

Using Haptic Cues to Aid Non-Visual Structure Recognition

Caroline Jay, Robert Stevens, Roger Hubbard and Mashhuda Glencross

Retrieving information presented visually is difficult for visually disabled users. Current accessibility technologies such as screen readers fail to convey presentational layout or structure. Information presented in graphs or images is almost impossible to convey through speech alone.

In this paper we present the results of an experimental study investigating the role of touch (haptic) and auditory cues in aiding structure recognition when visual presentation is missing. We hypothesize that by guiding users towards nodes in a graph structure using force fields, users will find it easier to recognise overall structure. Nine participants were asked to explore simple 3D structures containing nodes (spheres or cubes) laid out in various spatial configurations and asked to identify the nodes and draw their overall structure. Various combinations of haptic and auditory feedback were explored.

Our results demonstrate that haptic cues significantly helped participants to quickly recognise nodes and structure. Surprisingly, auditory cues alone did not speed up node recognition; however, when they were combined with haptics both node identification and structure recognition significantly improved. This result demonstrates that haptic feedback plays an important role in enabling people to recall spatial layout.

Categories and Subject Descriptors: H.1.2 [User/Machine Systems]: Human factors; H.5.2 [User Interfaces]: Haptic I/O; H.5.2 [User Interfaces]: Evaluation/methodology

Additional Key Words and Phrases: Multimodal Cues, Haptic Perception, Visual Disability, Accessibility

1. INTRODUCTION

For many visually disabled people, the primary means of accessing virtual information is using a screen reader. Whilst they work well for text, screen readers cannot provide access to a wealth of other information that relies on the visual presentation of the data to convey its meaning.

The problem exists at several levels. The text contained in Web pages can be presented via a screen reader, but without the meta-information provided by the layout, pages can be hard to navigate, and the required information can be difficult to find. The values contained in a table can be read by a screen reader, but the lack of spatial layout means that the relationships between them can quickly become unintelligible. Information contained in graphs, images or models is almost impossible to convey via speech alone. Even when a verbal description of the data can be generated, it is unlikely to capture the rich information contained in the visual presentation. The relationship between atoms in a complex

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0000-0000/20YY/0000-0001 \$5.00

molecule, for example, can easily be displayed and understood through a 3D model, but would be far more difficult to explain verbally.

Haptic feedback offers an alternative means of presenting this type of information to visually disabled users. Translating the visual model to a haptic model that can be accessed via touch preserves important properties of the data, including its spatial layout. However, a direct mapping is not a complete solution, as the limited ‘point-contact’ nature of current haptic technology makes successful haptic interaction very difficult in the absence of visual stimuli.

The following section describes previous work utilising proprioceptive, kinaesthetic and haptic feedback to convey spatial or structural information to visually disabled users, and discusses the limitations of current haptic technology. Section 1.2 then describes the focus of this paper: using haptic and audio cues to guide the user’s hand around the environment to aid recognition of 3D data at both the micro (individual nodes) and macro (overall structure) level.

1.1 Accessing Data Structure using Multimodal Cues

Tactile diagrams are the traditional means of providing non-visual access to data such as graphs, pictures and maps. Although electronic tactile displays are increasing in sophistication [Ando et al. 2006], they are currently very expensive and tend to be small or have a poor resolution. Typical tactile diagrams thus consist of raised lines on swell paper, produced using a special printer [Edman 1991]. Such diagrams are widely used in educational settings, but they are often unsuitable for showing complex graphics: people must spend time learning how to use them efficiently, and frequently suffer from information overload due to the much lower bandwidth of the haptic medium [Aldrich and Hindle 2003; Yu and Brewster 2002]. Tactile diagrams have other problems: they are static, and must be completely reproduced to incorporate an alteration; and they are unsuitable for displaying certain types of data, such as three dimensional models.

Increasingly, multimodal cues, in the form of force feedback and non-speech audio, are being used to improve non-visual access to graphical or pictorial data. Although both speech [Petrie et al. 1997] and non-speech audio has been used to supplement tactile diagrams [Wall and Brewster 2006a], more recent research has concentrated on presenting non-visual data electronically, avoiding the need for an ‘offline’, paper representation of the data.

