to move around the environment while receiving tactile directional cues on the left hand to guide the movements. There are also many examples where force-feedback is used to guide users around a trajectory.

Feygin *et al.* (2002) conducted a study into the possibility of providing training in moving along a trajectory using either visual only, haptic only, or visual-haptic guidance. Further to this, there were 2 conditions in which participants recalled the trajectories, which were haptic-visual (where the participants saw their cursor as they attempted to perform the trajectory), and haptic only (where the participants attempted the trajectory with no feedback of cursor position). Results showed significant improvement in recreating the trajectory in all conditions between the first and the last attempt. The haptic only training mode performed significantly worse than the haptic-visual training mode, but not significantly worse than the visual only training mode.

Dang *et al.* (2001) discuss a constraint-based surgery training system that provides guidance to users by restricting their movements from deviating from a path. This method allows a user to follow the path taken for a procedure by an expert surgeon, but allows the user to apply the forces to perform the surgery. Yokokohji *et al.* (1996) similarly examine haptic force playback for the purposes of training in a simple task. The system they studied actively dragged a user through the motions required to perform a task to provide training in performing that task. However, no significant training effect was noted here as the task chosen for this study proved as simple as to not require training.

There have been several systems designed to develop skills in calligraphy. Teo *et al.* (2002) demonstrate a system where the position of a teacher can be recorded and played back to a student to aid in forming Chinese calligraphy characters. A separate Chinese calligraphy training system developed by Wang *et al.* (2006) examines the performance of users when drawing characters both with and without force-feedback guiding their actions, with results suggesting that users were more accurate when drawing with the force-feedback. Henmi and Yoshikawa (1998) discuss a haptic system to allow learning of Japanese calligraphy. The method used is to record the position and rotation of the teacher's pen along with the forces used and replay this to the student.

Gentry *et al.* (2003) demonstrate a system which allows the user and computer to collaborate on a dancing task. Here, users had to synchronise their moves to music and with the movements of the device. Force-feedback is used to guide the user and improve the synchrony between the user and the computer.

Communication through Gesturing

The key aim of this work is developing a system to allow the computer (or a second user) to convey spatial information to a user. Two notable closely related works are that by Graham and Argyle (1975) and Oakley (2003). Both performed diagram perception studies (with sighted participants only) using a 'describe and draw' paradigm. Graham and Argyle studied empirically the effect that hand gestures had on the transmission of shape information. They examined participants describing complex abstract scenes using verbal description alone or verbal description and hand gesturing. Through independent rating of the closeness of the drawn pictures to the actual pictures, they were able to show significant improvements in communication when hand gesturing was allowed.

Similarly, Oakley (2003) studied the effect of gesturing on transmission of shape information but through a collaborative computer system. Computer mediated trajectory playback was used in this study. Oakley examined three conditions for communication of shape: through verbal discussion and haptic playback alone, through verbal discussion and visual feedback of the other user's cursor, and finally a combined verbal discussion, haptic feedback and visual feedback condition. The haptic only trajectory playback of the image was achieved by hiding both users' cursors during playback. The visual condition examined visual cursor playback combined with verbal description of the image. The participant drawing the image could see the describer's cursor as he/she was describing the image. Finally, a combined haptic, visual and verbal condition was also studied. The results suggested that visual only and visual/haptic conditions produced significantly better drawings than the haptic only condition although no significant differences were noticed between the visual only and the visual/haptic conditions. Oakley's results suggest that additional information to the haptic trajectory playback is important to understanding the image. As the visual channel is not available to our user group, we must attempt to compensate with additional feedback such as speech and non-speech audio.

Unimodal Trajectory Playback

In the previous study described in D6, the performance of a group of visually impaired users was compared to that of a group of sighted users in recreating shapes felt through haptic trajectory playback. A user felt a trajectory 5 times with no visual feedback available for either group. The task set to the user was to recreate that trajectory. A Neural Network trained on the trajectories was used to objectively measure the performance of the users. Specific findings of this study were that:

- The sighted group performed significantly better than the visually impaired group in the task. There was significant room for improvement for the visually impaired group.
- Users had problems segmenting the trajectories played back from forces used to drag the user to the centre of the workspace after finishing a playback.
- When the playback forces were turned off at the end of a playback, the user's hand tended to sink with gravity, perceptually adding on a downward tail to the trajectory,

2.3.4. Combining Audio and Haptic Trajectory Playback

A similar approach to the study described in D6 to was used for the haptic trajectory playback. The PID playback library developed for WP4 was used to drag the user through the path. This trajectory playback system is now available as an open source library (more details can be found in Crossan *et al.* (2006)).

However, to aid the user in determining the end point of a trajectory, the haptic feedback received at the end of the trajectory was changed. Once the trajectory playback was completed, users were held in place at the end point by the PHANTOM motors for two seconds before being returned to the central position. The goal here was to avoid users adding a tail to the trajectory as was noted in some traces from the previous playback experiment.

The goal of the auditory feedback was twofold: to help users segment the trajectory by providing more discernable start and end points, and to aid the user in discerning the shape of the trajectory. To help the user segment the trajectory, short distinct audio cues were played at the start and at the end of the trajectory. Further to this, during the drag back to the centre at the end of each playback, a wood scraping sound was played.

Audio cues were developed to help users to better determine the shape of the trajectory. Each trajectory in this study was limited to a two dimensional vertical plane. The audio cues were therefore developed to provide an indication of the current horizontal and vertical position in the environment. A sinusoidal tone was played to the user during the trajectory playback. The pitch of the tone was used to indicate vertical position. A higher pitch indicated a higher vertical position and a lower pitch a lower position. All shapes were either 12cm or 6cm in the vertical plane which corresponded to a one octave or a half octave change in pitch respectively. The frequency range over all the ideal trajectories was 200Hz to 400Hz with the midpoint (and trajectory start position) being mapped to a pitch of 300Hz. Audio pan was used to indicate horizontal position in the environment. As the user moved further left, the audio tone panned to the left and *vice versa*. The maximum pan was at 6cm horizontal distance from the centre (the maximum horizontal distance from the centre of the ideal shapes).

Two options were available when considering audio playback. The pitch and pan of the audio tone could either be related to the current controller position or the current user's position in the environment. The choice was made that the audio feedback corresponded to the user's position. Although the inevitable small deviations from the ideal trajectory ensured that a slightly different audio tone was heard during each playback, this ensured that the users could relate the sounds heard to their movements.

Methodology

Equipment

Similarly to the previous study, the device used for the trajectory playback in this study was the PHANTOM OMNI. In this study, the participants also had access to the keyboard for recording their drawing attempts. In the previous study the PHANTOM OMNI stylus buttons had proved problematic. To avoid these issues, participants were asked to hold down the spacebar while drawing (instead of holding down a stylus button) with their non-dominant hand and release when they had finished. Two speakers placed to the front left and front right of the participants were used to play the audio. As in the previously described experiment, PID Playback LIB (described in Crossan *et al.* (2006)) was used for this study. Furthermore, the preset PHANTOM OMNI control parameters were again used for the playback controller.

Conditions

Ten visually impaired participants from the Royal National College for the Blind in Hereford in the UK took part in the study. Four were female, with the age of the users ranging from 18 to under 45. Two participants were blind from birth and eight had some vision at some time in their life. Of these eight, six participants still had some residual vision.

A within-groups design was used with all users performing two conditions in a counterbalanced order. The two conditions for this experiment were the haptic playback condition (similar to the previous experiment) and the combined haptic and audio playback condition, referred to as C_H and C_{HA} respectively.

Training

All but two participants had some previous exposure to the PHANTOM. These two participants were given time before the experiment to feel the device and try a standard demo distributed with the device. Six of the nine participants from the previous study took part again (due to limited access to visually impaired participants), although with a gap of three months between the studies. Each participant initially went through a training period before starting the experiment. Three sample trials were conducted before the study was started with simple, nameable shapes – a circle, a square and a triangle. This training took place before each condition. Additionally in C_{HA}, participants took a short time (~1 minute) to freely explore the limits of the space to learn the mapping between their movements and the sound generated.

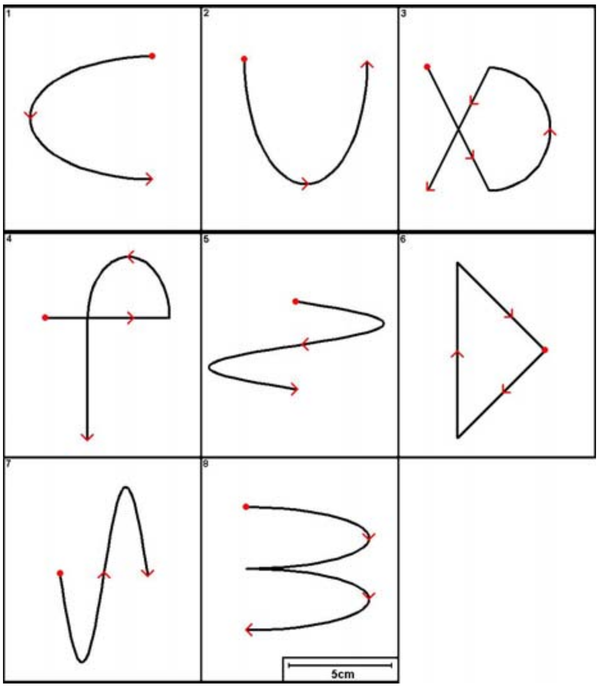

Figure 12: The trajectories used for the experimental task. The circle indicates the start point and the arrows the direction of movement.

Task

The task set for the participants was again to recreate the trajectories that were played to them. There were eight trajectories for each condition with the complexity of the trajectories in each condition being kept constant by using mirror image trajectories (flipped in the horizontal) in either of the conditions. The order of presentation was randomised. Each participant would experience a trajectory five times and draw it three times. However, unlike the previous study, the drawing attempts took place after the 3rd, 4th and 5th trajectory playbacks. This design was chosen to examine user improvement as the number of learning events increased. The eight shapes used in a condition are shown in Figure 12.

Analysis of the Results

Cursor trace data were recorded from this study at a rate of 100 samples per second. As in the previous study, the initial success of the participants in recreating the shapes was determined by a three layer MLP Neural Network. The Neural Network was trained on all sixteen shapes used in the experiment. A shape was considered successfully recognised if the Neural Network returned the correct shape as the highest recognition probability for all shapes. To ensure a fair comparison, the cursor traces from the mirror image trajectories were flipped in the horizontal. Recognition for the corresponding images in each condition was therefore relying on the probability of the same output of the Neural Network.

Further to Neural Network recognition, three independent human raters were asked to judge the closeness of fit of each trajectory to the ideal trajectory. Human raters were used to give an estimation of closeness of fit rather than a simple binary 'recognised' or 'unrecognised' result. All drawing attempts were grouped by shape with the order of presentation randomised. The raters were able to browse on paper all attempts at drawing each trajectory and assign them a rating between zero and ten for closeness of fit to the ideal trajectory. The limits of the scale were 'No Match' (zero) and 'Perfect Match' (ten). No indication of participant identifier, experimental condition, or repetition number was given to the raters. The measures for closeness of fit were left for each rater to decide. Drawings were pre-scaled (while maintaining aspect ratio) to the maximum size of shape that would fit inside a fixed box. This was to remove size of drawing as a variable to allow a more valid comparison to the Neural Network results.

Hypotheses

The main hypotheses for the study were:

- H1 – The recognition rates for C_{HA} would be significantly higher than those for C_H
- H2 – There would be a significant improvement in recognition rates between the first and third drawing attempts in both conditions.
- H3 – The rating given to the drawings in C_{HA} would be significantly higher than in C_H

Results

Neural Network Recognition Results

The data for recognition correctness were tested for normality and found not to be normally distributed ($p < 0.02$). A non-parametric Paired Wilcoxon test was therefore used to test for significant differences between C_H and C_{HA}. Initially, recognition of the raw trajectories was performed (without flipping the axes). Results show that the number of trajectories recognised by the Neural Network was significantly higher in the multimodal condition than in the unimodal condition ($W_9= 40.5$, $p < 0.04$). Of 24 possible trajectories for each participant, there were a median of 21.5 (mean = 19.4, std dev = 5.36) correctly recognised trajectories for C_{HA} compared to 18.0 (mean = 16.1, std dev = 6.45) for C_H.

This pattern was repeated when the mirror trajectories were flipped such that paired trajectories in each condition were compared to the corresponding eight trajectories in the Neural Network. Of 24 possible trajectories for each participant, there were a median of 21 (mean = 19.1, std dev = 5.40) correctly recognised trajectories for C_{HA} compared to a median 18.5 (mean = 17.0, std dev = 5.31) for C_H. A Paired Wilcoxon test was again used to show that this difference was significant ($W_9= 40.0$, $p < 0.05$). For all of the remaining Neural Network results, the flipped cursor data was used as it provides a fairer comparison between the conditions.

To test order effects during the study, success of the participants due to order of presentation was also tested. As the conditions were counterbalanced, five participants performed C_{HA} first, and five performed C_H first. The median number of correctly recognised trajectories for the first condition presented was 19.5 (mean = 17.6, std dev = 4.74) compared to 20.5, (mean = 18.5, std dev = 6.08) for the second condition. These data were tested using a Paired Wilcoxon test, and no significant effect was found.

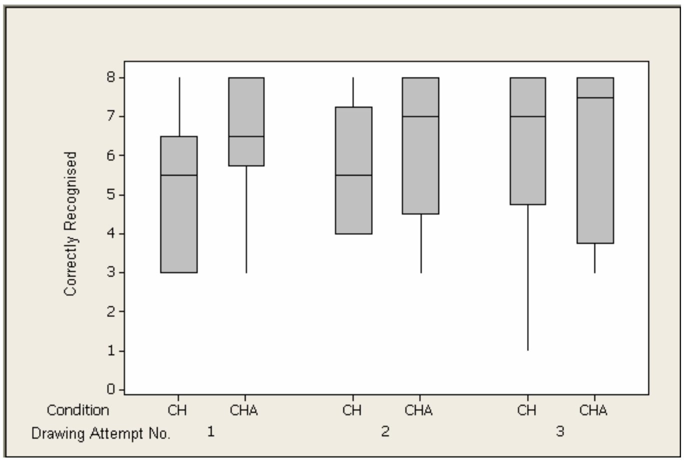

Figure 13: The median number of correctly recognised trajectories for all participants over all eight shapes for the three difference repetitions.

Learning due to multiple playbacks was also examined. At the first drawing attempt, a user had experienced a trajectory three times, at the second attempt four times and at the third five times. Figure 13 shows the median number of correctly recognised drawings for the eight trajectories performed in each condition over the three repetitions. A Paired Wilcoxon test was used to examine user performance for the different repetitions. There was no significant improvement in performance between the first repetition and the third repetitions in C_H ($W = 8.5$, $p = 0.11$) or C_{HA} ($W = 8.0$, $p > 0.9$).

The total number of trajectories correctly recognised for each participant in each condition is shown in Figure 14.

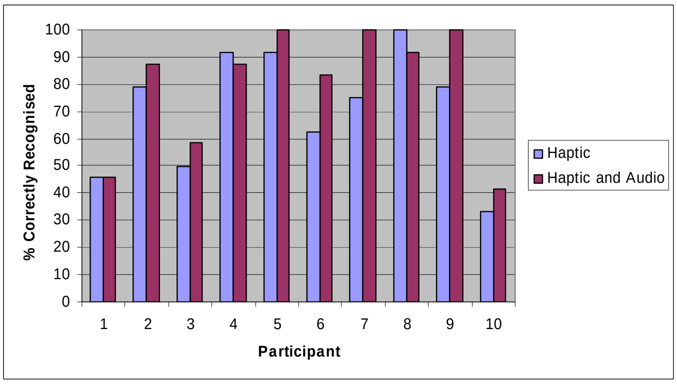

Figure 14: The percentage of trajectories correctly recognised for each participant in each condition.

Figure 15 shows the success of recognition by the Neural Network of each trajectory in either condition.

