GEOMETRIC SHAPE DETECTION WITH SOUNDVIEW

K. van den Doel[1], D. Smilek[2], A. Bodnar[1], C. Chita[1], R. Corbett[1], D. Nekrasovski[1], J. McGrenere[1]

Department of Computer Science[1]
Department of Psychology[2]
University of British Columbia
Vancouver, Canada

ABSTRACT

We present the results of user studies that were performed on sighted people to test their ability to detect simple shapes with SoundView. SoundView is an experimental vision substitution system for the blind. Visual images are mapped onto a virtual surface with a fine-grained color dependent roughness texture. The user explores an image by moving a pointer device over the image which creates sounds. The current prototype uses a Wacom graphics tablet as a pointer device. The pointer acts like a virtual gramophone needle, and the sound produced depends on the motion as well as on the color of the area explored. An extension of SoundView also allows haptic feedback and we have compared the performance of users using auditory and/or haptic feedback.

1. INTRODUCTION

SoundView is a system which allows a user to sense a static image synesthetically [1] through sound and touch. SoundView operates by constructing a virtual surface with a roughness texture corresponding to the image. Instead of feeling the roughness through touch, the user scrapes the surface with a virtual gramophone needle, which is moved with a pointing device such as a graphics pen.

For the details of the SoundView design and a review of related work on cross-modal vision systems for the blind we refer to [2].

In order to test the usability of the SoundView system we have performed user studies on the detectability of simple black and white shapes by sighted people. Sighted subjects were chosen for logistical reasons. We believe that if the results for sighted people are encouraging, blind users will most likely perform better, so tests on sighted people will provide us with a conservative estimate of the capabilities of SoundView. If the results are positive, then the next step can be taken in the form of clinical trials on blind subjects. We have also compared the performance of SoundView with Peter Meijers vision substitute for the blind "The vOICe" [3, 4, 5], which translates images from a camera on-the-fly into corresponding sounds.

Apart from measuring raw performance using the system, we are also interested in determining *how* people use the system to observe images as this will give us insights which will allow us to improve the usability of the system. In order to determine the importance of the nature of the feedback we have created an extension of SoundView which also allows haptic feedback and we have performed user studies aimed at comparing performance of shape detection using auditory, haptic, or combined auditory and haptic feedback.

The remainder of this paper is organized as follows. In Section 2 we describe two user studies on shape detection using the SoundView system. The first experiment asks user to draw the shapes they thought they were detecting, in order to get qualitative insight in the perception of shapes with the system. The second experiment is an eighteen alternative forced choice test. In Section 3 we perform two six alternatives forced choice experiments using SoundView and The vOICe in order to compare performance. In Section 4 we describe the extension to SoundView with haptic feedback and the result of user studies using sound only, haptic only, or both. Conclusions are presented in Section 5.

Figure 1: The six basic shapes used in Experiments 1 through 3.

2. EXPERIMENTS 1 AND 2: TESTING SOUNDVIEW

The goal of our initial tests of SoundView was to determine whether individuals could use SoundView to identify several basic geometric shapes such as those shown in Figure 1. In two experiments participants explored the shapes by moving a pen on a tablet and listening to the auditory feedback generated by SoundView. We tested people's ability to recognize the shapes in two different ways. In Experiment 1 participants were required to draw the shape on a sheet of paper and in Experiment 2 participants had to choose the correct shape from a set of 18 alternative shapes. Each of these experiments is described below.

2.1. Experiment 1

Methods

Eight undergraduate students at the University of British Columbia participated in a 20-minute session for course credit. All participants reported normal hearing and had normal or corrected to normal vision. Before commencing the experiment, each participant was given general instructions regarding how to interpret the auditory feedback provided by SoundView. Participants then explored a series of six shapes by moving a pen on a WACOM tablet (Pen PartnerTM) that measured 9.7 cm vertically and 13.8 cm horizontally. The six shapes that were used in the experiment are shown in Figure 1. Notice that half of the shapes contained a hole and half of the shapes did not contain a hole. Participants were not shown the shapes at any point in the experiment, nor where they told what the possible shapes would be. They were simply told that the stimuli were "simple" shapes. The shapes occupied roughly 60% of the active space on the tablet. The WACOM tablet was connected to a desktop computer driven by a 1.8 GHz Pentium III processor.