Much of the work in this field has focused on the haptic and audio presentation of tables and graphs. TableVis, developed by Kildal and Brewster, provides access to both an overview and individual elements of tabular data using sonification [Kildal and Brewster 2006]. The table is mapped to a WACOM graphics tablet, which the user browses using the tablet’s electronic pen. Each entry in the table is mapped to a MIDI note signifying its value: the higher the pitch, the greater the value. To gain an overview of the data, users can listen to the values of a row or column being played in order, allowing them to locate and compare entries without having to listen to and hold in memory spoken values. Although the set-up does not offer any haptic feedback, the unvarying position of the table on the tablet means that people can use proprioception and kinesthesia to help them navigate, a facility that is unavailable when using a standard mouse. Both the external memory provided by the graphics tablet and the overview gained from the sonification helped users to understand and browse the table more effectively.

Sonification, with the addition of haptic feedback, has also been used to provide access

to graphs. The ‘Tac-tiles’ application consists of a graphics tablet overlaid with a disk, providing access to pie chart data [Wall and Brewster 2006b]. Non-speech audio (MIDI notes mapped to values as in TableVis) again proved useful for providing an overview of the data. Tactile feedback, provided by a VTPlayer mouse to the non-dominant hand, allowed users to feel when they were moving between segments on the chart, but was only utilized by two participants.

Haptic feedback from a PHANToM device has proved more useful for browsing graphs than the tactile feedback from the VTPlayer mouse. In a bar chart browsing application, although participants made greater use of non-speech audio to gain an overview of the data, the majority also utilized the haptic feedback that allowed them to feel the spatial layout of the bars [McGookin and Brewster 2006]. In an application that allowed users to build graphs themselves, the haptic feedback from the PHANToM proved vital as a means of constructing and reviewing bars on a chart [McGookin 2006].

Audio feedback has successfully been used to provide an overview of ‘node-arc’ graphs [Brown et al. 2006], which convey relationships by specifying the spatial structure of a set of data, and it would initially appear that haptic feedback is also ideally suited to this task. In fact, with current haptic technology, interacting with 3D models in the absence of visual cues is not easy. A haptic model, which the user can touch and explore, appears to be a straightforward way of communicating the layout of the data, as the visual representation can be directly translated to the haptic model. Such a model preserves information about the connectivity and structure of the data, which may be difficult to convey using audio alone. Unfortunately, the point contact nature of the device (users normally interact with the virtual world using a single ‘pen’ probe or stylus) means that interacting haptically with data in the absence of visual feedback can be very difficult.

One of the primary problems can be locating the data to begin with. As Sjostrom (2001) reports: ‘For a blind person, locating an object with a point probe can be as hard as finding a needle in a haystack’ [Sjostrom 2001b]. An experiment by Faconti et al., examining how people interpret haptic versions of visual illusions, such as the Necker Cube, found that the free interaction they had originally intended proved unfeasible [Faconti et al. 2000]. If users were placed at a starting point on the relevant object (they were unable to find it without help), they could explore and recognize nearby parts of it, but they were unable to ‘jump’ from one object to another, thus missing considerable portions of the data.

Several studies report that even after locating an object, maintaining contact with it can be difficult, as the stylus has a tendency to fall off, particularly when the object is convex [Colwell et al. 1998; Yu et al. 2001; Sjostrom 2001a]. Colwell et al. describe users as becoming ‘lost in haptic space’ when they are unfamiliar with this type of interaction, as they frequently lose contact with an object they are trying to explore [Colwell et al. 1998].

The problems described above result when interaction with the virtual world is completely unconstrained. Without visual feedback, users are unsure where to direct the stylus. To aid haptic data exploration in the absence of vision it seems appropriate to provide some sort of guidance to users, at least until they are familiar with the data, or confident in the task [Colwell et al. 1998]. Sjostrom suggests using a virtual ‘magnet’ to pull users towards haptic items on a virtual desktop that would otherwise be hard to find, although he does not report an evaluation for this technique [Sjostrom 2001a]. This idea of drawing the user’s hand, and therefore attention, to the haptic data forms the basis of the work presented here.