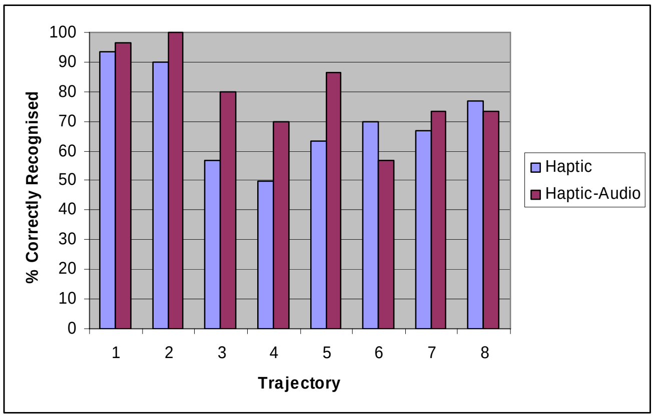

Figure 15: The percentage of correctly recognised trajectory's for all participants for each trajectory.

Rater Results

The results from three human raters were first tested for consistency between the raters to ensure the measures that the raters chose broadly agreed. A Pearson Product Moment test for correlation showed a strong correlation between the ratings given by each rater when rating the participant trajectories with $p < 0.01$ in all cases (correlation coefficients: R_1 and R_2 = 0.81, R_1 and R_3 = 0.83, R_2 and R_3 0.78). As the results from the raters correlated, the rating given to each trajectory was then given as the mean of the three ratings given to the trajectory.

Initially the ratings given to the trajectories that were recognised by the Neural Network were compared to those that were not recognised (see Figure 16). The recognised trajectories were rated at a median of 6.67 (var = 2.89) compared with a median of 3.67 (var = 3.33) for the unrecognised trajectories. This difference was shown to be significant with a Mann Whitney test ($p < 0.001$).

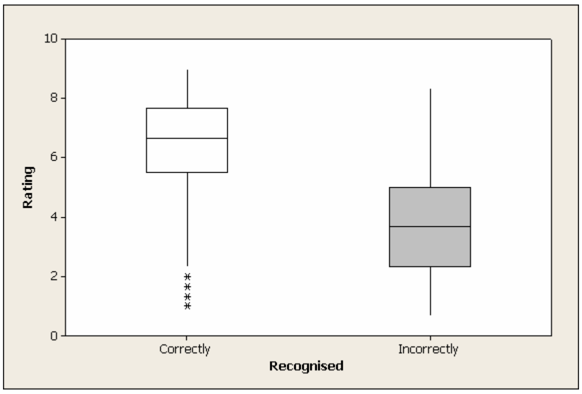

Figure 16: Graph of the ratings given to the trajectories separated by recognised and unrecognised trajectories.

When comparing ratings between conditions, a median rating of 6.33 was given to drawings in C_{HA} compared with 6.0 in C_H. A Friedman test was used to test for significance using paired comparisons of participants' ratings in either condition. There was no significant difference detected between C_H and C_{HA} (S = 0.40, p = 0.53).

Similar comparisons were made when looking at performance for the multiple drawing attempts. Figure 17 shows the median rating given to the first, second and third attempt at drawing a trajectory in both conditions. A Paired Wilcoxon test showed no significant difference in the rating between the first and third repetition for C_H (W = 1408.5, p = 0.23) or for C_{HA} (W = 1311.5, p = 0.85).

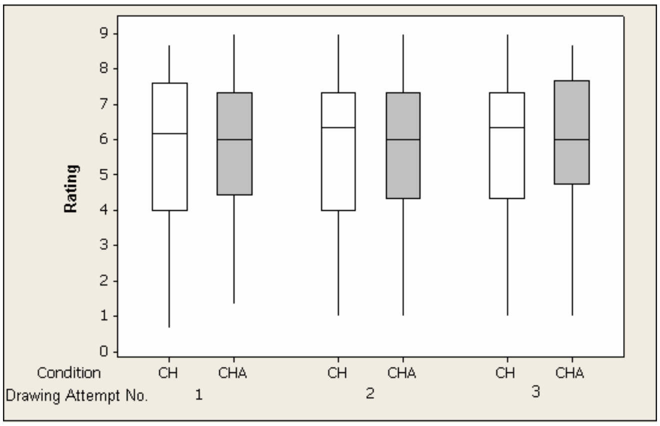

Figure 17: Median rating over the entire set of trajectories given to each attempt at drawing a trajectory in C_{HA}

Other Results

For both the Neural Network results and the human rater results, size was disregarded as a factor. The Neural Network recognised a small shape equally well as the same shape drawn larger and the drawings that the raters saw were pre-scaled (while maintaining the aspect ratio) to provide a valid comparison with the Neural Network results. However, sometimes the size of the trajectory is important for the task as well as the shape. The size of the trajectories drawn by each of the users was therefore compared to the ideal size. Cursor logs were analysed to determine the range of X (horizontal) and Y (vertical) values in each trajectory. The error in X and Y trajectory size was given as the percentage difference between the actual size and the ideal size with respect to the ideal size.

For C_{HA}, the median percentage X error with respect to the ideal size was 26% (conf. int. = 6.42) compared to 36% (conf. int. = 11.57) for C_H. This was tested using a Paired Wilcoxon test with no significant difference being detected ($W = 11.0$, $p = 0.36$). The median percentage errors in Y were 26.02% (conf. int. = 10.30) and 37.98% (conf int. = 13.65) for C_{HA} and C_H respectively. These data were again tested using a Paired Wilcoxon test and the trajectories in C_{HA} were found to be significantly closer to the ideal in Y ($W = 0.0$, $p<0.02$).

The additional audio cues proved successful for allowing users to segment the shape from the drag 'back phase'. From observation, only in one trajectory did one user confuse the 'drag back' section of a trial with the playback. The participant noticed the mistake at the second drawing attempt of the trajectory and reported it to the experimenter. Visual analysis of the cursor traces also indicated that the tail noticed in the previous study (due to gravity affecting the user's movements at the end of a trajectory) was resolved by sticking the user temporarily to the end point of the shape.

Discussion

The results from the Neural Network recognition of trajectories indicate that combining haptics and sound can lead to significant advantages when learning the shape of trajectories non-visually. After three times experiencing the trajectory, visually impaired users were able to recreate the trajectory to such an extent that approximately 80% of trajectories drawn were recognised compared to 70% without the additional audio playback. These results support H1. Analysis of the multiple repetitions shows that user performance did not significantly change between the first and third drawing attempt in C_{HA} or C_H. This, combined with the recognition results, suggests that users had a good idea of the trajectory by the third time that it was played back. It was therefore not possible to support H2 with these data.

Similarly to the previous study, there was a high level of variability between participants. The two participants who had the least number of trajectories correctly recognised were the two participants who had never had vision. This suggests that the task is most difficult for participants with the least experience of writing and of visualising shapes.

Despite the fact that mirror image trajectories were used in either condition, there was no significant improvement in performance detected between the first condition presented and the second. This suggests that little learning effect between the conditions took place due to the similarities of the shapes.

When comparing the Neural Network recognition results with the human rater results, there is agreement between more accurate and less accurate attempts at drawing a trajectory. This is indicated by the fact that the rating given to trajectories that were recognised wrongly by the Neural Network were given a significantly worse rating than those that were recognised correctly. However, no significant differences were detected between the ratings given to the drawings in C_H and C_{HA}. There are many possible reasons for this difference. The Neural Network will be robust to some errors that may cause the raters to reduce their rating such as sharp corners being rounded or straight lines being slightly skewed or wavy. Conversely, the information given to the raters was different from that received by the Neural Network. The raters' task was a pattern matching task as opposed to a trajectory recognition task. The dynamic aspects of the trajectory were removed from the representation presented to the raters. With these results, it was not possible to support H3.

Analysis of the cursor trace shows that users could more accurately represent the size of the trajectory using the audio and haptic feedback combined. Significantly greater accuracy was shown in the vertical dimension but not the horizontal. This can be explained by *post hoc* discussion with participants. The pitch was found to be far easier to determine than the pan for all users. Most users reported entirely disregarding the audio pan information as they found it difficult to concentrate on both pitch and pan at the same time. Although the addition of audio feedback was shown to improve performance here, future work should examine alternative auditory designs.

2.3.5. Guidelines

The following guidelines can be extracted from this work:

- Audio cues can be combined with trajectory playback to allow the users to segment trajectory playback events.
- Combining audio with haptic trajectory playback can significantly improve users' performance in learning and recreating trajectories then haptic trajectory playback alone.
- Pitch can be used successfully by a user for one dimension in a haptic trajectory playback learning task.
- When using pitch and pan for the audio playback, designers should be aware that it is difficult to concentrate on both pitch and pan at the same time. Pitch tends to be easier to distinguish and therefore dominates.

2.3.6. Conclusions and Future Work

This section has described a study to examine the performance of visually impaired people in trajectory learning task. It has shown how the incorporation of multimodal haptic and auditory cues can improve performance over haptic cues alone both for recreating the trajectory and segmenting trajectory playback events. This work has since been integrated into a collaborative drawing environment. The environment uses these trajectory playback techniques to allow a sighted teacher to describe a diagram to a visually impaired student. This system is currently being evaluated as part of WP3.

3. Tactile External Memory Aids

With tactile external memories, we must take a less direct approach when interacting with the user. Tactile cues, similarly with audio, can be used to present information to users but cannot directly affect their movements through the environment. It is, however, currently unknown how best to present the information in an efficient manner. Previous studies described in D6 examined various designs of tactile cue, presenting information to the user through a tactile array. Here we examine ways of increasing the amount of information in a tactile cue that can be interpreted by the user.

3.1. Introduction

This section describes an evaluation of multi-dimensional tactile cues. This is an extension of previous work in MICOLE examining uni-dimensional Tactons (see Pietrzak *et al.* (2006), and Crossan and Brewster (2006)). Here, we examine how we can increase the flow of information to a user. It is envisaged that these cues will be used in a task involved in browsing or navigating non-text or complex data. This task can place a strain on the user's memory as it is difficult to replace the visual glance with a non-visual equivalent. Particularly when there could be large quantities of data or data that the spatially relationships are important to the understanding. Tactile cueing could be one potential solution allowing users to locate previously visited objects in the workspace or to display data about an object while still allowing their auditory sense to process speech or audio data simultaneously.

3.1.1. Background information and relevant literature

<u>Tactile Messages</u>

Previous work in the MICOLE project has looked at the problem of non-visual navigation around a two dimensional environment. Touch and in some cases audio have been used to guide the user around a path. The Multimodal maze environment developed and evaluated by the University of Glasgow show how tactile messages can be used guide users through an environment. In this section we attempt to formalise the design of a more general form of tactile messages with respect to related tactile research described by Brown *et al.* (2005). They define Tactons as "structured tactile messages" and are analogous to Earcons in audio (Brewster (1994)).

This work builds from the previous work on Tactons displayed through vibration actuators. Brown *et al.* (2005) initially examine one dimensional Tactons and show how a high degree of accuracy can very quickly be reached for simple Tactons. They then demonstrate how using multiple independent identifiable parameters - that do not interfere with the recognition of other parameters - can be used to increase the flow of information to the users. Multiple vibrotactile actuators were attached to different locations on a user's forearm. Information was displayed to the user through three parameters of the tactile signal: The rhythm at which the actuator is vibrated, the roughness of the vibration and the body location of the actuator stimulated. Brown *et al.* (2006) demonstrate that users could achieve a high level of accuracy with three varying parameters of the tactile signal. However, they note that a careful choice of parameters is required such that there are no interactions between the parameter when varying their value. For example, they note that increasing the frequency of a tactile signal can also lead to a perceived increase in amplitude in the signal. Further work by Hoggan and Brewster (2006) have examined the cross-modal equivalence between haptic Tactons and auditory Earcons. This would allow the same information presented in different modalities in different situations.

Other examples of structured vibrotactile messages include Vibetonz from Immersion). They suggest the incorporation of tactile messages into a mobile device. Similarly Motorola have studied the addition of vibrotactile feedback into a mobile device in order to enhance the audio feedback Chang and O'Sullivan (2005).

Related work into the parameters of tactile perception is discussed by MacLean and Enriquez (2003). They describe a series of experiments examining usable parameters for haptic icons or 'Hapticons'. To display the Hapticons, their studies used a one dimensional force feedback wheel which the user grasped between the thumb and a finger. The parameters examined to display data included frequency of vibration, waveform shape and force magnitude.

Pin Array Tactons

The above examples examine usable parameters for displaying information tactilely through vibrotactile devices. Pietrzak *et al.* (2006) have since examined user performance in distinguishing tactile messages through a small 4x4 tactile array. In this study the authors present the message on one 4x4 tactile display on the VTPlayer tactile mouse (shown in Figure 18). The authors study one dimension of information display using three styles of Tacton: static, dynamic and blinking. In the static condition, the Information was presented using the shape and position of the pattern of raised pins. The dynamic condition used Tactons composed of several frames of equal length presented serially. In each frame, a pattern on the grid would shift in a direction and the flow of the pattern was used to present information. The blinking style of Tactons was composed of two frames of equal length. The first frame presented a pattern similar to that of the static patterns, while the second frame had no pins raised. These two frames were alternated to give the impression of a pattern blinking off and on.

Similarly, Crossan and Brewster (2006) examined performance of users in a tactile guidance task using static and dynamic. The task was to navigate a maze by moving a PHANTOM OMNI device (from SensAble Technologies) in their dominant hand while interpreting Tactons indicating direction to the exit through their non-dominant hand. Participants were significantly more successful and faster in the static condition.

Results in both studies by Pietrzak et al. and Crossan et al. demonstrated that that the poorest performing style of Tactons was dynamic. For this study, dynamic Tactons were therefore disregarded. There were no significant performance differences demonstrated between the blinking and the static Tactons. Blinking patterns will be considered for this study as it allows us to consider blink speed to be used as a parameter for transferring information to the user.

Figure 18. The VTPlayer tactile mouse.

3.1.2. How it fits into the deliverable

This work is building towards the goals of 2.2.2. These cues will be used as part of a larger multimodal interface to guide the user to certain points in the interface or around a path. One potential use of these cues is therefore to provide external memory aids to users in a multimodal environment. The user may receive speech and non-speech sounds as well as haptic effects through his or her dominant hand while feeling the Tactons through a tactile array in the non-dominant hand. The non-dominant hand will therefore receive Tactons giving information about the position or type of objects of interest in the environment.

3.2. Designing Multidimensional Tactons

3.2.1. Choosing Appropriate Parameters

When developing Tactons on a pin array device, there are many possible parameters that can be used to transfer information to the user. It is key to the usability of the system however, that an appropriate group of parameters is chosen. When many dimensions of information are present in the signal, it is important that they not interfere such that the user finds it difficult to perceive the individual dimensions. Braille is one example

Brown *et al.* (2006) describe their experiments to develop multidimensional Tactons that use dimensions suitable for vibrotactile feedback. There are many potential dimensions available for use for pin array Tactons. For example, we may consider using pattern shape, pattern size, pattern movement, blink tempo, or blink rhythm. However, not all of these will provide benefit. Previous work has shown for example that dynamic shifting patterns can be more difficult to distinguish than static patterns (Pietrzak *et al.* (2006), Crossan and Brewster (2006)). For these experiments, we therefore disregard shifting patterns as a parameter. The three parameter chosen for this initial study were pattern shape, pattern size and blink speed. Pattern shape has previously proved to be a successful method of transferring information through pin array devices. Similar patterns were chosen as one parameter of the Tactons used in this study.

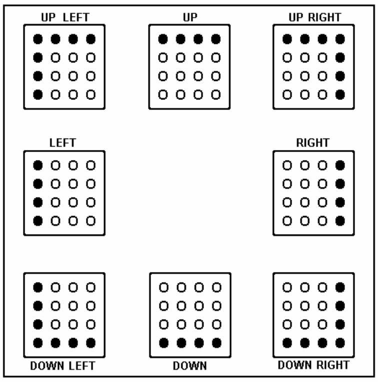

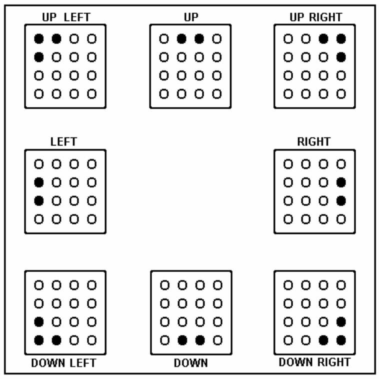

Figure 19. The eight large patterns (left) and corresponding eight small patterns (right) used in this work.