Each of the six shapes depicted in Figure 1 was presented on a separate trial of the experiment. Each shape was presented only once in the experiment for a total of six trials in the experiment. The order of presentation of the shapes was randomized and thus differed across participants. Each trial of the experiment was preceded by an auditory message instructing the participant to begin exploring the shape. Participants were given 90 seconds to explore the shape. During the exploration time, participants were allowed to view their hand and the pen. After the 90 seconds of exploration elapsed, participants were given auditory instruction to record their answer by drawing the shape on a sheet of paper. Participants were given 90 seconds to record their answer. The next trial was initiated automatically following the 90-second response period.

Results

The accuracy of free drawings was assessed in three different ways. First, we assessed the accuracy with which participants correctly reported the presence or absence of a hole in the shape. Second we assessed the accuracy with which participants drew the external contours of the shapes. Finally, we assessed the overall accuracy with which participants reported both the presence or absence of a hole and the external contours of the shape. The results showed that participants reported the presence or absence of a hole with 68.8% accuracy. Furthermore, participants accurately drew the external contours of the shapes 35.4% of the time. Finally, the overall percentage of trials on which participants accurately depicted both the presence and absence of a hole and the shape of the external contours was 30.0%. These initial results suggest that detecting whether or not a shape contained a hole using SoundView was relatively easier than ascertaining the precise contours of the shape.

Further inspection of the drawings revealed that our measure of the accuracy with which participants recorded the contours of the shapes was likely a very conservative estimate of performance because even small deviations from the actual contours were coded as being incorrect. Because of the general difficulty of analyzing freehand drawings, we conducted another experiment (Experiment 2) using an 18-alternative forced choice procedure. By using a forced choice procedure it was possible to evaluate shape recognition in a more objective manner.

2.2. Experiment 2

Methods

Thirty undergraduate students at the University of British Columbia participated in a 20-minute session for course credit. None of the participants in this experiment participated in the previous experiment. As in the previous experiment, all participants reported normal hearing had normal or corrected to normal vision.

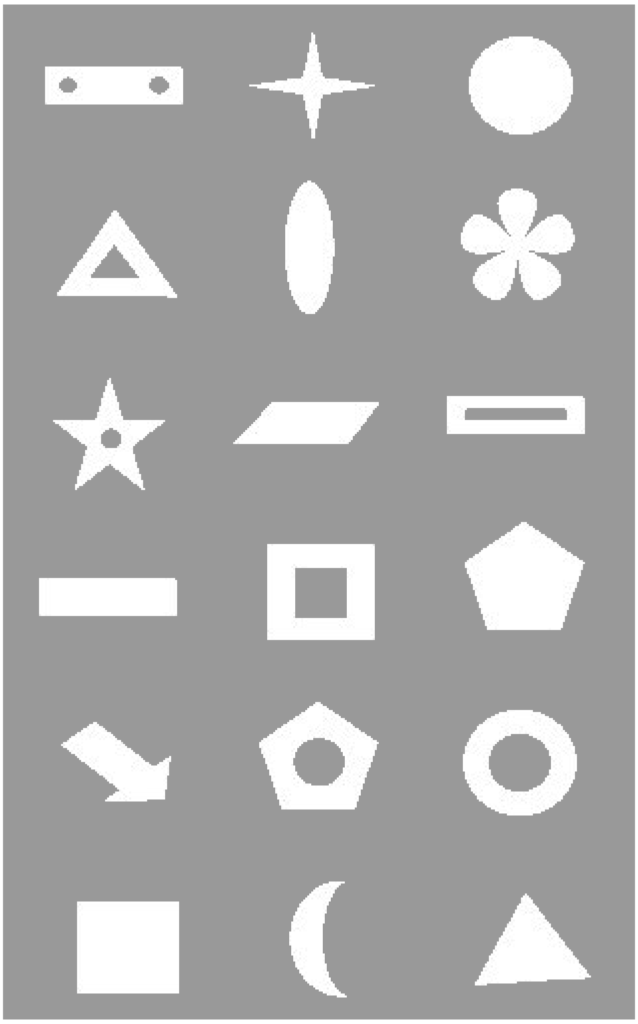

Figure 2: The 18 shapes from which participants had to choose the correct answers in Experiment 2.