1.2 Aiding Data Exploration with Haptic and Auditory Cues

In this paper we examine whether using haptic cues to guide the user's hand around the environment aids recognition of abstract data at both the micro (individual nodes) and macro (overall structure) level. Examples of the data used in the experiment are shown in Figure 1. Individual nodes (either cubes or spheres) are arranged in a particular pattern within the environment. A haptic spring directs the stylus to the leftmost node, and a constraint keeps the stylus attached to the surface of the object, whilst the user feels its shape. The user can move through the nodes in sequence by pressing the space bar: the spring moves the stylus from the first node to the second, the second to the third, the third to the fourth and from the fourth back to the first again. The user can repeat this process until he or she is confident of the spatial layout of the data in the environment. It is important to note that the haptic cues necessarily constrain the user's ability to freely explore the data. Previous research has shown that constraining exploration results in poorer object recognition than allowing free exploration [Lederman and Klatzky 1987]. Given the immense constraints imposed by the point probe haptic device itself, however, it is hypothesized that people will perform better with the cues than they will when they are allowed to freely explore the environment.

Previous research has shown that audio feedback is better than haptic for comparing data values [McGookin and Brewster 2006], and that non-verbal cues are preferable to speech when quickly describing data, reducing both mental workload and frustration [Ramlohl and Brewster 2002]. In the current implementation, when the user is in touch with an object, he or she can request an audio cue that indicates the 'type' of that particular node (cube or sphere).

Cues that are able to convey accurate information about shape and structure have enormous potential: such technology could help visually disabled users to understand and navigate any visual electronic medium, including maps, graphs, diagrams and even Web pages. In the remaining part of the paper, we describe an experiment that tests the efficacy of these haptic and audio cues in aiding the recognition of individual nodes and overall structure when interacting haptically with abstract data.

2. METHOD

2.1 Equipment

The experiment was run on a Sony VAIO VGN-FS315S. Participants interacted with the environment haptically using a SenseAble PHANToM Omni and listened to audio cues provided by the VAIO on-board sound and speakers. The investigator watched the participants' interaction with the environment on the VAIO monitor. No visual display was available to the participants; however, they were also blind folded, so they could not see the position of their hand during the experiment (see Figure 2).

2.2 Environment

The environment was rendered using OpenGL and the OpenHaptics Toolkit. Figure 1 shows examples of the environment used in the experiment. A total of 16 data sets, designed to be equivalent to one another, were used. Each data set contained 4 nodes, a mixture of cubes and spheres 5cm in diameter, which were rendered haptically throughout the experiment. The nodes were arranged in a particular pattern on the same plane, constituting an overall structure.




Fig. 1. Examples of the stimuli

The surfaces of the objects were rendered using impedance control, with a surface stiffness coefficient of 0.7. When haptic cues were available, pressing the space bar guided the stylus around the objects using a spring, as described in section 1.2. The spring force, F , was calculated using Hooke's Law, $F = k(p - x)$, where p is the spring anchor (set at the centre of the object), x is the current stylus position and k is the gain (set at 0.6). When the stylus touched the target object, the spring was deactivated and a constraint held the stylus tip against the object's surface, allowing the stylus to move freely over the object's surface without losing contact with it. It was possible to leave the surface of the object by pulling away sharply. If this happened by mistake, the user could reproduce the spring to guide the stylus back to the object by pressing button 1 (the blue button) on the stylus (see Figure 3). Pressing button 2 (the white button) whilst touching a node played an audio cue that indicated the node type: a wind-chime sound indicated a sphere; and a bicycle bell indicated a cube.



Fig. 2. Performing the experiment.



Fig. 3. The PHANTom stylus.

2.3 Participants

3 female and 6 male right-handed, sighted participants between the ages of 25 and 30 took part. Each participant was paid £5. Although the application is intended to aid non-visual data exploration, and as such will be of chief benefit to visually disabled users, sighted participants were used in this instance because we wished to understand whether the mental representation users gained of the data matched the actual layout of the data – did their drawings look like the graphically rendered version of the environment, shown in Figure 1? If the technique is successful in achieving this, there is good evidence it should be developed further and tested with visually disabled users.