In this instance, patterns representing the cardinal points of the compass are made up by a single line of pins at the top, bottom, left or right of the display. The diagonal patterns consist of two lines of pins, for example at the top and left for a diagonally up and left pattern. The pattern size parameter could be distinguished by the number of pins raised to make the pattern. Lines of pins for the large patterns consisted of four raised pins for large patterns compared to two raised pins for small patterns. Only two pattern sizes are considered here as for a 4x4 array, symmetric patterns can only be formed by these lengths on lines. Figure 19 (left) shows the eight large patterns used for this study. They have been derived from previous work by Pietrzak *et al.* (2006). Figure 19 (right) shows the 8 small patterns used.

To test the suitability of pattern size and blink speed as useful parameters for transferring information to the user, pilot studies were first conducted individually.

3.2.2. Developing Usable Parameters

Pattern Shapes

Previous research by Pietrzak *et al.* (2006) influence the patterns chosen for this study. Eight patterns different patterns were chosen to represent the four cardinal compass points and four diagonals. The patterns chosen were those determined to be the most successful in Pietrzak's study. The shapes used for both large and small patterns are shown in Figure 19. The large patterns have been chosen from the Pietrzak study, with the small patterns being based on the same shape of pattern with shorter lines. Using the same patterns for the large Tactons as a previous one dimensional study allows basic comparisons to be made between one and multi-parameter Tactons. This allows us to gain insight into how the extra complexity brought on by the additional information being presented to the user affects user performance.

Pattern Size

Given the shapes chosen from Pietrzak's study and the limited number of pins available in the Virtouch array (4x4), only two sizes of shape were possible. A short pilot study was run with four participants was run to determine whether size was a viable parameter. Participants were able to distinguish between large and small patterns with minimal training over 90% of the time over 96 trials each. These results suggest that pattern size potentially a useful parameter that should be investigated further.

Blink Speed

When choosing blink speed values, there are a number of different methods that were considered. For P_u being the frame length that the frame where the pattern is displayed for and P_d being the frame length for the frame where no pattern is displayed we could choose values such that: P_u is always equal to P_d, P_u varies and P_d is constant, P_u is constant and P_d varies, or the sum of P_u and P_d is constant. Participants performed each condition for trials with 6 potential values for the blink speed ranging from 0.04s to 0.5s per frame. They were presented with two stimuli and had to reply whether the stimuli were the same or different. Only blink speed was varied and only one stimulus could be experienced at the same time. Confusion matrices were plotted of the results showing extremely similar performance for all four methods of varying the blink speed. P_u being set equal to P_d was chosen as the method for future studies as no preference was suggested by the users and using this method allowed comparisons to be made with the previous study by Pietrzak *et al.* (2006).

When selecting what values to use, there will be a trade-off between message length and distance between the values. If the message length is too long, this will slow the user's interpretation of the message whereas if the distance between the values is small, users will have trouble distinguishing between values. The confusion matrices showed there were no confusions between 5 of the 6 values chosen.

3.2.3. Experiment methodology

Participants

A between groups study was run with 20 sighted participants, who were students from the University of Glasgow. The age range of these participants was 18 to 29. There were four left handed participants split evenly between conditions. Sighted participants are used to provide a baseline performance and to inform the design of future studies with visually impaired participants.

Experimental Set-up

In all instances, participants felt the Tactons through the index finger of their non-dominant hand. The Tactons were displayed through the Virtouch VTPlayer tactile mouse with a 4 x 4 array of pins were used to display the patterns. The three parameters varied for each Tacton were 'Pattern Shape', 'Pattern Size' and 'Blink Speed'.

There were two conditions tested: a three speed condition (S_3) and a two speed condition (S_2). Participants were randomly assigned to one condition or the other, with balancing to ensure that an equal number of participants performed each.

Due to the fact that a different S_3 has an extra level of blink speed, there is a larger of range of Tactons available in this condition. For an experiment comparing S_3 to S_2 performance, a decision therefore needs to be made as to make a fair comparison. Three possibilities are to present:

- an equal number of trials during each condition.
- an equal number of repetitions of each Tacton in each condition.
- only a subset of the available Tactons in S_3.

If an equal number of repetitions on each Tacton is used, it will lead to participants in S_3 performing more trials and taking longer to complete the study. User fatigue and learning effects or potential factors here. Allowing only a subset of S_3 to be presented to the participants complicates analysis of the results for the other Tacton parameters. Therefore, here we choose to maintain an equal number of trials in each condition. There are 32 unique Tactons for S_2 and 48 for S_3. Each participant was given 96 Tactons to identify. Participants in S_2 and S_3 were therefore presented with two or three repetitions of each Tacton respectively.

Participants placed their non-dominant hand on the VTPlayer mouse and held down a key on the keyboard with their dominant hand to feel the Tacton. Once they release the key, the Tacton was stopped. They then gave their answers verbally to the experimenter. For Pattern shape, the answers were given as ‘up’, ‘down’, ‘left’, ‘right’, ‘up and left’, ‘up and right’, ‘down and left’ and ‘down and right’. Timing data was measure as the length of time the Tacton was felt for (the length of time the key was held down). The maximum time the Tacton was felt for was capped at 10 seconds.

Information Transmission

Information Transmission is a measure of the correlation between the amount of information in the stimuli and the amount of information in participants’ responses (see Miller (1956)). Due to the different number of levels of the Blink Speed parameter in each condition, there are a different number of available messages in S_2 and S_3 – 32 messages for S_2 and 48 messages for S_3. This leads to a different amount of information being presented to the user in stimuli in each condition. In S_3, there are more available messages so more information present in each stimulus. The number of bits of information can be given as $\log_2(N)$ where N is the number of available messages. This gives values of 5 bits of information per Tacton in S_2 and 5.59 bits of information per Tacton in S_3. The amount of information eventually transmitted to the user will depend on the error rate. This will lead to a decrease in the information present in the users’ responses and therefore a corresponding decrease in the information transmitted to the users. Miller (1956) provides a fuller discussion of this as well as specific details for calculating Information Transmission.

Hypotheses

For this study, we hypothesise:

- Participants in S_2 will make significantly less errors in identifying Tactons. We hypothesise this due to the fact that participants should find it easier to differentiate two as opposed to three blink speeds.
- Participants in S_2 will take significantly less time to identify Tactons than participants in S_3. Again, this hypothesis is due to the fact that participants should find it easier to differentiate two as opposed to three blink speeds.
- There will be significantly more information transferred during S_3 when compared to S_2. Although more errors may be made by participants in S_3, there is also the potential for a greater amount of information to be transferred.

3.2.4. Evaluation results

S_3 vs S_2

Due to the discrete and bounded nature of the data, non-parametric analysis of the data was carried out. The Mann Whitney test was therefore used to test for significant differences in the independent measures between groups. For comparing data within groups for the individual parameters, Paired Wilcoxon tests were used.

When looking at correct identification of all three parameters together, there were a median of 12 (mean = 12.3, std. dev = 8.4) errors in S_2 compared to a median of 27 (mean = 24.2, std. dev = 13.0). This difference was shown to be significant ($W = 135.0$, $p < 0.03$). When analysing each of the individual parameters, the difference was shown to be due to the blink speed parameter. Participants in S_3 made a median of 9.5 (mean = 11.5, std dev. = 8.0) errors in S_3 compare to 0.5 (mean = 0.9, std. dev. = 0.9) in S_2. No significant differences were found in the Pattern shape or Pattern size data. Figure 20 shows the percentage correct for each parameter as well as the total number of correctly identified Tactons. There were no significant difference recorded in time taken to identify the Tactons in S_3 and S_2 ($W = 107.0$, $p = 0.91$). The median time to identify a Tacton for participants in S_3 was 2.09 seconds (mean = 2.7, std. dev. = 1.6) compared with 2.31 seconds (mean = 2.73, std. dev. = 1.95).

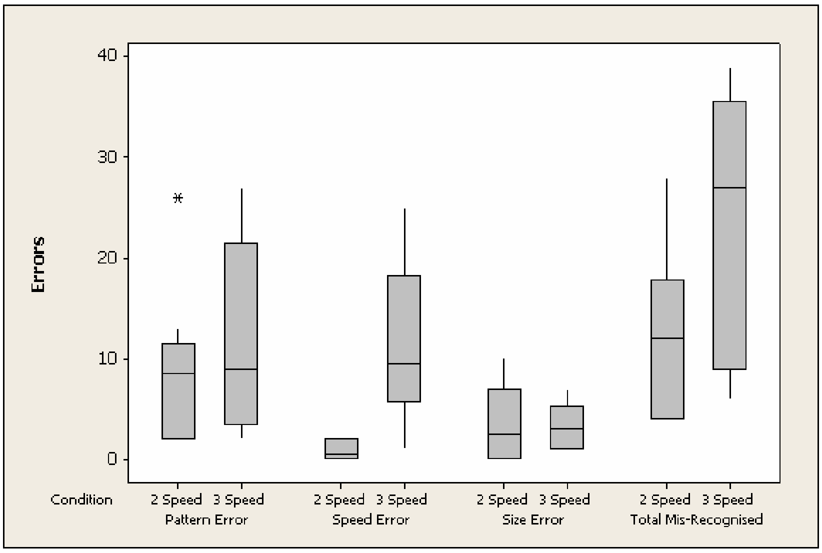

Figure 20. A box plot showing errors for each parameter and overall, and for both conditions.

When examining the information transmitted, a median 5.0 (mean = 5.0, std. dev. = 0.32) bits per Tacton was transmitted in S_3, compared with a median of 4.6 (mean 4.6, std. dev. = 0.26) bits per Tacton in S_2. These data were tested using a Mann Whitney test and this difference was found to be significant ($W = 75.0$, $p < 0.03$).

Significant Interactions Between Parameters

Pattern size had a significant effect on participants' performance in both S_3 and S_2. The large patterns were significantly easier to identify than the smaller patterns. There were significantly more Tactons of small size mis-recognised in both S_3 ($W = 0.0$, $p<0.01$) and S_2 ($W = 0.0$, $p<0.01$). The median number of mis-recognised small Tactons in S_3 was 13.5 (mean = 15.8, std. dev. = 9.1) compared with a median of 7.5 (mean = 8.4, std. dev. = 5.9) for large Tactons. This pattern was repeated in S_2 with medians of 7.5 (mean = 8.4, std. dev. = 5.9) and 13.5 (mean = 15.8, std. dev. = 9.1) mis-recognised Tactons for large and small patterns respectively. By analysing the individual Tacton parameters, it is shown that this difference is due to significantly more pattern shape errors being present in the data in both S_3 ($W = 0.0$, $p<0.01$) and S_2 ($W = 0.0$, $p<0.01$). In S_3, there were medians of 1 (mean = 1.3 std. dev. = 1.6) and 8 (mean = 11.0 std. dev. = 8.4) errors in distinguishing the pattern shape for large and small patterns respectively. Similarly, there were median errors of 1 (mean = 1 std. dev. = 1.2) and 7 (mean = 7.9 std. dev. = 6.5) for large and small patterns respectively in S_2.

Participants, also required longer feeling the stimuli to answer when identifying small Tactons and large Tactons in S_3 ($W = 6.0$, $p<0.04$) and S_2 ($W = 6.0$, $p<0.04$). For S_3, participants felt the stimuli for medians of 1.94s (mean = 2.45, std. dev. = 1.39) and 2. 27s (mean = 2.94, std. dev. = 1.77) before answering for large and small patterns respectively. Similarly for S_2, participants felt the stimuli for medians of 1.97s (mean = 2.55, std. dev. =2.0) and 2.62s (mean = 2.92, std. dev. = 1.92) before answering for large and small patterns respectively.

3.2.5. Discussion

Firstly, it can be seen that the method used to analyse the data affects how we view the results. Looking simply at error rates, there were significantly fewer errors in the two speed condition when compared to the three speed condition. This is unsurprisingly due to an increase in errors in identifying the speed parameter. This is an interesting result as it to some extent enforces the findings of Brown *et al.* (2006). Participants are making relative judgements for the blink speed and Brown's finding suggested that adding a third level of a parameter that required relative judgements could significantly increase complexity for the user. These data would suggest that it is better to use only two blink speeds.

However, if we examine the data in this manner, we ignore the benefits of including an extra level of this parameter. It allows us to transfer more information to the user in one Tacton. The benefits are brought out by the Information Transmission results showing that significantly more information per Tacton was transmitted to the users despite these extra errors.

The significant interactions shown between pattern size and pattern shape suggest that these parameters might not be compatible multi-parameter Tactons on a small tactile pin array. Performance was shown to be significantly worse for the small Tactons indicating. When choosing parameters for multi-dimensional Tactons, they should not interfere. However, decreasing the pattern size affect how participants perceived the shape of the pattern in both S_2 and S_3.

3.2.6. Conclusions

This section has presented a study on multi dimensional Tactons presented through a small pin array. It has compared three dimensional Tactons with different levels of information in the signal. It has shown how although increasing the information flow in each Tacton can lead to more errors by the participant, it can also lead to a greater amount of information being transmitted to the user and therefore a potentially a more efficient interface. Pattern shape and pattern blink speed have both been shown to be viable parameters for displaying Tactons on a pin array. However, interactions between pattern size and pattern shape suggest that designers should think carefully before combining these in an interface using a small tactile pin array.

This study was conducted with sighted users instead of blind users. However, the results from this study will allow us to refine our experimental design. Future work will also include a thorough comparison of the results from this multi parameter Tacton study with the study by Metz described in D4 looking at single parameter Tactons. Information transmission will be one factor that is taken into account, although we will also examine the drop off in performance of the one parameter used in the Metz study (pattern shape) when additional parameters are added.

3.3. Guidelines

The following guidelines can be drawn from this study:

- Designers can increase the amount of information presented through a Tacton by increasing the number of independent parameters in the Tacton.

- Pattern shape and blink speed are two factors that can be used independently.
- For small arrays, using pattern size will affect the user's perception of the pattern shape. Smaller patterns are harder and take longer to distinguish.
- The range of blink speeds used in this study (0.04s per frame to 0.5s per frame) did not significantly effect the time it took the user to recognise the Tacton.

3.3.1. Future developments

A late cancellation from the college that our user group is drawn from has meant that the corresponding study with visually impaired users has been postponed until after the submission of this deliverable. The major area of future work will be to run a similar study with visually impaired participants as there are no guarantees that similar results will be obtained. A factor such as level of experience in Braille use for example could have a significant effect on the results obtained.

4. Audio External Memory Aids

Providing external memory aids involves the support for finding and re-finding objects, without having to rely on memory. Auditory beacons are often described as a tool for this, and 3D non-speech audio is often the sonification of choice, but is 3D audio the most effective way to provide localisation cues? This study addresses the essential task of getting an overview and finding objects in a haptic virtual 3D environment, essentially what is involved in re-finding a marked auditory beacon.

4.1. Designing an Auditory Display to Facilitate Object Localization in a Haptic Virtual 3D Environment

4.1.1. Introduction

Finding objects in space is difficult for people with a visual impairment. Without vision, the ability to identify, localize and grab for objects is limited. This problem exists both in real and in virtual environments. In virtual environments, however, it is possible to fully control the information presented, and to use abstractions and a non-realistic presentation to enhance the interaction. Sound and haptic feedback could be combined to enable visually impaired persons to successfully interact with complex and dynamic virtual 3D environments. Research on this topic is scarce, but very necessary. Furthermore, the field of application for this research is larger than enabling access to virtual environments for visually impaired persons. The audio or haptic modality can be the most suitable modality for presenting different types of information (Brewster, 2003; Snibbe et al., 2001), for example, an auditory signal representing a fire alarm, or force feedback in a steering wheel. The auditory and haptic modality are also of interest because it can reduce visual overload (Brewster, 2003; Tang et al., 2005), for example, by reducing the number of warning lights on a airplane control panel by having a computer read aloud the warnings. Furthermore, the combination of modalities can improve the experienced realism and immersion with virtual environments (Alcañiz, Lozano & Rey, 2004).