The apparatus and procedures used in the present experiment were similar to those of Experiment 1. As in Experiment 1, participants were presented with each of the six shapes shown in Figure 1. Each shape was presented once on a separate trial of the experiment for a total of six trials. The exploration and report durations were identical to those used in Experiment 1. Experiment 2 differed from Experiment 1 in an important way. In Experiment 2, rather than drawing the shape, participants had to choose the correct shape from among 18 alternative shapes. Participants reported their choice by circling one of 18 shapes printed on sheet of paper. The 18 alternatives used in the experiment are shown in Figure 2.

Results

The mean percent correct shape discrimination, averaged across participants and shapes, was 38.3% (standard deviation (SD) = 24.8%). A single sample t-test revealed that this overall mean accuracy was significantly greater than that expected by chance alone (i.e., 5.6%), t(29) = 7.242, $p < 0.001$. These results indicate that

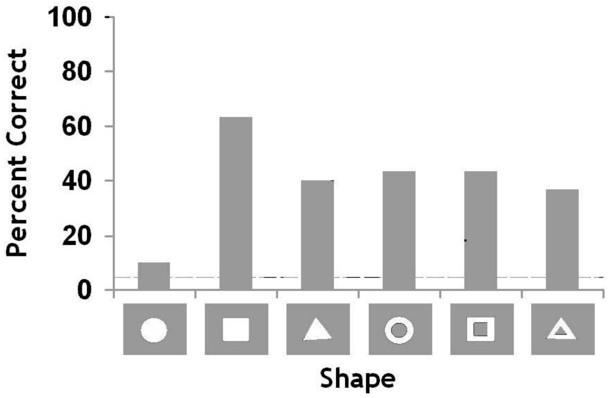

Figure 3: The mean percent correct discrimination performance for each of the six shapes in Experiment 2. The broken line represents chance performance.

participants were able to use SoundView to discriminate among the shapes.

The mean percent correct for discriminating each of the six shapes, averaged across participants is shown in Figure 3. A one-way repeated measures analysis of variance (ANOVA) revealed that discrimination accuracy differed across shapes, F(5, 145) = 4.795, MSE = 0.188, $p < 0.001$. Inspection of Figure 3 reveals that this overall difference in discrimination between the shapes was likely due to the poor discrimination of the circle, which did not differ from chance, t(29) = 1.225, $p = 0.230$. Apart from the circle, discrimination of the each of the other shapes was substantially above chance performance as indicated by a series of single sample t-tests, all $ts > 3.168$, all $ps < 0.004$. These results further support the general conclusion that participants were able to discriminate the shapes at above chance levels.

3. EXPERIMENT 3: COMPARING SOUNDVIEW WITH THE vOICe

Another set of experiments was conducted to compare the usability of SoundView with the usability of The vOICe. The vOICe [3, 4, 5] is a vision substitute for the blind which translates images from a camera on-the-fly into corresponding sounds. This is done by sweeping a vertical scan line periodically over the image. The scan line generates sounds depending on the brightness of the pixels it is crossing and the height is mapped to pitch. Though the sounds thus created are not easily interpreted at first it is hoped that the brain can learn to map the information in the sounds to images, either through induced synesthesia, or simply by providing equivalent information through the auditory channel. In [6] acquired synesthesia was reported to appear in a patient several years after vision loss. The patient experienced visual sensations evoked by tactile stimuli on the hands. The main differences in design philosophy between SoundView and The vOICe are first that SoundView uses active exploration with sound, whereas The vOICe passively produces the soundscape of an image. Second, SoundView has been designed in order to make the correspondence between images and sounds as intuitive and easy to learn as possible, whereas the relation between sound and image in The vOICe is probably more difficult to learn. The results of the studies presented here can therefore not readily be used to draw conclusions about the performance difference between the two systems by trained users and are more indicative of novice behavior.

Methods

Sixty undergraduate students at the University of British Columbia participated in a 20-minute session for course credit. None of the students who took part in the experiment participated in any of the pervious experiments. All participants reported normal hearing and had normal or corrected to normal vision.

The technique used to convert visual information into sound was varied across participants resulting in two conditions. In one condition (the SoundView condition) participants discriminated among shapes using auditory feedback from SoundView. In the other condition (The vOICe condition) participants discriminated among shapes using auditory feedback from The vOICe.