2.4 Design

Participants had to identify both the individual nodes and the structure (spatial layout) of the data under four conditions: no additional cues; audio cues only; haptic cues only; both audio and haptic cues. For the first 8 data sets, participants completed 2 tasks under each condition in a random order. For the remaining 8 data sets, the order of the conditions was reversed.

2.5 Procedure

Participants sat in front of the Omni, holding the stylus in their right hand, and pressed the space bar to start the session. First, participants completed a practice session in which they could use both haptic and audio cues. In the experimental sessions, if haptic cues were available, the space bar would guide them to the leftmost node, which they would then identify, speaking their answer out loud to the investigator. The investigator noted down whether the answer was correct. Pressing the space bar again would guide the stylus to the next node. If haptic cues were not available, participants had to search the environment independently. They were still asked to press the space bar immediately after identifying a node in order to time how long it took to identify the node. When participants had identified all of the nodes, they were free to continue exploring the environment until they were confident of the spatial relationships between the nodes. They then exited the virtual environment, and drew the overall structure on a piece of paper. The experiment took approximately 30 minutes to complete.

2.6 Questionnaire

Participants were asked to rank the conditions in order of how difficult they found them, from 4 (most difficult) to 1 (least difficult). They were also asked whether they found haptic and/or audio cues useful for identifying both the individual nodes and the overall structure.

2.7 Metrics

Four measurements were taken in each condition: the time to identify a node (the time from first touching the node to pressing the space bar); the time to identify the structure of the data (the time from first pressing the space bar to exiting the environment); the percentage of nodes correctly identified; and the percentage of structures correctly identified. When participants drew the structure, they were asked, if possible, to draw the node shapes as well. However, they were told that it was more important to identify the spatial relationship between the nodes than the nodes themselves. If they misidentified or could not

remember any node, it did not affect whether they were judged to have correctly identified the structure.

2.8 Hypothesis

It is anticipated that guiding the user's hand around the nodes in the environment will make it easier to recognise the overall structure, or spatial layout, of the data, and that audio cues will provide the most efficient way of recognising the node type, as it is quicker to listen to a single sound than to feel the surface of the object. However, there is likely to be some interaction between the two types of cues. Constraining the stylus to the surface of the object may help node recognition, by making it easier to produce an audio cue (these can only be played when touching an object), and making it easier to feel the surface of an object if audio cues are not available. Node recognition is thus hypothesized to be fastest and most accurate when both types of cues are present. If only one type of cue is available, audio cues are hypothesized to result in faster and more accurate node recognition than haptic cues, but haptic cues are still expected to result in better performance than no cues at all. If audio cues speed up node recognition, they may also speed up structure recognition, as the user has to spend less time identifying each object. Structure recognition is thus hypothesized to be fastest and most accurate when both cues are available. However, haptic cues are likely to provide the greatest benefits, and if audio cues are not available performance is only expected to be slightly worse. As audio cues do not give any spatial information, they are not expected to aid structure recognition in the absence of haptic cues, so performance when only audio cues are available is not expected to be any better than performance in the absence of both sets of cues. The hypotheses can be summarized as follows:

- (1) Both node and structure recognition will be fastest, most accurate and easiest when both types of cue are available.
- (2) Haptic cues alone will aid structure recognition but audio cues alone will not.
- (3) Once a node has been located, audio cues will provide a more efficient means of identifying it than haptic cues, but the haptic surface constraint will still provide an advantage over no cues at all.
- (4) Participants will rate haptic cues as more useful for identifying structure, and audio cues as more useful for identifying nodes.