Various devices have been developed providing haptic feedback. In some cases tactile feedback was added to an already existing device, such as the FEELit mouse introduced in 1997 (Burdea, 2000). In other cases a device was specially designed to provide cutaneous information, for example, the refreshable Braille display that was introduced during the 1970's. This study considers a tool-handling force display, as this currently is the most practical and unrestrained device for exploring haptic virtual 3D environments (Iwata, 2003).

The first tool-handling force display that became commercially available was the PHANToM Haptic interface, developed by Massie and Salisbury in 1994 (Massie & Salisbury, 1994). Since then, various PHANToM haptic devices have been introduced, having three or six degrees of freedom for input (x, y, z, yaw, pitch, and roll) and three or six degrees of freedom in presenting force feedback. The PHANToM devices are currently one of the most popular commercially available haptic interfaces (Iwata, 2003).

Limitations of using a tool-handling force display

Haptic feedback is perceived through direct contact with the immediate environment. With the exception of heat, which can travel, no haptic information is available for objects that we are not in direct contact with. Without the other senses, an environment has to be physically explored to know what is there. With the PHANToM haptic device users have to grab a stylus or place one of their fingers in a thimble, a holder for the fingertip, in order to 'touch' a virtual environment. Force feedback is applied to the tip of the stylus or to the tip of the thimble. With the PHANToM haptic device a virtual environment is explored with only one point of contact. Picking up and enclosing objects is thus not possible when using only one device. Inside a virtual environment the point of contact is represented by a cursor. Normally, this cursor is very small, comparable to the tip of a ballpoint. With only a single point of contact that also is small, it is very hard to find small objects in a virtual environment and objects are easily missed (Sjöström, 2001a; 2001b).

Then, when an object is found, the haptic information perceived needs to be made sense of; the object needs to be identified. Both in vision (Goldstein, 1999) and in touch (Klatzky & Lederman, 1999), objects are identified mainly by their geometric properties. Compared to vision, the haptic system is slow and less accurate in the identification of objects (Klatzky & Lederman, 1999). Generally, the exploratory procedure of enclosure is used to quickly assess the geometric properties of an object. But with one-point interaction systems, enclosure is not possible, which further limits the identification of objects.

Next, when objects are found and identified, they need to be memorized and placed in a mental representation of the virtual environment. It takes a lot of time and effort to build an accurate mental representation of a virtual environment using haptic exploration only (Wall & Brewster, 2004). And when objects are moving, it becomes increasingly difficult to form an accurate mental representation. Since users lack this continuous overview of a virtual environment, it is important that reference points are presented to facilitate navigation and the forming of a mental representation (Sjöström, 2001a).

Overcoming limitations

To full exploit the advantages that haptic devices offer to visually impaired persons, the limitations have to be dealt with. To overcome some of the limitations, researchers have suggested several virtual haptic tools, for example, an enlarged cursor or attractive forces exerted by objects (Sjöström, 2001b). But virtual haptic tools can only partly compensate the limitations of unimodal haptic interaction. A different approach is to create a multimodal virtual environment, combining the haptic modality with another modality available to visually impaired persons, the auditory modality. We will investigate the auditory modality to overcome the limitation of finding objects in a haptic virtual 3D environment.

Using Sound to overcome haptic limitations

Speech is very suitable for naming objects and presenting absolute values. However, the slow rate of presentation, due to the serial nature of speech (Brewster, 1994), makes it less suitable for presenting continuous feedback on the position of objects. Therefore, non-speech audio is further investigated for facilitating object localization in haptic virtual 3D environments.

A general advantage of sound is that it can be heard over distance. Auditory information is perceived from sounds sources in all directions. Sounds are also rapidly detected and processed by the human brain (Kramer, 1994). Multiple sounds can be processed a once, called parallel listening (Kramer, 1994). When we hear our name being called in a conversation we are not attending, is an example of parallel listening. Some sounds can be attention grabbing, for example, the horn of a car, yet other sounds can blend into the background, such as the sound of the air-conditioner (Kramer, 1994). But when a sound that blended into the background suddenly changes, this change is often noticed.

Another advantage of non-speech sound is the wide variety of sound characteristics, such as pitch and roughness, that humans can discriminate. When a sequence of sounds is played, we perceive more than individual sounds. Sounds are grouped and discriminated according to auditory Gestalt principles like proximity and similarity. When hearing a sequence of sounds, we perceive characteristics as tempo, rhythm and melody. These perceptual dimensions can be used to present information.

Previous research on auditory displays

Researchers have designed various auditory displays that give visually impaired persons access to a graphical user interface. Winberg (2001) discriminated three main models on which auditory interfaces are based; the linear, hierarchical and spatial model. With the linear model, auditory information is presented sequentially. Screen readers are examples of auditory interfaces that are based on the linear model. With a hierarchical model, users step through a tree-structured list of which the items are read aloud. Examples of a hierarchical model can be found in the Mercator project (Edwards, Mynatt & Stockton, 1994) and in the auditory interface designed by Savidis et al. (1996). With the spatial model, objects are placed in a virtual area similar to a graphical user interface. The spatial model is the only model that represents a virtual environment in its natural form. While the linear and hierarchical models ensure that all information is easily accessible, the spatial model is the only model that allows a natural interaction with a virtual 3D environment.

One of the first systems that added an auditory user interface to the graphical user interface was the SonicFinder by Gaver (1989). With the SonicFinder, auditory icons were mapped to system events, such as selecting a file and copying a file. Auditory icons are everyday sounds that metaphorically represent actions and events, for example, the sound of a glass being filling with water used to represent a progress bar. The SonicFinder was only designed to augment the graphical user interface. The SonicFinder did not address the navigational issues when visually impaired persons try to use the interface.

Soundtrack, created by Edwards (1989), was one of the first spatially based auditory user interfaces specially designed for visually impaired persons. Soundtrack was a text editor that was fully accessible by auditory feedback. Edwards (1989) combined non-speech sound for navigating the user interface with speech synthesis for menu's and written text. Finding objects on the user interface was simple, because the workspace of the interface was divided into eight blocks and each whole block represented an object. The use of the non-speech sound did not work as expected. Users counted the sounds that were played when entering a new object, instead of inferring which object the sound indicated (Pitt & Edwards, 1991).

To create a more general application, Pitt & Edwards (1991) investigated how objects could be located by sound with the use of a mouse. First, a visual navigation experiment was conducted; investigating how well users could use distance-, directional- and combined information to find an object. The results show that objects are approached faster using directional information than using distance information. They also found that when directional information and distance information was presented simultaneously but separately, the localization of objects took much longer than when the directional and distance information were presented in their natural combined form, when presenting a cursor and target object. Pitt & Edwards (1991) then replaced the visual cues with auditory cues. An object was represented by a synthesized cello tone. Volume was used to represent distance and was used independently or combined with stereo panning, which represented the horizontal direction, "the intended analogy was that of the cursor acting as a microphone" (Pitt & Edwards, 1991). Pitt and Edwards (1991) also experimented with a non-linear volume increase, having the volume increase faster when the cursor comes closer to the target. This was done because, in a pilot, users were fast in moving the cursor close to an object but took a longer time to finally locate the object. Their results show that both the use of a non-linear increase and the use of stereo panning reduced the time necessary to locate an object.

Mereu and Kazman (1996) experimented with audio enhanced 3D interfaces. The x, y, and z-coordinates of the cursor and of the target were mapped, in the same order, to panning, pitch and volume. Standard, the user would hear a sound, which represented the position of the cursor, but when pressing the left mouse button he or she could hear a sound, which represented the position of the target. Mereu and Kazman (1996) found that users could accurately locate a target by comparing sounds.

Magnusson and Rassmus-Gröhn (2005) investigated two different auditory user-interfaces integrated in a haptic virtual 3D environment. First, they investigated a 3D version of the virtual microphone method from Pitt and Edwards (1991), mapping virtual sounds sources to the position of objects and using the cursor as a microphone to listen to the virtual environment. And second, they investigated the use of virtual "sound beams" extending from objects along the z-axis. The virtual microphone method was reported as intuitive to use, but nothing was reported on the use of virtual sound beams (Magnusson & Rassmus-Gröhn, 2005).

4.1.2. The application

Only a small amount of projects have combined haptic and auditory feedback for exploring non-visual virtual 3D environments. Haptic devices, like the PHANToM Desktop, offer a natural and intuitive way of interacting with a virtual 3D environment. To enable visually impaired users to take advantage of this new interaction device, we need to overcome the limitations of unimodal haptic feedback. The addition of sound opens up many possibilities, but only a very limited amount of research is available on how sound can be used to overcome these problems. Therefore a qualitative investigation, on how sound can facilitate the localization of objects, was carried out.

Five different auditory displays were implemented in a virtual haptic environment developed for the purpose of this investigation. Each of the five different auditory displays used different combinations of auditory dimensions to represent spatial information. An auditory display maps information onto one or more perceptual dimensions of sound. Creating a complete and consensual classification is almost impossible due to the great number and ambiguousness of the perceptual dimensions. Though, a reasonable consensus exists for several basic dimensions. Most books on sound perception differentiate *pitch*, *loudness* and *timbre* (Goldstein, 1999; Sekuler & Blake, 2002). While pitch and loudness are two well-defined concepts, timbre seems to be a collection of everything not defined by pitch and loudness (Brewster, 1994), including, for example, roughness and brightness. Two generally used additions to these three dimensions are *duration* and *location*. Also, when sounds are played over a period of time, various other dimensions can be discriminated. This study considered the *start*, *tempo*, *dynamics*, and *composition* of sounds (Figure 21); the start of a sound refers to its origin in time, tempo refers to the time between sounds, dynamics refer to the perceived changes between sounds over time, for example, a rising pitch or decreasing loudness, and composition refers to the combination of different sounds.

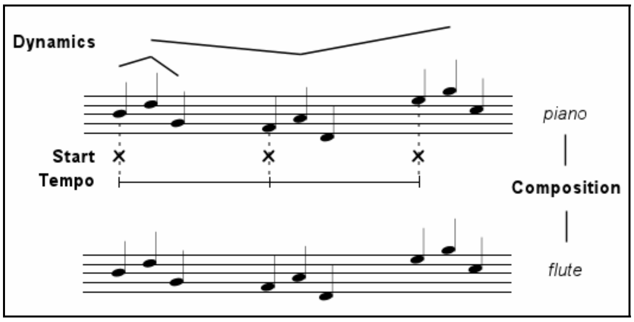

Figure 21. Dimensions in a sound sequence.

The Virtual Environment

The haptic 3D virtual environment designed for this investigation was bounded by a box measuring 100 x 100 x 80 millimeters (see Figure 22). The objects in the virtual environment were cubes measuring 8 x 8 x 8 millimeters.

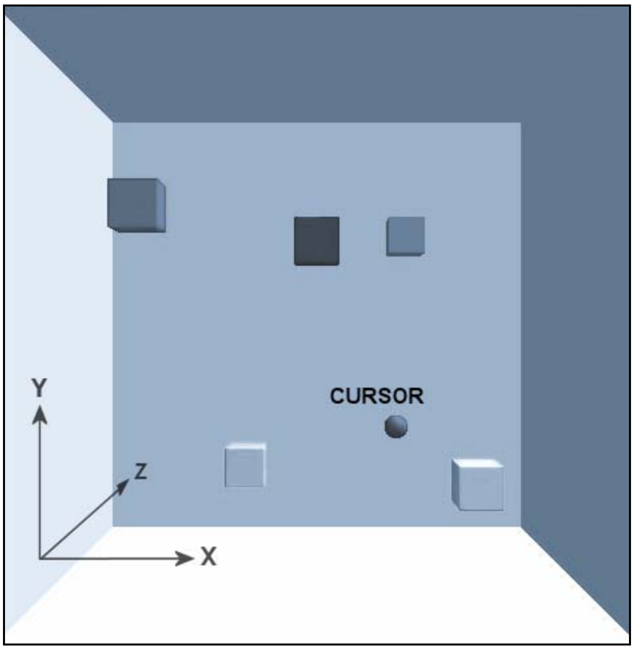

Figure 22. Visual representation of the virtual environment.

The cubes could be moved by making contact with the cursor and subsequently pressing the button on top of the stylus. The cubes could not rotate and no gravity was simulated in the virtual environment. The cubes stayed in the exact position where they were created by the system or placed by the user. The initial position of a cube was randomly generated by the system.

AD1: Coordinate-based Auditory Display

The first of the five auditory displays was the only auditory display that sonified the absolute position of objects. The auditory presentation only changed when objects were moved and not when the cursor was moved. A virtual auditory scanner, represented by a low pitched buzzing sound, continuously swept from left to right through the virtual environment in approximately 2.5 seconds. When the scanner passed an object, a bell sounds was played to represent that object. Characteristics of the bell sound were transformed to represent the position of the object. Five panning values (left, left/center, center, right/center, and right) were mapped to the x-coordinate. Pitch was mapped to the y-coordinate and a cross-fade from an original to a high-pass filtered sample was mapped to the z-coordinate. The high-pass filter made the bell sample sound thin and subdued. Volume, as used in the experiment of Mereu and Kazman (1996), was intentionally not chosen to represent the z-coordinate, because it interfered with panning.

AD2: Nearest Path Auditory Display

The second auditory display sonified the direction and distance from the cursor to the nearest object. Again a bell sound was used to represent the objects. When a different cube became the closest cube to the cursor, a short double bell sound was played to indicate that a new path was being sonified. Discrete transformations were used to represent the path from the cursor to the object. Panning (left, right and center) was used to indicate if the cursor should be moved to the left, to the right or kept in the same horizontal position. Pitch (high, low, and medium) was used to indicate if the cursor should be moved upward, downward or kept on the same height. And a high-pass filter (on or off) was used to indicate whether the cursor should be moved backward or not. The bell sounds were played repetitively. The time between repetitions decreased logarithmically as the distance from the cursor to the object decreased. This created an auditory space with a low resolution far from the objects and a high resolution close to the objects, as suggested by the investigation of Pitt and Edwards (1991).

AD3: Virtual Microphone Auditory Display

In the third auditory display the haptic cursor acted like a virtual microphone through which the user can listen to the virtual room, similar to an auditory display investigated by Magnusson and Rassmus-Gröhn (2005). The objects, represented by different drum loops, were positioned in 3D auditory space corresponding to their position in the virtual environment. This was done by applying Head Related Transfer Functions (HRTFs) in real-time to the drum loops. A HRTF imposes an interaural time difference and integrates absorption, reflection and diffraction effects caused by the head, torso and pinna of a general user. The HRTFs, used in this investigation, were programmed into the hardware of the soundcard. The 3D auditory space was set-up in such a way that objects close-by were heard loudly, objects at a small distance were still audible and objects far away could not be heard. The drum loops used were selected in advance based on the quality of synthesis, broadness of spectrum and distinctiveness of rhythm.

AD4: Prioritized Notes Auditory Display

The fourth auditory display was based on the idea that objects close to the haptic cursor are of larger interest than distant objects. This auditory display represented each object with a different piano note. Together these notes formed a pleasant sounding chord. Distant objects were sonified simultaneously once or twice every 1.8 seconds depending on their distance. Objects close to the cursor were sonified independently of the tempo of the far objects. In this manner, objects far from the cursor blended together, and objects close to the cursor stood out. For objects close-by, the time between repetitions decreased logarithmically as the distance from the cursor to the object decreased identically to how distance was represented by the nearest path auditory display. Similarly to the virtual microphone auditory display, all piano notes were positioned in 3D auditory space. The 3D auditory space was however modified, in such a way that all notes played were audible, even when the distance between the cursor and an object was at its maximum.