3. RESULTS

The time to identify both individual nodes and the overall structure, and the accuracy with which participants were able to do this, varied significantly according to the cues that were available. Haptic cues enabled participants to move efficiently between nodes, and identify the structure quickly and accurately. The addition of audio cues improved participants' ability to identify individual nodes, and as such speeded up their recognition of the overall structure. However, contrary to expectation, audio cues alone did not speed up node recognition. The results are reported below in three sections, the first looking at structure recognition, the second looking at node recognition, and the third considering participants' ratings of the extent to which the cues contributed to their performance.

3.1 Structure Recognition

The participants' drawings were compared with a key illustrating the structure of the data in each condition. The experimenter and an independent observer made judgements as to




Fig. 4. Mean time to identify a structure in each condition. Error bars show standard deviation.

whether the drawing accurately reflected the structure. The most important criterion was that the angles between the nodes (180, 90 or 45 degrees) were accurately illustrated; the second criterion was that the drawings reflected the relative distance between the nodes (each node was one ‘unit’ from adjacent nodes in the X and/or Y directions). If either observer thought these criteria were not fulfilled, the drawing was regarded as incorrect.

Figure 4 illustrates the mean time taken to identify the structure in each condition. Three participants failed to identify any structures in the no cue condition, and five failed to identify any in the audio only condition, making statistical analysis of all four conditions inappropriate. However, a t-test shows that providing both types of cues results in significantly faster structure recognition than providing haptic cues alone ($t_8 = 2.61$, $p < 0.01$).

A repeated measures ANOVA with pairwise comparisons shows that haptic cues had a clear effect on the number of structures accurately identified in each condition (see Figure 5). The accuracy rate was more than twice as high in the conditions where haptic cues were provided than in the conditions where no cues or only audio cues were provided ($F_{3,24} = 16.36$, $p < 0.001$). Table I shows the p values from the pairwise comparisons where these are significant at the 5% level. The partial Eta squared value was 0.87, showing a very large effect size.

3.2 Node Recognition

Figure 6 shows how long it took participants to recognize a node in each condition. A repeated measures ANOVA shows a main effect of condition ($F_{3,24} = 3.38$, $p < 0.05$), but pairwise comparisons reveal that whilst the combination of both cues results in significantly faster selection times, providing either cue alone does not improve performance (see Table II). Again, the effect size was large, with the analysis showing a partial Eta

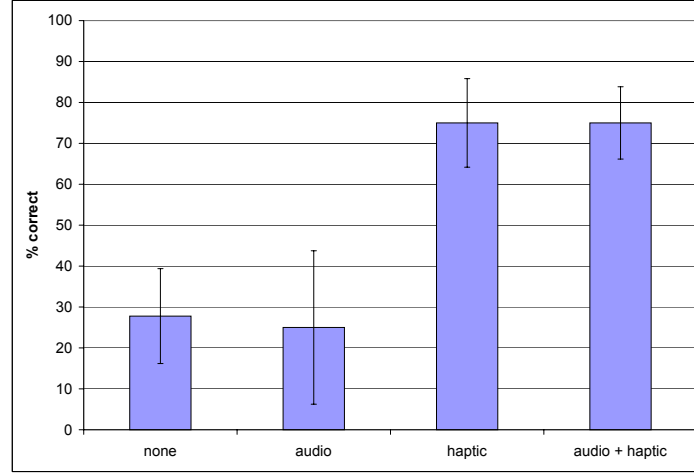


Fig. 5. Mean percentage of structures correctly identified in each condition. Error bars show standard deviation.

Pairwise Comparison	p
no cues to haptic cues	0.001
no cues to both cues	0.000
audio cues to haptic cues	0.002
audio cues to both cues	0.005

Table I. Structure recognition accuracy pairwise comparisons significant at the 5% level.

Pairwise Comparison	p
no cues to both cues	0.001
audio cues to both cues	0.029
haptic cues to both cues	0.031

Table II. Node recognition time pairwise comparisons significant at the 5% level.

squared value of 0.84. That constraining the stylus to the surface of the object did not speed up node recognition runs contrary to our hypothesis. However, more surprising is the fact that audio cues do not help people to recognise nodes in the absence of haptic cues (see Figure 8): once the user has found a node, using an audio cue would presumably be the fastest way to identify it. The time to identify a node in this case can be divided into two distinct periods: the time from first touching the object to requesting the cue; and the time from requesting the cue to identifying the node. The relative lengths of these two periods in the audio only and audio and haptic conditions are shown in Figure 7. T-tests show that both the time to request an audio cue ($t_8 = 5.76$, $p < 0.001$) and the time to identify the node following the cue ($t_8 = 2.61$, $p < 0.05$) are significantly higher in the audio only condition.