AD5: Pointing-based Auditory Display

The fifth auditory display transformed the stylus of the PHANToM Desktop into a pointing device. The user had to point the stylus through the virtual environment and when the stylus was pointing at an object, a bell sound was played repetitively. The bell sound was transformed, similarly to the coordinate-based auditory display, to represent the absolute position of the object. The time between repetitions, similarly to the nearest path auditory display and the prioritized notes auditory display, decreased logarithmically as the distance from the cursor to the object decreased. When not pointing directly at but close to an object, a sound was played that indicated how users should rotate the pen to point directly at the object. This guiding sound can be described as an ‘ahhh’ sound. Three pitch values (high, middle, and low) were mapped to rotation around the x-axis, also called pitch rotation. And three panning values (left, center, and right) were mapped to rotation around the y-axis, also called yaw rotation. The guiding sound was designed to make it easier to find objects by giving auditory feedback on a larger area than the cube itself.

The five auditory displays compared

Each of the five auditory displays relied on different combinations of auditory dimensions to represent spatial information. The mapping of information to auditory dimensions is presented for each auditory display in table 1. This table also shows if, and by which auditory dimensions, objects could be identified and discriminated from each other. Since each auditory display was based on a different principle and relied on different auditory dimensions, it was expected that some auditory displays would be more intuitive to use than others. However, all of them were expected to be learned within a relative short amount of time. When learned, it was expected that objects were detected quickly with auditory displays that presented multiple objects at once; AD1, AD3 and AD4. With AD2 and AD5, where only one cube could be head at a time, it was expected that it would be easy to move towards a detected object, but detecting them would be more difficult.

The haptic magnetism tool

Developers of haptic applications generally resorted to haptic tools or widgets to aid target location. Crosshairs, magnets, and haptic cursor enlargements are examples of suggested haptic tools (Sjöström, 2001b). It is still unknown how auditory displays compare to haptic tools, when both are designed for facilitating target location. It is also still unknown whether an auditory display can be usefully combined with a haptic tool. To get a general answer to those questions, a form of haptic magnetism was implemented along with the auditory displays. With haptic magnetism enabled, objects exerted small attractive forces upon the haptic cursor, pulling the hand of the user towards the objects. The added forces were intentionally kept small, to avoid interference with the movements of the user.

4.1.3. Method

Experimental design

This research investigates how auditory feedback can support the detection and location of objects in a haptic virtual 3D environment. The five different auditory displays and one haptic tool, implemented in a haptic virtual environment, were evaluated in two qualitative user studies. In the first user study, participants used all five auditory displays and the haptic tool. In the second user study, participants used only two auditory displays and the haptic tool.

Both user studies had, for each auditory display, an introductory phase, a question phase, a task phase and an interview phase. During the introductory phase, participants got instructions and tried out a randomly selected auditory display. During the question phase, participants used this auditory display to answer questions about the global layout of the objects within the virtual environment. During the task phase, participants were asked to move cubes to various places within the virtual environment. During the interview phase, participants were asked to comment on the auditory display they just used. In the first user study participants were additionally asked to rank the auditory displays on usefulness and pleasantness. When there still was time available after exploring and, in the first user study, ranking the auditory displays, the haptic tool was evaluated.

The whole procedure was captured on video. The following information was extracted through video analysis: (1) the percentages of incompleteness and the time to complete questions and tasks, (2) remarkable events, abilities and disabilities of the participants, and (3) relevant remarks articulated by the participants during the evaluations and the interview phase. The rankings on usefulness and pleasantness were immediately written down when given by the participants during the first user study.

We will use the quantitative data to indicate how well the different auditory displays performed. The qualitative data will be used to identify the characteristics of each auditory display and the haptic magnetism tool that are responsible for its performance.

Participants

One female and two male participants, ages 30, 54, and 55, took part in the first user study. All three participants were born with a visual impairment, and have been legally blind for a minimum of 10 years. One of the three participants reported some previous experience with a haptic device similar to the PHANToM Desktop. Four additional participants, two male, and two female took part in the second user study. Except one female, who still had low vision (10/100 unaided), all four participants were legally blind from birth. Their age varied between 26 and 59. One of the four participants reported some previous experience with a haptic device similar to the PHANToM Desktop.

Apparatus

The virtual environment was run on a modern desktop computer with Microsoft Windows XP Professional installed as operating system. A Sound Blaster X-Fi XtremeMusic from Creative was installed and used to play audio. Its Creative MultiSpeaker Surround 3D feature (CMSS-3D) was used to produce 3D Surround Sound over a set of Sennheiser headphones. The PHANToM Desktop haptic device of SensAble Technologies was used as a 3D input device that simultaneously gives kinesthetic feedback.

Figure 23. Video-frame of a mixed audio/video-file.

A digital video camera and web cam were used to capture the upper body of the participant, his or her use of the haptic device, a visual representation of the virtual environment, the audio generated by the computer and the conversation between the participant and the observer. The video and audio streams were mixed together in a single audio/video file (Figure 23).

Procedure User Study 1:

Haptic Introduction

To get familiar with haptic interaction, participants first explored, for approximately two minutes, a virtual haptic environment containing a large pyramid shape and a large doughnut shape. Participants were asked to find and identify these two shapes without the help of an auditory display or the haptic magnetism tool.

Auditory Display Introduction

Next the participant got, in random order, an introduction to one of the five auditory displays. After the introduction, participants used the auditory display in a virtual environment containing two cubes. During their exploration, participants received additional instructions on how to use the auditory display. Participants were allowed to ask questions about the auditory display at all times. After participants could successfully locate and manipulate the two cubes in the virtual environment, the evaluation was continued.

Question phase

After the introduction on an auditory display, an environment was started with a randomly three to eight cubes present. Participants were asked to describe the global position of the cubes, or how many cubes were present in the virtual environment. Two of these questions were scheduled per auditory display.

Task phase

After the question phase, a new environment was started with five cubes present. The participants were asked to move the cubes to a random wall of the room, for example left, or to move half of the cubes to a random wall and the other half to the opposite wall, for example front and back, or to build a tower of the cubes in the middle of the floor. When the participants could not find the next cube within a reasonable time, the task was ended by the observer. Three tasks were scheduled per auditory display, but generally only one or two tasks were completed due to time constraints.

Interview phase

After completing the task phase, participants were asked several questions about their experience with the auditory display they just used, for example, "Can you tell me why the tasks were easy or difficult?"

Ranking phase

The process of auditory display introduction, question phase, task phase and interview phase, was repeated until all five auditory displays had been evaluated. A ten minutes break was held between evaluating the third and fourth auditory display. After evaluating all five auditory displays, participants ranked the auditory displays on usefulness and pleasantness.

Haptic tool

When there still was time remaining after testing all five auditory displays, participants evaluated the haptic magnetism tool. Haptic magnetism was enabled independently or combined with one of the auditory displays. The participants were given questions and tasks identical to the ones asked during the evaluations of the auditory displays.

The total procedure was completed within approximately two hours.

Procedure User Study 2:

The start of the second user study was identical to the first evaluation. But instead of randomly testing five auditory displays, two out of four auditory displays were pre-selected. The coordinate-based auditory display was excluded from the second user study, because participants could hardly use this auditory display to come into contact with a cube. In the second user study, the question phase was swapped with the task phase to see if experience from the task phase increased performance during the question phase. Six tasks and two questions were scheduled per auditory display in the second user study. Participants were given a break of ten minutes between testing the first and second auditory display. When time remained, participants also evaluated the haptic magnetism tool independently and combined with the two previously evaluated auditory displays.

The total procedure was again completed within approximately two hours.

Description of External Memory Aids

4.1.4. Results

A total of 33 questions on global layout were answered during the first user study; 26 while evaluating an auditory display and 7 while evaluating haptic magnetism. In addition, a total of 26 tasks were completed; 18 while evaluating an auditory display and 8 while evaluating haptic magnetism. During the second user study, a total of 15 questions were answered, all while evaluating an auditory display. In addition, a total of 55 tasks were completed; 46 while evaluating an auditory display and 9 while evaluating haptic magnetism.

This section reports about two types of results, quantitative and qualitative. The quantitative results indicate how well the different auditory displays aided participants in answering questions and completing tasks. It does not **prove** that one auditory display is 'better' than the other. For haptic magnetism, only qualitative results are reported, because too few trials were completed for a quantitative comparison.

Quantitative results

Percentage of incompleteness and time taken per object

The percentage of incompleteness was determined by the relative number of cubes that the user did not count, locate, or move. For questions, the time per object was determined by dividing the time it took to complete a question by the number of cubes present. For tasks, the time it took to locate and move a single object was directly measured during the video analysis. For this reason, more observations are available on time per object than on percentage of incompleteness.

Figure 24 summarizes the quantitative measurements for both user studies. The figure is split into four graphs, one for the percentages of incompleteness of the first user study, one for the percentages of incompleteness of the second user study, one for the time taken per object during the first user study, and one for the time taken per object during the second user study. Each graph is divided in a question section and a task section. Per auditory display, a bar is displaying the percentage of incompleteness or the time taken per object, is given. The absolute percentages and times are displayed above each bar. The standard deviations are displayed in parenthesis next to the absolute percentages and times. The numbers of observations are given under each bar.

In the graphs of the second user study, the first auditory display is missing, because the coordinate-based auditory display (AD1) had been excluded from the second user study.

Answered Questions

During the first user study, the participants answered the questions quickest and most accurate while using the coordinate-based auditory display (AD1). With this auditory display the participants did not have to explore the environment with the haptic device answer the questions. Participants also did well, though slower, using the virtual microphone auditory display (AD3). Participants had a little more difficulty answering the questions using the prioritized notes (AD4) and pointing-based auditory display (AD5). The least useful auditory display for answering the questions was the nearest path auditory display (AD2); the average percentage of incompletion was 83% (figure 5).

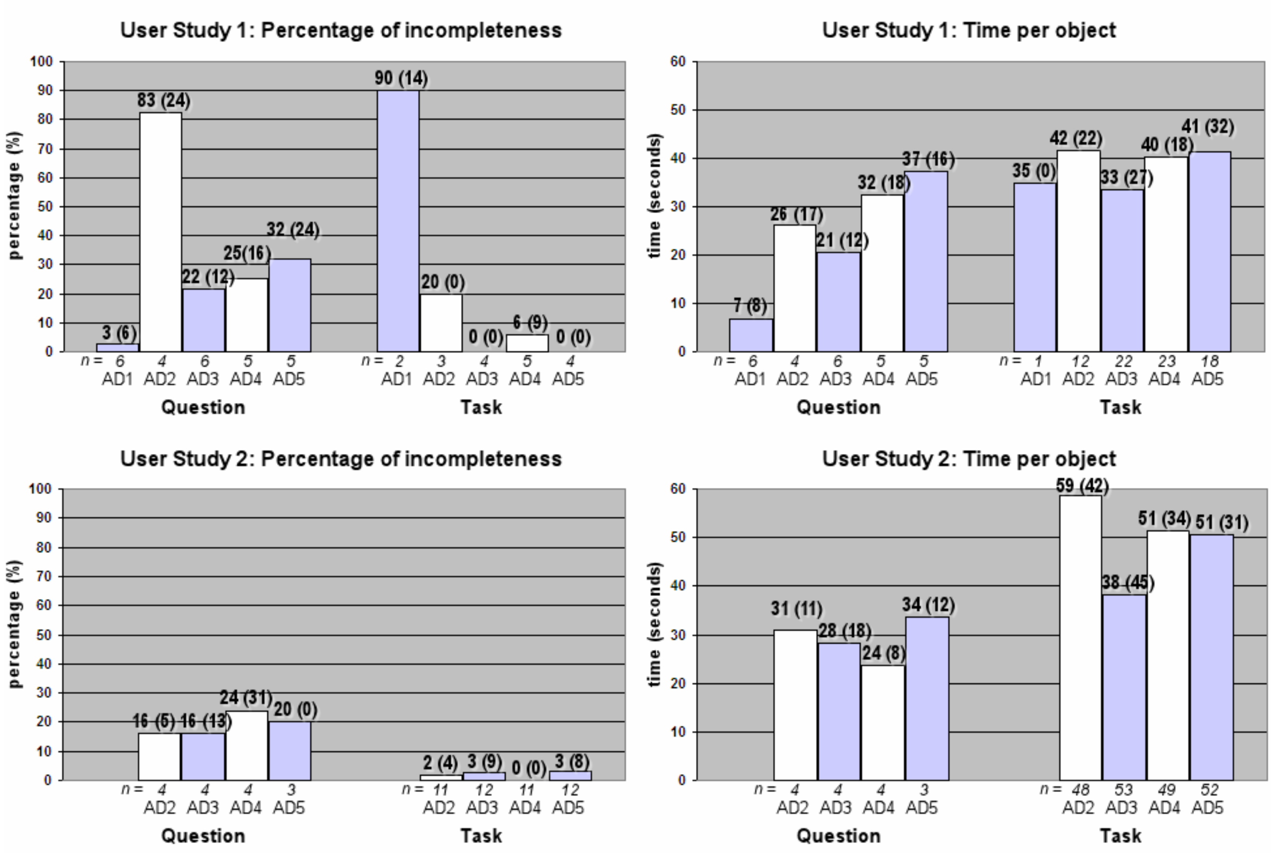

Figure 24. Percentage of incompleteness and time per object values.

During the second study, the participants had more experience using the auditory displays before answering the questions. The differences in percentage of incompletion and time per object became very small between the auditory displays. Questions were answered accurately, regardless of which auditory display was being used, even with the nearest path auditory display (AD2).

Completed Tasks

The participants were unable to successfully complete the tasks while using the coordinate-based auditory display (AD1). Only one cube was successfully found and moved, illustrated by the 90% incompleteness bar in Figure 24. For this reason the coordinate-based auditory display was excluded from the second user study. With the four other auditory displays, almost all tasks were successfully completed. No significant differences in the percentages of incompleteness between these auditory displays were found.

During both user studies, the participants located and moved the cubes fastest while using the virtual microphone auditory display (AD3). Second, in time per object, came the prioritized notes auditory display (AD4) and the pointing-based auditory display (AD5). Their average performance was similar to each other, despite being based on very different principles. Out of the dynamic auditory displays, hereby excluding the coordinate-based auditory display (AD1), the participants took the longest time to locate and move the cubes using the nearest path auditory display (AD2).

The average time per object was higher during the second user study compared to the first user study. During the second user study, participants had to complete the relatively difficult task of building a tower from the cubes fourteen times, compared to two times during the first user study. A relatively large increase in time per object for the nearest path auditory display (AD2) was noticed between the user studies. The time per object increased from 42 seconds in the first user study to 59 seconds in the second user study. This increase can be explained by the extra time given to the participants in the second user study, when having trouble finding the next cube. Logically, this also led to a decrease in the percentage of incompleteness.

Perceived usefulness and pleasantness

During the first user study, the three participants ranked all five auditory displays on usefulness and on pleasantness. Rank 1 was reserved for the most, and rank 5 was reserved for the least useful or least pleasant auditory display. The average rankings for each auditory display are given in table 2.

	Perceived Usefulness	Perceived Pleasantness
AD1: Coordinate-based	3,8	4,2
AD2: Nearest Path	3,8	3,5
AD3: Virtual Microphone	1,7	1,0
AD4: Prioritized Notes	3,7	3,3
AD5: Pointing-based	2,0	2,3

Table 2: Avg. rankings on usefulness and pleasantness.

The virtual microphone and pointing-based auditory display were ranked as the two most useful auditory displays. While the use of the pointing-based and prioritized notes auditory display resulted in similar percentages of completion and times per object, the pointing-based auditory display was clearly perceived as more useful. When rating the auditory displays on pleasantness, one person remarked that it is more important that sounds are informative. Nonetheless, participants unanimously ranked the virtual microphone auditory display as the most pleasant auditory display to listen to. Most appealing to the participants was the sound of the drum loops and the naturalness of the auditory display. The pointing-based auditory display was ranked second on pleasantness. One of the reasons mentioned by a participant, on why this auditory display sounded pleasant, was that the auditory display could be made silent by pointing the stylus away from the cubes. While the coordinate-based auditory display was ranked as the least pleasant sounding auditory display, none of the three participants reported it as annoying.