Fig. 6. Mean time to identify a node in each condition. Error bars show standard deviation.

Pairwise Comparison	p
no cues to both cues	0.037
no cues to haptic cues	0.011
audio cues to both cues	0.030
audio cues to haptic cues	0.008
haptic cues to both cues	0.017

Table III. Node recognition accuracy pairwise comparisons significant at the 5% level.

Although it is not immediately clear why participants take longer to perform both these actions in the absence of haptic cues, there are possible explanations. The increased request time could result from a difficulty initiating and maintaining contact with the object. The increased identification time could result from a lack of confidence in answering, but it could also be due to a delay in pressing the space bar, due to the fact that the action has no purpose when it does not guide the stylus to the next node.

Figure 8 shows the mean percentage of nodes correctly identified in each condition. A repeated measures ANOVA with pairwise comparisons shows that audio cues alone did not significantly improve accuracy ($F_{3,24} = 7.9$, $p < 0.001$). However, haptic cues did significantly increase the number of nodes recognised, presumably because they made it easier to locate them, and adding audio cues provided a further significant advantage (see Table III). The partial Eta squared value in this case was 0.77.

3.3 Questionnaire results

All participants rated the condition with both sets of cues as easiest, and the condition with none as hardest. Only one participant rated the audio only condition as easier than the haptic only condition. The reason for this is not entirely clear, as this participant had far

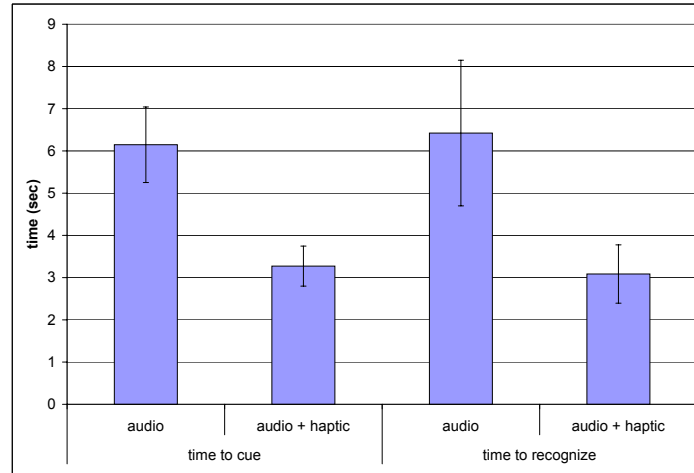


Fig. 7. The time from first touching an object to requesting an audio cue, and the time from requesting the cue to identifying the node in the audio only and audio and haptic conditions. Error bars show standard deviation.




Fig. 8. Mean percentage of nodes correctly identified in each condition. Error bars show standard deviation.

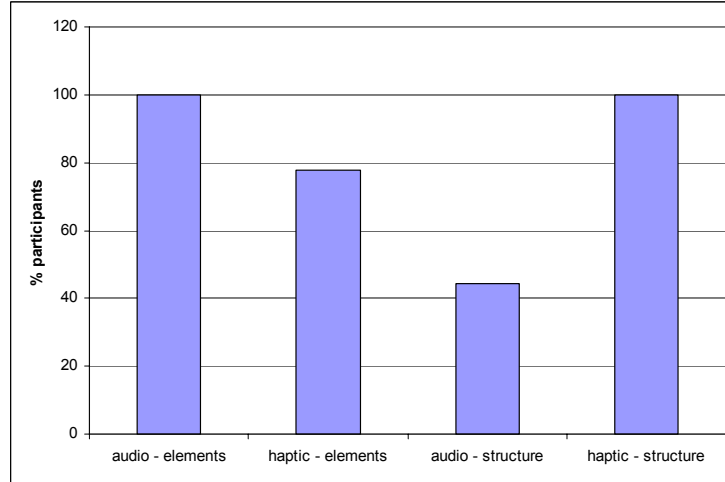


Fig. 9. Percentage of participants who found cues helpful for recognising nodes (left) and structure (right)

higher accuracy rates for both node and structure recognition in the haptic only condition.