Qualitative results

AD1: Coordinate-based Auditory Display

Counting the number of cubes in the virtual environment was easy for the participants. They could quickly identify the start and end of a sequence, by the low buzzing sound being panned from left to right, and by the recurring rhythm of the bell sounds. When two or more cubes had similar horizontal positions (x-values), participants were, generally speaking, still able to distinguish the different bell sounds. When asking to describe the global position of the objects, participants always started by describing the horizontal positions of the cubes. The temporal aspect of the virtual auditory scanner made it easy to estimate the horizontal position, as shown in the following participant remark; “If you only want to know if it is to the left, right or in the middle then I like the scanning method, because it moved from left to right. It was harder to listen to the pitch and the bells to find out if it was back or forward or high or low”. Participants indeed had more trouble estimating the vertical position of the cubes. It was easier to relate the vertical position of one cube to another cube than to relate the vertical position of one cube to the room. Estimating the front-back position of a cube was the most difficult for the participants. In several cases no estimation at all was given for the front-back position; “It is hard to say if they are close to the front or to the back wall”, was one of the remarks. In general, participants could only distinguish depth between cubes being to the far back of the room and cubes being to the very front of the room. During the user studies, it was also noticed that the high-pass filtered sample interfered with the pitch transformation of the bell sound. This led to several wrong estimations on the vertical position of cubes.

The participants were unable to complete a single task using the coordinate-based auditory display within a reasonable time. Participants did have a general idea of where cubes were located, but the cubes could not be found with the haptic device. Participants were moving the cursor in large scanning motions, or along the walls, in the hope of bumping into a cube. During the haptic exploration, participants were relying more on luck than on the information from the auditory display.

AD2: Nearest Path Auditory Display

It took participants quite a while to understand, and correctly use, the nearest path auditory display. The sound played when a different cube became the closest cube to the cursor was often misunderstood. For example, some participants mistakenly thought that the cursor was very near to an object when hearing this sound. Participants asked a variety of questions about this sound and, for three of the five participants who used this auditory display, its meaning had to be explained several times. It was also observed that participants had trouble distinguishing the different cubes; some cubes were counted twice and some cubes skipped, because participants mistakenly thought they had already counted them. During the second user study, participants did a lot better in answering the questions. They had learned how to find the cubes with the auditory display, and went to touch all cubes in order to count them and describe their position.

Both in the first and in the second user study, tasks took the longest time to complete with the nearest path auditory display compared to the other three dynamic auditory displays. Nonetheless, participants were able to successfully accomplish all tasks when given enough time. Participants were quick in using tempo as an indicator of distance. It took longer to learn and use the discrete directional cues. To four of the five participants who used this auditory display, the meaning of the directional cues had to be explained not only during the introduction, but also once or twice during the evaluation of the auditory display. Two participants reported that they wanted a more distinct difference between the levels of the directional cues.

AD3: Virtual Microphone Auditory Display

Participants needed little explanation to start working with the virtual microphone auditory display. Participants could answer the questions using the virtual microphone auditory display well, though they needed time to explore and find the different drum loops in the virtual environment. Occasionally a cube was missed. Some participants did not always touch the cubes to count them or describe their location. This did almost never happen with the other auditory displays or with the haptic magnetism tool. In approximately six cases, when a cube was near to a wall, the wall was mistaken for the cube.

Participants reported that they could easily distinguish most of the different drum loops. But when having to compare drum loops from memory it was more difficult. One participant, for instance, had to start over counting, because he forgot which sounds he had already heard. Some participants learned the drum loops after a few tasks, and reported during a task to be searching for a specific drum loop, for example, “I knew what sound I was looking for, for the fifth one”. Other participants, however, reported they wanted the drum loops to sound more differently. It was also noticed that some drum loops could be heard over greater distances, which annoyed one of the participants. She reported that, because of this, it was harder to find drum loops that did not carry as far.

With the virtual microphone auditory display the cubes were quite easily discovered. But from discovering to coming in contact with the cube took longer. This was expressed by one of the participants: “I hear it all right, but to get it with touch is more difficult”. Participants noticed when they got nearer to a cube, but the direction to move in was more difficult. Hearing if a cube was to the left, middle or right of the cursor was relatively easy with 3D Surround Sound. Hearing if a cube was to the front-back of the cursor, or above-below the cursor was difficult. Nevertheless, one participant reported he could hear and use these cues from the 3D Surround Sound.

AD4: Prioritized Notes Auditory Display

It took some time for the participants to get a sense for the tempos and the blending of the notes, but during the reflection phase the prioritized notes auditory display was reported as being easy and intuitive to use. Participants were not able to distinguish all individual notes when played in a chord. This made it necessary to explore the virtual environment in order to identify all the different cubes. Notes were easy to compare directly, but when having to compare them from memory, it became more difficult. Participants were occasionally unsure if they had already come across the cube they just found. A note played in an asynchronous faster tempo was easily distinguished from the other notes. When participants noticed a cube, they were relying more on the distance information, than on the directional information provided by the 3D Surround Sound. Participants were observed moving the cursor back and forth in multiple directions, listening to the change in tempo and then finding the right direction to move closer to a cube.

AD5: Pointing-based Auditory Display

Generally, participants needed some time and extra explanation to understand and use the pointing-based auditory display, but afterwards several reported this auditory display as being easy to use. Participants had to actively scan the virtual environment room to find the cubes. No major strategy could be discovered in the scanning motions of the participants. One person remarked: “Some times you find them and sometimes you don’t. I don’t know if there is any reason for it”. A few participants tried to scan the virtual environment using a probing motion, or forward longitudinal motion, which is useless for discovering new objects. One participant rotated the stylus sideward when scanning the virtual environment and was able to successfully locate several cubes. Without a structural strategy participants were still able to scan a great deal of the virtual environment and find most cubes. Some places in the front of the virtual environment were harder to point at, because the stylus needed to be pulled to the very front of the workspace or rotated into an unusual angle. Participants also had trouble differentiating cubes that were positioned close together or in one line.

The guiding tone, when pointing near a cube, was not used in the same manner by all participants. Most of the participants began by using the guiding tone as a border around a cube. As time progressed, two of the five participants started to use the directional cues given by the guiding tone. This was demonstrated in the way participants closed in on a cube when they located one. Since moving the stylus in a straight line towards a cube was not easy, participants occasionally had to re-orientate the stylus to point directly at a cube again. While three participants pulled back on the stylus, two started to use the directional cues to rotate the pen, to point directly at the cube again.

Haptic Magnetism

In total, 24 trials (7 questions, 17 tasks) were carried out with haptic magnetism enabled; 8 trials with haptic magnetism independently and 16 trials with haptic magnetism combined with an auditory display. In general, participants were positive about the haptic magnetism. One of the participants remarked for example, “I liked it, it was easy”. Participants were particularly positive when haptic magnetism was combined with one of the auditory displays. Several positive remarks were made about using an auditory display to get near to a cube and then letting the haptic magnetism pull you in. For example, “When you hear you are not so far you let the pen show you were the object is” and “I heard the sound getting stronger so I knew what direction to move, and then when I got close enough I got sucked in”. Three tasks were completed remarkably fast in which haptic magnetism was combined with the virtual microphone auditory display, for example, finding and moving five cubes to the left wall of the virtual environment in 30 seconds.

There were also problems when using the haptic magnetism tool. Without an auditory display, cubes can only be distinguished by their position. Participants, therefore, could not always remember if they had touched a particular cube already, especially when several cubes were located close together. Combining haptic magnetism with an auditory display did not always speed up a task either, for example, only being able to move three out of four cubes to the bottom of the virtual environment in a time of 215 seconds. At times haptic magnetism interfered with the desired movements of the participants. Participants remarked: “In this case I would probably appreciate not to have the magnetism interfere with my search” and “the magnetic force perhaps wants me in another direction than I want”. With haptic magnetism enabled, some participants also started to make more errors while trying to grab a cube in order to move it. One participant, for instance, failed to grab a cube, six consecutive times.

The haptic device

For all participants, the haptic device created a realistic haptic experience. Two of them even reported that they felt compelled to use their free hand to reach in and touch the virtual objects. The way of using the haptic device differed among the participants. Some were using the haptic device with confidence, making large and forceful movements, and others were using the haptic device in a more cautious way, making small movements and keeping the cursor in contact with a wall for continuous haptic feedback.

The stylus of the haptic device was not always grabbed in the intended way. Holding the stylus similar to holding a regular ballpoint is not self-evident. Four out of seven participants grabbed the stylus as one normally grabs a stick or a door handle. These four participants were instructed to hold the stylus in the intended way, but two of them, one immediately and one later on, switched back to their initial gripping style. The different grip did not seem to interfere much with the exploration. However, when moving the stylus to the top-back of the virtual environment, the haptic device produced a mechanical force that led to some confusion.

During the evaluations one participant experienced physical discomfort in the upper arm from using the haptic device. After half an hour this participant reported that using the haptic device was “tiring” and added: “I would suggest a support for your arm”. After the scheduled break in the middle of the user study the participant started supporting the right arm with her left arm. Soon after the break she reported: “my arm is hurting a little”. The position of the haptic device was changed in a way that the left hand could be used to explore the virtual environment. It was observed that using the left arm also caused strain. The participant did not seem to perform worse after switching arms, though in general, this participant had more difficulty answering questions and completing tasks.

4.1.5. Discussion

All results indicate that the virtual microphone auditory display was the best auditory display to be used in the, here explored, virtual environment. The principle of a virtual microphone is intuitive and is easy to implement in other virtual environments. The pointing-based auditory display also performed well. With this auditory display it was harder to detect objects, but, as expected, it was easy to move in and touch the objects. A pointing-based auditory display could be very useful in virtual environments with larger objects, because they are easier to find. The prioritized notes auditory display preformed equal to the pointing-based auditory display. The addition of tempo as an indicator of distance should have made this auditory display outperform the virtual microphone auditory display, which it clearly did not. This method needs further development, before it has advantages over the virtual microphone method. The nearest path auditory display also needs further development before it becomes a useful auditory display. The coordinate-based display is not suitable as an auditory display that facilitated object localization.

Designing an auditory display

All sounds in an auditory display should have a purpose. Too many sounds will clutter the auditory space and, therefore, sounds should be presented parsimoniously. The information conveyed in an auditory display should also be easily understood by the user-group. This was, for example, not the case with the nearest path auditory display. A sound that is informative to the user is also less likely to be perceived as annoying. Also important to the pleasantness of an auditory display is the ability to make it silent.

Most perceptual auditory dimensions have a relative low resolution. Still, a lot of information can be communicated by using characteristics of sound efficiently. It is known, and seen during the evaluations, that tempo and pitch have a relatively high resolution. Panning and a cross fade between an original and a high pass filtered sample had a relatively low resolution. The necessary resolution needed for a variable should be assessed and then mapped to a fitting aspect of sound. Care should be taken with using different characteristics of sound, as they could interact. The lack of orthogonality is a known problem for designing auditory displays. An example of two aspects of sound interacting can be seen in the coordinate-based auditory display. In this auditory display the cross fade with a high pass filtered sample interfered with the change in pitch.

One should also look for intuitive representations of variables. Tempo, for instance, is an intuitive characteristic of sound for representing distance. Panning is an intuitive characteristic for representing left, middle and right. When no intuitive characteristic can be found, one should use a characteristic that possesses easy distinguishable values/levels to which a variable can be mapped. Users should be given time and explanation to learn and internalize this mapping. Because no absolute values are perceived through characteristics of sound it is advisable to present reference values.

All possible events in a virtual environment should be identified and assessed on importance to the user. Providing feedback on important events confirms the actions of the user and makes it easier to understand what is happening. For example, by sonifying the event of coming into contact with a cube, users would less likely identify the wall as a cube.

Detecting objects in a virtual environment

To be able to detect an object in a virtual environment, it should somehow be distinguishable. This can be done in many ways. Sonifying objects separately in time is a good method as demonstrated with the coordinate-based auditory display. The different drum loops are another good example. The drum loops could be distinguished by their rhythm and composition of instruments. In the virtual microphone auditory display, the drum loops propagated in spherical shapes from objects. The distance the drum loops propagate (the size of the spherical shapes around the objects) is important to the functionality of the auditory display. When the distance over which drum loops are carried is large, one can identify objects from afar. On the other hand, more sounds are bound to reach the cursor, which complicates the perception of the individual sounds. Furthermore, the resolution of the distance information will decrease, since the range is spread over a larger distance. A third good example of how objects can be distinguished is the discriminative tempo assigned to a sequence of unique notes in the prioritized notes auditory display. A note played in a fast tempo asynchronous to the tempo of the other notes, which also changed fast in relation to the distance between the cursor and the sonified object, was easily distinguished from the other notes. The blending of notes in the prioritized notes auditory display was intended to give an impression of the amount and position of objects that were not close to the cursor. Unfortunately, participants were not able to take advantage of this information during the relatively short evaluation. It ought to be seen, whether more experience with the prioritized notes auditory display will change this.

Locating objects in a virtual environment

Results from the coordinate-based auditory display show that feedback on the position of an object, relative to the cursor, is necessary to successfully move towards and make contact with an object. The path to an object is determined by the distance and direction from the cursor to this object. Again there are many ways of sonifying the variable representing distance, and the vector representing direction, in an auditory display. See table 1 in the introduction section, for how the different auditory displays sonified the relative distance and relative direction. Both tempo and 3D Surround Sound, as used in the auditory displays, were intuitive measures of distance. Changes in tempo and 3D Surround Sound were easily perceived and associated with relative distance. Since the direction from the cursor to an object is not a singe variable, but a vector with three components, the direction is more difficult to sonify. In the nearest path auditory display, the direction vector was represented by three different characteristics of sound. To understand the relative direction the user had to learn eight different sounds (up, center, down; left, center, right; forward combined with centre, and backward). Participants were not able to fully internalize these directional cues during the evaluations. The directional cues from the 3D Surround Sound should not take long to understand, since these cues are continuously used during daily life. However, the cues produced by the 3D Surround Sound were not effective enough to represent direction well. A completely different way of representing the relative direction to an object, is giving auditory feedback on the orientation of the stylus when it is pointing at an object, as implemented in the pointing-based auditory display. Kinesthetic feedback and auditory feedback are thus combined to represent a direction. Though the principle of pointing is based on visual cues, the visual impaired participants had little problems understanding this auditory display. The pointing principle was by far the most effective way of representing direction during the evaluations.

Haptic magnetism and auditory displays

Haptic magnetism is easy to understand and intuitive to use. It does not give a continuous overview of a virtual environment, but it makes the user find an object when the cursor comes close to it. It is questionable whether haptic magnetism does a good job in representing distance and direction. Users did not take time to feel in which direction the force was pulling and then move along that direction, but when users felt a force, they let the force drag the cursor to the object. Haptic magnetism was easy to combine with an auditory display. Haptic magnetism is expected to combine well with an auditory display that is good in getting the cursor close to an object, because the magnetism will then pull the cursor towards this object. Care should be taken, however, with implementing haptic magnetism. The forces involved should be rather weak, to limit interference with intended cursor movements and to limit interference with perceiving other forces, for example, object contact forces. As haptic magnetism is only one of several haptic tools for finding objects, also other haptic tools should be considered for combination with an auditory display.

Some final remarks on the haptic device

During the evaluation it was found that the intended grip of the stylus of the PHANToM Desktop haptic device is not natural for visual impaired persons. The natural grip was not fully supported by the haptic device, as mechanical forces appeared in a few specific places of the workspace. SensAble technologies, the producer of the PHANToM Desktop haptic device, also supply a thimble, which is interchangeable for the stylus. A comparison between the thimble and stylus should clarify which one is better suited for use by visual impaired person.

Use of a haptic device can cause strain in the upper arm. One participant reported strain in the upper arm during the user studies. The strain was caused by having to carry the weight of ones arm and some extra weight from the haptic device. A support for the elbow of the arm using the haptic device is expected to solve this potential problem. An extensive research on the ergonomics on the PHANToM Desktop and comparable haptic devices is advised.

4.1.6. Guidelines

The following guidelines have been extracted from this work:

- A local sonification model where the listening point is located at the cursor/proxy is necessary when supporting object localisation
- A global sonification model supports the overview of a complete auditory space
- Using tempo for distance and stereo for left and right is intuitive and effective ways of representing directions
- When no intuitive characteristic can be found, one should use a characteristic that possesses easy distinguishable values/levels to which a variable can be mapped
- All possible events in a virtual environment should be accounted for and presented in a clear and unambiguous way to the user

- Using pointing as a way of finding and locating objects is easy to use, but takes some time so get used to and learn

4.1.7. Future developments

The use of audio to support a virtual haptic 3D display shows many promises, but needs to be further investigated. The use of 3D audio, however intriguing, still a lot of problems, and as shown in this study the only variables that really worked were loudness to represent distance, and stereo panning to represent left-right (the x-axis). Apart from that, other more effective ways of sonifying the remaining variables needs to be found.

The use of gestures in a haptic environment is something that is presently being investigated by many different researchers, and in this study the simple act of pointing shown a novel and effective way of locating objects in a virtual 3D space. The relationship between gestures and auditory displays needs to be further investigated.

5. Demonstrators

Demonstrations of several of our work package prototypes are available on the MICOLE project website (http:// micole.cs.uta.fi).

6. Other Activities

The partners in the WP have published 25 papers on the work in the last year. Glasgow's work on the Maze application was accepted for the Royal Society Summer Exhibition, a prestigious science exhibition in the UK. We showed our work to over 4000 people at events in London and Glasgow.

UGLAS also organised the First International Workshop on Haptic and Audio Interaction Design at Glasgow in August/September. Over 50 people attended from all over the world, representing universities and industry. Much of the work of MICOLE partners were presented, along with demonstrations or our applications. For details see www.dcs.gla.ac.uk/~mcgookdk/multivis/workshop.html. The proceedings were published in the Springer LNCS series.



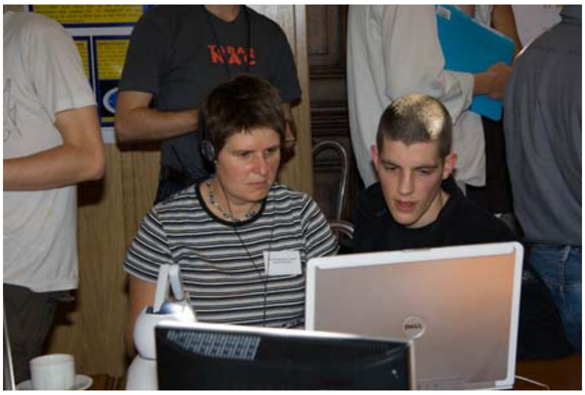


Jansson and Brewster have taken the work done in the WP and fed it into the new ISO guidelines on tactile and haptic interaction. Wall and Brewster at UGLAS edited a special issue on Haptic Interfaces and applications that appeared in the Springer journal Virtual Reality (Vol 9, issues 2-3). This featured work from some of the MICOLE partners.

Crossan publicised UGLAS and the whole project's work on the VIP On Air radio station for blind and visually impaired people (www.viponair.com: VIP ON AIR is Europe's first radio station for blind people, producing programming made by and for blind people. The station broadcasts news and information - such as daily newspapers - which can be difficult for a blind person to get hold of immediately).

7. Guidelines

As with our previous deliverable (D6) we have developed some preliminary guidelines from the research. These will be formalised in WP5 (there is a meeting to discuss this on 5-6 February 2007) but we present a summary of those related to external memory aids.

7.1. Force-feedback external memory aids

7.1.1. Design

- We recommend constant type forces, possibly with some 1/sqrt(r) like snapping behaviour in the vicinity of a point to make sure users actually reach it. This snapping has to be weighed against the possible disadvantage of interfering with user exploration close to the points.
- The initial force must be possible to feel, but it cannot be too strong. For short distances some increase in the force is acceptable – but this increase should not be too strong. These considerations point towards the type of force recommended above.
- Different users need different strengths of the forces, so it should be possible to adjust the strength of the force individually.
- To avoid vibration the force needs to be continuous towards r=0, but it is ok to use a force curve that is not smooth everywhere (at least for the tasks in the present test).

7.1.2. Multimodal line drawing

- Both positive and negative relief is possible to feel and to work with. Negative relief is preferred when working with simple line shapes. There are indications that negative relief shortens examination times.
- Both vertical and horizontal work areas can be used.
- Simple shapes can be recognized when they are kept in a specific context.
- The sound feedback can be used to get information about the program mode. The sound effects for the drawing tools are considered to add to the experience.
- The PHANToM Premium is hard to use especially for blind users who also don't handle an ordinary pencil very easily.

7.2. Trajectory playback

- Audio cues can be combined with trajectory playback to allow the users to segment trajectory playback events.
- Combining audio with haptic trajectory playback can significantly improve users' performance in learning and recreating trajectories then haptic trajectory playback alone.
- Pitch can be used successfully by a user for one dimension in a haptic trajectory playback learning task.

- When using pitch and pan for the audio playback, designers should be aware that it is difficult to concentrate on both pitch and pan at the same time. Pitch tends to be easier to distinguish and therefore dominates.

7.3. Tactile external memory aids

- Designers can increase the amount of information presented through a Tacton by increasing the number of independent parameters in the Tacton.
- Pattern shape and blink speed are two factors that can be used independently.
- For small arrays, using pattern size will affect the user's perception of the pattern shape. Smaller patterns are harder and take longer to distinguish.
- The range of blink speeds used in this study (0.04s per frame to 0.5s per frame) did not significantly effect the time it took the user to recognise the Tacton.

7.4. Audio external memory aids

- A local sonification model where the listening point is located at the cursor/proxy is necessary when supporting object localisation
- A global sonification model supports the overview of a complete auditory space
- Using tempo for distance and stereo for left and right is intuitive and effective ways of representing directions
- When no intuitive characteristic can be found, one should use a characteristic that possesses easy distinguishable values/levels to which a variable can be mapped
- All possible events in a virtual environment should be accounted for and presented in a clear and unambiguous way to the user
- Using pointing as a way of finding and locating objects is easy to use, but takes some time so get used to and learn

8. Published Papers

A range of papers have been published on the research done in the work package since deliverable D6.

8.1. UTA

T.V. Evreinova, G. Evreinov and R. Raisamo, Video as Input: Spiral Search with the Sparse Angular Sampling. In Proc. of ISCIS 2006 A. Levi et al. (Eds.), Istanbul, Turkey, 1-3 Nov., LNCS 4263, Springer-Verlag Berlin Heidelberg 2006, pp. 542 – 552. http://dx.doi.org/10.1007/11902140_58

Yfantidis, G. and Evreinov G. The Amodal Communication System through an Extended Directional Input. In Proc. of the 10th Int. Conf. on Computers Helping People with Special Needs. ICCHP 2006 Austria, Linz, 10-14 July 2006, LNCS 4061, Springer-Verlag Berlin Heidelberg 2006, pp 1079-1086.

Evreinova, T.G., Evreinov G. and Raisamo R. Evaluating the Length of Virtual Horizontal Bar Chart Columns Augmented with Wrench and Sound Feedback. In Proc. of the 10th Int. Conf. on Computers

Helping People with Special Needs. ICCHP 2006 Austria, Linz, 10-14 July 2006, LNCS 4061, Springer-Verlag Berlin Heidelberg 2006, pp. 353-360.

8.2. KTH

Crommentuijn, K., & Winberg, F. (2006). Designing auditory displays to facilitate object localization in virtual haptic 3D environments. In *Proceedings of the 8th international ACM SIGACCESS conference on Computers and accessibility, ASSETS 2006*. ACM Press, pp255-256.

8.3. Glasgow

Crossan, A. and Brewster, S. "Multimodal Trajectory Playback for Teaching Shape Information and Trajectories to Visually Impaired Computer Users". Submitted to ACM Transactions on Accessible Computing.

Crossan, A. and Brewster, S.A. *Two-Handed Navigation in a Haptic Virtual Environment*. In Vol II Proceedings of CHI 2006 (Montreal, Canada), ACM Press, pp 676-681.

Forrest, N. and Wall, S. ProtoHaptic: Facilitating Rapid Interactive Prototyping of Haptic Environments, In Proc of First International Workshop on Haptic and Audio Interaction Design, Volume II, 31st August – 1st September 2006, pp. 21-25.

Johnston, J., Crossan, A., and Brewster, S.A. . Multimodal Interaction and Control in a Virtual Environment, In Proc of First International Workshop on Haptic and Audio Interaction Design, Volume II, 31st August – 1st September 2006, pp. 13-16

Plimmer, B., Crossan, A. and Brewster, S.A., *Computer Supported Non-Visual Signature Training*, In Vol II Proceedings of the First International Workshop in Haptic and Audio Interaction Design (HAID2006). University of Glasgow, UK, 31st August - 1st September 2006

8.4. Forth

Anthony Savidis, Apostolos Stamou, Constantine Stephanidis "An Accessible Multimodal Pong Game Space", 9th ERCIM Workshop "User Interfaces for All"

8.5. Metz

Nicolas Noble and Benoît Martin. Shape discovering using tactile guidance. In EuroHaptics 2006, Paris, France, July 3-6, pp. 561-564, 2006.

Thomas Pietrzak, Isabelle Pecci, and Benoît Martin. Static and dynamic tactile directional cues experiments with VTPlayer mouse. In EuroHaptics 2006, Paris, France, July 3-6, pp. 63-68, 2006.

Pietrzak, T., Noble, N., Martin, B & Pecci, I. Evaluation d'un logiciel d'exploration de circuits électriques pour déficients visuel. In 3rd french meeting of young researchers on Human Computer Interaction (3èmes Rencontres Jeunes Chercheurs en Interaction Homme-Machine), RJC-IHM, 4 pages, published on CD-Rom and presented as poster, 2006.

Benoît Martin , Isabelle Pecci and Thomas Pietrzak. Angle Recognition Cues using a new API dedicated to the VTPlayer Mouse. In Proceedings of HuMaN'07 (Human Machine iNteraction Conference), Conférence Internationale sur l'Interaction Homme-Machine, TIMIMOUN, Algerian Sahara, March 12-14, 2007, 6 pages (accepted).

8.6. Uppsala

Jansson, G. & Pedersen, P. (2005b). Obtaining geographical information from a virtual map with a haptic mouse. Paper at the International Cartographic Conference (Theme “Maps for Blind and Visually Impaired”) 9-16 July, A Coruna, 2005, Spain. Available on conference CD-ROM.

Juhasz, I. (2006). A haptic mouse used for reading of virtual maps. In M.A. Hersh (Ed.), *Conference and Workshop on Assistive Technology for People with Vision & Hearing Impairments. Technology for Inclusion, Kufstein, Austria, 19th-21st July, 2006.*

Jansson, G., Juhasz, I. & Cammilton, A. (2006). Reading virtual maps with a haptic mouse: Effects of some modifications of the tactile and audio-tactile information. British Journal of Visual Impairment, 24, 60-66.

8.7. Lund

Kirsten Rassmus-Gröhn. Enabling Audio-Haptics. Licentiate Thesis, Certec 2:2006, Department of Design Sciences, Lund University, September 2006, Lund, Sweden

User evaluations of a virtual haptic-audio line drawing prototype Kirsten Rassmus-Gröhn, Charlotte Magnusson, Håkan Eftring Workshop on Haptic and Audio Interaction Design University of Glasgow, 31st August - 1st September 2006

Non Visual Haptic Audio Tools for Virtual Environments Charlotte Magnusson, Henrik Danielsson, Kirsten Rassmus-Gröhn Workshop on Haptic and Audio Interaction Design University of Glasgow, 31st August - 1st September 2006

Test of three different audio-haptic navigational tools Charlotte Magnusson, Kirsten Rassmus-Gröhn, Henrik Danielsson, Håkan Eftring 2nd Enactive Workshop May 25-27 2006 - McGill University, Canada

Two haptic-auditory applications for persons with visual impairments Kirsten Rassmus-Gröhn, Joakim Eriksson, Charlotte Magnusson 2nd Enactive Workshop May 25-27 2006 - McGill University, Canada

9. Conclusions

This deliverable has discussed the design and evaluation of external memory aids. These allow users to mark points of interest and to re-find them. They can be a powerful way of reducing the working memory load when using a non-visual application. If the working memory load of manipulating the application is reduced that leaves more cognitive resources to be applied to the task the user is trying to solve. There has been some work on using such memory aids for visually-impaired people, but no detailed studies. We conducted a range of studies into the different modalities and combinations to really understand how to design and use external memory aids.

An initial study looked at the design of force-feedback memory aids using force beacons. A key issue is the amount of force that should be used; too much would cause interaction problems and too little would not be noticed. An investigation was carried out into the profile of the forces used. The results showed that 1/sqrt(distance) was the best profile: it gets users reliably to the point of interest but does not interfere too much.

The second area of work was related to this and focused on force trajectory playback, or dragging a user around an object shape. This is useful as users could mark an interesting object and then be dragged back around it at a later point. Based on work presented in D6, haptic playback on its own was problematic. The multimodal combination of audio and haptic feedback resolved many of these issues and made such shape gestures much more effective.

We also investigated the tactile modality for presenting external memory aids. These are less intrusive as users are not dragged to the target, but can present useful information as a user moves over part of the display. A detailed study was undertaken to investigate the different aspects of pin-array based displays that could be manipulated to give information. It is more difficult than one expects to create effective patterns. We tried many different designs to come up with a set that work well. In the experiments reported here we studied shape, size and blink speed. The results showed that shape and blink speed were independent parameters that could be used effectively. Pattern size was more difficult so the recommendation is to use only one (larger) size. This work has also influenced others outside of the project, such as Kildal who is also now looking at tactile memory aids (Kildal and Brewster, 2006).

Non-speech auditory external memory aids have much promise as they do not conflict with interaction, as force-feedback might if not designed correctly. Within the context of a haptic display our study looked at the design of five different types of audio feedback to support marking locations. This would allow someone to mark a particular location and then know where it was in the future, so that he/she could return to it. The solution that was most effective was the virtual microphone. Here the cursor acts like a microphone, though which the user can hear objects in the scene. Objects were presented in 3D sound with a loudness cue to indicate distance.

The last study looked at a more complete multimodal drawing application that included some of the audio and force-feedback memory aids discussed above. For example, the user can be dragged to objects that have been drawn, rather than having to search for them. A user study was conducted to see how such an application would be used by blind participants. Results showed that the haptic guidance was used, especially in the case of sighted and blind people working together.

The results of the work presented here are from some of the most in-depth studies into the design of external memory aids. We have answered many basic questions about how to design such aids in several different modalities. We have also begun to study them in real applications. This work sets the foundations for others to use these techniques to significantly improve the interactions of visually-impaired people with complex multimodal environments.

10. References

Art Beyond Sight Art Education for the Blind. (2003). Inc. and AFB Press.

Alcañiz, M., Lozano, J., & Rey B. (2004). Technological background about VR. In G. Riva, C. Botella, P. Légeron, & G. Optale (Eds.), Cybertherapy: Internet and Virtual Reality as Assessment and Rehabilitation Tools for Clinical Psychology and Neuroscience (pp. 199-214). Amsterdam: Ios Press.

Brewster, S. (2003). Nonspeech auditory output. In J. A. Jacko & A. Sears (Eds.), The Human Computer Interaction Handbook (pp. 220-239). Mahwah: Lawrence Erlbaum Associates.

Brewster, S. A. (1994). Providing a structured method for integrating non-speech audio into human-computer interfaces. Unpublished PhD Thesis, University of York, York, UK.

Brown, L. M., Brewster, S. A., & Purchase, H. C. (2005). A First Investigation into the Effectiveness of Tactons. Paper presented at the In Proceedings of WorldHaptics, Pisa, Italy.

Brown, L. M., Brewster, S. A., & Purchase, H. C. (2006). Multidimensional Tactons for Non-Visual Infomation Display in Mobile Devices. Paper presented at the In Proceedings of MobileHCI, Espoo, Finland.

Burdea, G. (2000). Haptic feedback for virtual reality. Special issue on Virtual Prototyping, International Journal of Design and Innovation Research, 2(1), 17-29.

Bussell, L. (2006). Touch Tiles: Elementary Geometry Software with a Haptic and Auditory Interface for Visually Impaired Children.