All participants said that audio cues were useful for recognising nodes, and haptic cues were useful for recognising structure (see Figure 9). 78% found haptic cues useful for identifying nodes, and 44% reported that audio cues aided structure recognition.

4. DISCUSSION

Prior to conducting the experiment, we had four hypotheses:

- (1) Both node and structure recognition will be fastest, most accurate and easiest when both types of cue are available.
- (2) Haptic cues alone will aid structure recognition but audio cues alone will not.
- (3) Once a node has been located, audio cues will provide a more efficient means of identifying it than haptic cues, but the haptic surface constraint will still provide an advantage over no cues at all.
- (4) Participants will rate haptic cues as more useful for identifying structure, and audio cues as more useful for identifying nodes.

The results provided strong support for the first hypothesis. In the condition where haptic and audio cues were available together, participants identified both nodes and structures fastest, and identified nodes more accurately, than in any of the other conditions. The advantage the cues provided was highly significant, resulting in 71% faster structure recognition and 46% faster node recognition than the no cue condition. 100% of nodes were correctly identified in this condition - a result not achieved in any other condition. All of the participants rated this condition as the easiest overall.

The second hypothesis, that haptic cues will aid structure recognition, but audio cues will not, was also supported. The percentage of structures correctly identified stood at 75% for the haptic only and haptic and audio conditions, compared to 28% and 25% for the no cue and audio only conditions.

Haptic cues alone resulted in a higher recognition accuracy rate for nodes than audio cues, simply because they made it easier to locate the nodes. However, contrary to expectations, the third hypothesis—that audio (and to a lesser extent, haptic) cues would result in faster identification once a node had been located—was not supported. The times to identify a node in the audio only, haptic only and no cue conditions were identical.

Both the time to request a cue, and the time to identify the node following a cue, were longer in the audio only than the audio and haptic cue condition. The longer identification time could have been an artifact of the measurement procedure: pressing the space bar in the audio only condition served no purpose other than to record the recognition time, so participants may not have done it as automatically as they would have in the haptic cue conditions. However, this cannot explain the longer time to request the cue initially. The fact that participants appeared to hesitate could have occurred because establishing contact with an object is much harder when haptic cues are not available. If you are guided to an object, and then attached to the surface with a constraint, you can be sure you are touching it. If this facility is not available, it may take much longer to be confident you are in contact with an object, and therefore to request an audio cue.

Despite this, 100% of participants rated audio cues as useful for recognizing nodes, compared to 78% rating haptic cues as useful. 100% of participants also rated haptic cues as useful for recognizing structure, compared to 44% rating audio cues as useful, supporting hypothesis four. The large proportion of participants who found haptic cues useful for recognizing nodes indicates that help finding and maintaining contact with an object significantly improves the user's perception of how easy it is to identify it. The increased speed with which participants could identify nodes meant that several participants also rated audio cues as useful for identifying structure.

5. SUMMARY

Without visual feedback, understanding structure and layout is very difficult and time-consuming. The results reported here show that simple haptic cues, in the form of springs that guide the user's hand from one node to the next, greatly improve his or her ability to accurately identify the structure, or spatial layout of the data. The results also show that adding audio cues—tones corresponding to data type—reduces the time it takes to recognize individual nodes, and as such speeds up structure recognition as well.

However, the results also provide evidence of an important cross-modal interaction. Whilst haptic cues increased structure recognition speed and accuracy in the absence of audio cues, providing audio cues alone was not sufficient to improve node recognition. Being able to listen to the data type, rather than having to go through the time-consuming process of determining it through touch, improves users' ability to identify data, and their perception of task difficulty, but this only occurs when audio cues are combined with haptic. The reason for this is not entirely clear, but it may well result from the reduced confidence that users have when interacting with the environment without the haptic cues. It is clear that the spatial information provided by the haptic cues is central to improving both performance and the interaction experience, and without it participants are left, literally,

hanging in mid-air.