Chang, A., & O'Sullivan, C. (2005). Audio-Haptic Feedback inMobile Phones. Paper presented at the In Proceedings of CHI Portland, Oregan, USA.

Clark, F., & Horch, K. (1986). Kinesthesia. In K. Boff, L. Kaufman, & J. Thomas (Eds.), Handbook of Perception and Human Performance v.1 Sensory Processes and Perception (pp. 13.1-13.62). New York: Wiley.

Colwell, C., Petrie, H., Hardwick A. and Furner, S. (1998) Haptic Virtual Reality for Blind Computer Users, Proceedings of the ACM Conference on Assistive Technology - ASSETS 98, pp 92-99

Crossan, A., & Brewster, S. (2006). Two-Handed Navigation in a Haptic Virtual Environment. Paper presented at the In Extended Proceedings of CHI, Montreal, Canada.

Crossan, A., Williamson, J., & Brewster, S. (2006). A General Purpose Control-Based Playback for Force Feedback Systems. Paper presented at the In Proceedings of Eurohaptics Paris, France.

Dang, T., Annaswamy, T. M., & Srinivasan, M. A. (2001). Development and Evaluation of an Epidural Injection Simulator with Force Feedback for Medical Training. Paper presented at the Medicine Meets Virtual Reality, USA.

Edwards, A. (1989). Soundtrack: An Auditory Interface for Blind Users, Human-Computer Interaction, 4(1), 45-66.

Edwards, W., Mynatt, E., & Stockton, K. (1994). Providing Access to Graphical User Interfaces — Not Graphical Screens. In Proceedings of the first annual ACM conference on Assistive technologies, ASSETS ’94 (pp. 47-54). New York: ACM Press.

Feygin, D., Keehner, M., & Tendick, F. (2002). Haptic Guidance: Experimental Evaluation of a Haptic Training Method for a Perceptual Motor Skill. Paper presented at the Haptic Interfaces for Virtual Environment and Teleoperator Systems, Florida.

Gaver, W. (1986). Auditory icons: Using sound in computer interfaces. Human-Computer Interaction, 2, 167-177.

Gaver, W. (1989). The SonicFinder: An Interface that Uses Auditory Icons, Human-Computer Interaction, 4(1), 67-94.

Gentry, S., Wall, S., Oakley, I., & Murray-Smith, R. (2003). Got Rhythm? Haptic-only Lead and Follow Dancing. Paper presented at the In Proceedings of Eurohaptics, Dublin, Ireland.

Goldstein, E. (1999). Sensation and perception (5th ed.) Chapters 7, 11 & 12. Pacific Grove, CA: Brooks/Cole.

Graham, J. A., & Argyle, M. (1975). A cross-cultural study of the communication of extra-verbal meaning by gestures. International Journal of Psychology, 10(1), 57-67.

Hansson, C. (2003). Haptic Drawing Program. Unpublished Master Design Sciences, Lund University, Lund.

Henmi, K., & Yoshikawa, T. (1998). Virtual lesson and its application to virtual calligraphy system. Paper presented at the In Proceedings of IEEE International Conference on Robotics and Automation.

Hoggan, E., & Brewster, S. A. (2006). Crossmodal Spatial Location: Initial Experiments. Paper presented at the In Proceedings of NordiCHI, Oslo, Norway.

Iglesias, R., Casado, S., Gutierrez, T., Barbero, J. I., Avizzano, C. A., Marcheschi, S., et al. (2004). Computer graphics access for blind people through a haptic and audio virtual environment. Paper presented at the In Proceedings of Haptic, Audio and Visual Environments and Their Applications, HAVE 2004.

Immersion. VibeTonz System. http://www.immersion.com/mobility/ [Electronic Version].

Iwata, H. (2003). Haptic interfaces. In J. A. Jacko & A. Sears (Eds.), The Human Computer Interaction Handbook (pp.206-219). Mahwah: Lawrence Erlbaum Associates.

Kamel, H. M. (2003). The Integrated Communication 2 Draw (IC2D). Unpublished PhD, University of California.

Kennedy, J. M. (1993). Drawing and the Blind - Pictures to Touch. New Haven and London: Yale University Press.

Kildal, J. and Brewster, S.A. VibroTactile External Memory Aids in Non-Visual Browsing of Tabular Data. In Vol II Proceedings of First International Workshop on Haptic and Audio Interaction Design 2006 (Glasgow, UK), pp 40-43.

Klatzky, R., & Lederman, S. (1999). The haptic glance: A route to rapid object identification and manipulation. In D. Gopher & A. Koriats (Eds.) Attention and Performance XVII. Cognitive regulations of performance: Interaction of theory and application (pp. 165-196). Mahwah: Erlbaum.

Kramer, G. (1994). An introduction to auditory display. In G. Kramer (Ed.), Auditory Display: Sonification, Audification and Auditory Interfaces (pp. 1-77). Reading: Addison-Wesley.

Langdon, P., Hwang, F., Keates, S., Clarkson, P. J., and Robinson, P. (2002), Investigating haptic assistive interfaces for motion-impaired users: Force-channels and competitive attractive-basins. In Proceedings of Eurohaptics 2002 (Edinburgh UK, 2002), pp. 122-127

Lederman, S., & Klatzky, R. (1987). Hand movements: A window into haptic object recognition. Cognitive Psychology, 19, 342-368.

Lemmetty, S. (1999). Review of speech synthesis technology, Masters thesis, Helsinki University of Technology, Espoo.

MacLean, K. (2000). Designing with haptic feedback. in Proceedings of IEEE Robotics and Automation ICRA’2000 (pp. 783-788). Retrieved November 15, 2006 from http://www.cs.ubc.ca/~maclean/publics/icra00-DesignWithHaptic-reprint.PDF

MacLean, K., & Enriquez, M. (2003). Perceptual Design of Haptic Icons. Paper presented at the in Proceedings of Eurohaptics, Dublin, Ireland.

Magnusson, C., & Rassmus-Gröhn, K. (2005). A Virtual Traffic Environment for People with Visual Impairments. Visual Impairment Research, 7(1), 1-12.

Magnusson, C., & Rassmus-Gröhn, K. (2005). Audio haptic navigational tools for non visual environments, In Proceedings of the first ENACTIVE Workshop. Retrieved November 15, 2006 from http://www.certec.lth.se/haptics/papers/ENACTIVE05_AudioHapticToole_CMagnussonl.pdf

Magnusson, C., & Rassmus-Gröhn, K. (2005). Audio haptic tools for navigation in non visual environments, 11-17-0050.

Magnusson, C., Danielsson, H., Rassmus-Gröhn, K., Non Visual Haptic Audio Tools for Virtual Environments, Presented at the Workshop on Haptic and Audio Interaction Design University of Glasgow, 31st August - 1st September 2006

Magnusson, C., Rassmus – Gröhn, K. (2005), A Virtual Traffic Environment for People with Visual Impairment, Visual Impairment Research, Vol 7, nr 1, pp 1-12

Magnusson, C., Rassmus-Gröhn, K., Sjöström, C., & Danielsson, H. (2002). Navigation and Recognition in Complex Haptic Virtual Environments - Reports from an Extensive Study with Blind Users. Paper presented at the In Proceedings of Eurohaptics, Edinburgh, UK.

Magnusson, C., Rassmus-Gröhn, K.,Audio haptic tools for navigation in non visual environments ENACTIVE 2005, the 2nd International Conference on Enactive Interfaces, Genoa- Italy on November 17-18, 2005.

Mereu, S., & Kazman, R. (1996). Audio enhanced 3D interfaces for visually impaired users. In Human Factors in Computing Systems, CHI’96 Conference Proceedings (pp 72-78). New York: ACM Press.

Miller, G. A. (1956). The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. The Psychological Review, 63, 81-97.

Noble, N., & Martin, B. (2006). Shape discovering using tactile guidance. Paper presented at the In Proceedings of EuroHaptics, Paris, France.

Oakley, I. (2003). Haptic Augmentation of the Cursor: Transforming Virtual Actions into Physical Actions. Unpublished PhD, University of Glasgow, Glasgow.

Oakley, I., McGee M., Brewster, S., & Gray, P. (2000) Putting the feel in ’look and feel‘. In Proceedings of the ACM CHI 2000 Human Factors in Computing Systems Conference (pp. 415-422). New York: ACM Press.

Pakkanen, T., & Raisamo, R. (2005). Perception strategies in modal-redistributed interaction. Paper presented at the In Proceedings of 31st Annual International Symposium on Computer Architecture.

Pietrzak, T., Pecci, I., & Martin, B. (2006). Static and dynamic tactile directional cues experiments with VTPlayer mouse. Paper presented at the In Proceedings of Eurohaptics, Paris, France.

Pitt, I., & Edwards, A. (1991). Navigating the Interface by Sound for Blind Users. In D. Diaper & N. Hammond (Eds.), People and Computers VI, Proceedings of the CHI ’91 Conference (pp.373-383). Cambridge University Press.

Pokluda, L., Sochor, J., (2005) Spatial Orientation in Buildings Using Models with Haptic Feedback, Proceedings of WorldHaptics 2005 (Pisa, Italy, 2005), pp 523-524

Prada, R., Payandeh, S. (2005), A Study on Design and Analysis of Virtual Fixtures for Cutting in Training Environments, Proceedings of WorldHaptics 2005 (Pisa, Italy, 2005), pp 375-380

Rassmus-Gröhn, K., Magnusson, C., & Eftring, H. (2006, 8-31-2006). User evaluations of a virtual haptic-audio line drawing prototype. Paper presented at the In Proceedings of Haptic and Auditory Interaction Design Workshop, Glasgow, UK.

Riedel, B., & Burton, A. M. (2002). Perception of Gradient in Haptic Graphs by Sighted and Visually Impaired Subjects Paper presented at the In Proceedings of Eurohaptics, Edinburgh, UK.

Roberts, J.C., Franklin, K., Cullinane, J. (2002), Virtual Haptic Exploratory Visualization of Line Graphs and Charts. In Mark T. Bolas, editor, The Engineering Reality of Virtual Reality 2002, volume 4660B of Electronic Imaging Conference, page 10. IS&T/SPIE, January 2002.

Saga, S., Kawakami, N., Tachi, S. (2005), Haptic Teaching using Opposite Force Presentation, Poster presentation (CD proceedings) WorldHaptics2005

Sallnäs, E. (2004). The effect of modality on social presence, presence and performance in collaborative virtual environments, Doctoral thesis, Royal Institute of Technology, Stockholm.

Savidis, A., Stephanidis, C., Korte, A., Crispien, K., & Fellbaum, K. (1996). A generic direct-manipulation 3D-auditory environment for hierarchical navigation in non-visual interaction. In Proceedings of the second annual ACM conference on Assistive technologies, ASSETS ‘96 (pp. 117-123). New York: ACM Press.

Sekuler, R., & Blake, R. (2002). Perception (4th ed.) Chapter 11. New York: McGraw-Hill.

Sjöström, C. (1999) “The IT Potential of Haptics – touch access for people with disabilities”, Licentiate thesis, http://www.certec.lth.se/doc/touchaccess/

Sjöström, C. (2001a). Designing Haptic Computer Interfaces for Blind People, In Proceedings of the sixth international symposium on signal processing and its applications (pp. 68-71). Malaysia: Universiti Teknologi Malaysia.

Sjöström, C., (2001b). Virtual Haptic Search Tools - The White Cane in a Haptic Computer Interface, In Proceedings AAATE 2001, 6th European conference for the advancement of assistive technology (pp. 68-71). Retrieved November 15, 2006.

Sjöström, C. (2002). Non-Visual Haptic Interaction Design - Guidelines and Applications. Unpublished PhD, Lund University, Lund, Sweden.

Sjöström, C., Danielsson, H., Magnusson, C., & Rassmus-Gröhn, K. (2003). Phantom-based haptic line graphics for blind persons. Visual Impairment Research, 5(1), 13-32.

Snibbe, S., MacLean, K. Shaw, S., Roderick, J., Verplank, W., & Scheeff, M. (2001). Haptic Techniques for Media Control. In Proceedings of the 14th annual ACM symposium on user interface software and technology, UIST '01 (pp. 199–208). New York, ACM Press.

Srinivasan, M., Basdogan, C. (1997). Haptics in Virtual Environments: Taxonomy, Research Status, and Challenges. Computer and Graphics, 21(4), 393–404.

Stevens, R. (1996). Principles for the design of auditory interfaces to present complex information to blind people, Ph.D. Thesis, University of York, UK.

Subramanian, S., Gutwin, C., Sanchez, M.N., Power, C., & Liu, J. (2005). Haptic and Tactile Feedback in Directed Movements. In Proceedings of GOTHI'05 guidelines on tactile and haptic interactions. Retrieved November 15, 2006 from http://www.cs.usask.ca/research /research_groups/userlab/GOTHI/Sriram.pdf

Tang, A., McLachlan, P., Lowe, K., Saka, C., & MacLean, K. (2005). Perceiving ordinal data haptically under workload, In Proceedings of the 7th international Conference on Multimodal interfaces, ICMI '05 (pp. 317-324). New York: ACM Press.

Teo, C., Burdet, E., & Lim, H. (2002, March). A Robotic Teacher of Chinese Handwriting. Paper presented at the In Proceedings of the Symposium for Haptic Interfaces for Virtual Environment and Teleoperator Systems.

Wall, S., Brewster, S. (2004), Providing External Memory Aids in Haptic Visualisations for Blind Computer Users, In Proceedings of International Conference Series on Disability, Virtual Reality and Associated Technologies (ICDVRAT), New College, Oxford, UK, 20-22 September 2004, pp. 157-164.

Wall, S.A., Paynter, K., Shillito, A.M, Wright, M., Scali, S. (2002) The Effect of Haptic Feedback and Stereo Graphics in a 3D Target Acquision Task, In Proceedings of Eurohaptics 2002, Univeristy of Edinburgh, 8-10th July, 2002, pp. 23-29

Wang, D., Zhang, Y., & Yao, C. (2006). Stroke-based modeling and haptic skill display for Chinese calligraphy simulation system. Virtual Reality, 9, 118-132.

Wies, E., Gardner, J., O'Modhrain, S. and Bulatov, V. (2001). Web-based touch display for accessible science education. In Haptic human-computer interaction. (Brewster, S. A. and Murray-Smith, R. Eds.). Berlin, Springer-LNCS, 2058, pp. 52-60.

Winberg, F. (2001). Auditory direct manipulation for blind computer users, Licentiate thesis, Royal Institute of Technology, Stockholm.

Wood, J., Magennis, M., Arias, E.F.C., Helen Graupp, H., Gutierrez, T., Bergamasco, M. (2003), The Design and Evaluation of a Computer Game for the Blind in the GRAB Haptic Audio Virtual Environment, EuroHaptics 2003, pp 147-158

Yokokohji, Y., Hollis, R. L., Kanade, T., Henmi, K., & Yoshikawa, T. (1996). Toward Machine Mediated Training of Motor Skills -Skill Transfer from Human to Human via Virtual Environment. Paper presented at the In Proceedings of RO-MAN, Japan.

Yu, W. and Brewster, S. A. (2003). Evaluation of multimodal graphs for blind people, Journal of Universal Access in the Information Society 2 2,pp. 105-124.

Yu, W., & Brewster, S. A. (2002). Comparing Two Haptic Interfaces for Multimodal Graph Rendering. Paper presented at the In proceedings of IEEE VR2002, 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, Florida, USA.

Yu, W., Brewster, S., Ramloll, R., & Riedel, B. (2001). Exploring computer-generated line graphs through virtual touch. Paper presented at the In Proceedings of Sixth International, Symposium on Signal Processing and its Applications.

Yu, W., Kangas, K., & Brewster, S. (2003). Web-based haptic applications for blind people to create virtual graphs. Paper presented at the In Proceedings of 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems.

Zecca, M., Micera, S., Carozza, M., & Dario, P. (2002). Control of Multifunctional Prosthetic Hands by Processing the Electromyographic Signal, Critical Reviews™ in Biomedical Engineering, 20(4-6), 459-485.

Zhang, J. and Norman, D. A. (1994). Representations in distributed cognitive tasks, Cognitive Science, pp. 87-122.