REFERENCES

- ALDRICH, F. K. AND HINDLE, Y. 2003. Tactile graphics – problem or solution? *RNIB Visibility* 39, 25–26.
- ANDO, H., WATANABE, J., AMEMIYA, T., AND NAKATANI, M. 2006. Siggraph '06: Acm siggraph 2006 emerging technologies. ACM Press, New York, NY, USA. Conference Chair-John Finnegan and Program Chair-Tom Craven.
- BROWN, A., STEVENS, R., AND PETTIFER, S. 2006. Audio representation of graphs: A quick look. In *Proceedings of the International Conference on Auditory Displays 2006*.
- COLWELL, C., PETRIE, H., KORNBROT, D., HARDWICK, A., AND FURNER, S. 1998. Haptic virtual reality for blind computer users. In *Assets '98: Proceedings of the third international ACM conference on Assistive technologies*. ACM Press, New York, NY, USA, 92–99.
- EDMAN, P. K. 1991. *Tactile graphics*. American Foundation for the Blind.
- FACONTI, G. P., MASSINK, M., BORDEGONI, M., ANGELIS, F. D., AND BOOTH, S. 2000. Haptic cues for image disambiguation. *Comput. Graph. Forum* 19, 3.
- KILDAL, J. AND BREWSTER, S. A. 2006. Non-visual overviews of complex data sets. In *Proceedings of ACM CHI*. ACM Press, Montreal, Canada, 947–952.
- LEDERMAN, S. J. AND KLATZKY, R. L. 1987. Hand movements: A window into haptic object recognition. *Cognitive Psychology* 19, 342–368.
- MCGOOKIN, D. K. 2006. Graph builder: Constructing non-visual visualizations. In *ICAD 2006 (London, UK)*.
- MCGOOKIN, D. K. AND BREWSTER, S. A. 2006. Multivis: Improving access to visualisations for visually impaired people. In *CHI '06: CHI '06 Extended Abstracts on Human Factors in Computing Systems*. ACM Press, Montreal Canada, 267–270.
- PETRIE, H., JOHNSON, V., STROTHOTTE, T., RAAB, A., MICHEL, R., REICHERT, L., AND SCHALT, A. 1997. Mobic: An aid to increase the independent mobility of blind travellers. *The British Journal of Visual Impairment* 15, 2, 63–66.
- RAMLOLL, R. AND BREWSTER, S. 2002. A generic approach for augmenting tactile diagrams with spatial non-speech sounds. In *CHI '02: CHI '02 extended abstracts on Human factors in computing systems*. ACM Press, New York, NY, USA, 770–771.
- SJOSTROM, C. 2001a. Designing haptic computer interfaces for blind people. In *Proceedings of the 3rd International Workshop on Website Evolution*. IEEE.
- SJOSTROM, C. 2001b. Virtual haptic search tools - the white cane in a haptic computer interface. In *AAATE*.
- WALL, S. AND BREWSTER, S. 2006a. Feeling what you hear: tactile feedback for navigation of audio graphs. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM Press, New York, NY, USA, 1123–1132.
- WALL, S. A. AND BREWSTER, S. A. 2006b. Tac-tiles: multimodal pie charts for visually impaired users. In *NordiCHI '06: Proceedings of the 4th Nordic conference on Human-computer interaction*. ACM Press, New York, NY, USA, 9–18.
- YU, W. AND BREWSTER, S. 2002. Multimodal virtual reality versus printed medium in visualisation for blind people. In *ACM Assets*. Edinburgh, Scotland, 57–64. ISBN 1-58113-464-9-02/07.
- YU, W., RAMLOLL, R., AND BREWSTER, S. A. 2001. Haptic graphs for blind computer users. In *Proceedings of the First International Workshop on Haptic Human-Computer Interaction*. Springer-Verlag, London, UK, 41–51.