The apparatus and procedures used in the SoundView condition were similar to those of Experiment 2. Participants once again explored the six shapes shown in Figure 1 using feedback from SoundView. The exploration and report durations were the same as those in Experiment 2. In this experiment, however, we measured discrimination of the shapes by having participants choose the correct shape from six alternative shapes rather than the 18 alternatives used in Experiment 2. The six alternatives were printed on a sheet of paper and on each of the six trials in the experiment, participants had to circle the correct shape. The six alternative shapes that participants had to choose from were the six shapes used as stimuli in the experiment (see Figure 1).

The methodology used in The vOICe condition was closely matched to that of the SoundView condition. As did the participants in the SoundView condition, participants in The vOICe condition listened to the sounds that corresponded to each the six shape shown in Figure 1 for a duration of 90 seconds and were then given 90 seconds to choose the correct answer from six alternative shapes printed on a sheet of paper. However, there were also several critical differences between the conditions. One important difference concerned the instructions given to participants. Whereas in the SoundView condition participants were told how to interpret auditory feedback from SoundView, in The vOICe condition, participants were taught to interpret the auditory output form The vOICe. Specifically, before starting the experimental trials, participants in The vOICe condition completed five simple examples, each of which involved viewing a shape for ten seconds while listening to the corresponding output from The vOICe program. The shapes that were used in the examples were different than the shapes used on experimental trials. Another critical difference between the conditions involved the nature of the auditory feedback. In The vOICe condition, participants did not explore the shape with a pen as they did in the SoundView condition, but passively listened to the sounds generated from The vOICe that corresponded to each of the six shapes shown in Figure 1.

Results

The overall percent correct shape discrimination in the SoundView condition was 66.2% (SD = 26.8%). A single sample t-test revealed that this overall discrimination accuracy was significantly greater than chance, which in this experiment was 16.6%, t(29) =

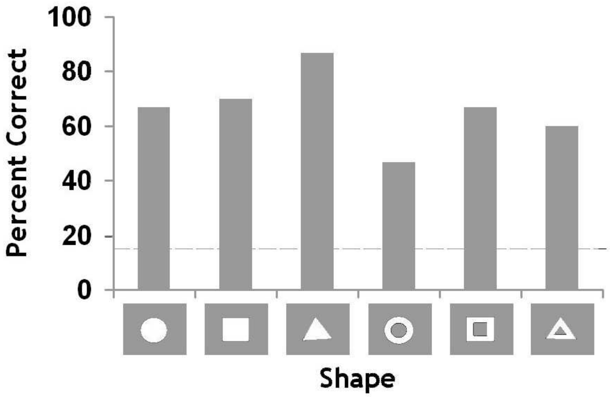

Figure 4: The mean percent correct discrimination performance for each of the six shapes in the SoundView condition of Experiment 3. The broken line represents chance performance.

10.149, $p < 0.001$. As such these results corroborate the general findings of the previous experiments by indicating that participants were able to discriminate between the shapes. Note that the overall discrimination accuracy was much higher in the present experiment (66.2%) than in Experiment 2 (38.3%). This difference in performance across the two experiments can be explained by the fact that accuracy typically increases as the number of alternatives in a forced choice test decreases.

The mean percent discrimination score for each of the six shapes, averaged across participants is shown in Figure 4. A one-way repeated measures ANOVA revealed that discrimination performance was not equivalent across the shapes, F(5, 145) = 2.939, MSE = 0.174, $p = 0.016$. The difference in discrimination performance among the shapes was likely due to the relatively high discrimination of the triangle and the relatively low discrimination of the circle with a hole. However, a series of single sample t-tests revealed that discrimination of the each of the shapes was substantially above chance, all $ts > 3.245$, all $ps < 0.004$.

The overall percent correct for discriminating the shapes using The vOICe was 31.0% (SD = 28.3%). This overall average performance was significantly greater than chance performance (16.6%), t(29) = 2.512, $p = 0.02$. The mean percent correct discrimination scores for each of the six shapes, averaged across participants, are shown in Figure 5. Inspection of Figure 5 suggests that some of the shapes were more difficult to discriminate than others, F(5, 145) = 6.015, MSE = 0.136, $p < 0.001$. A series of single sample t-tests revealed that only the square and the square with a hole were discriminated significantly above chance, $ts > 2.242$, $ps < 0.034$. The discrimination of the remaining shapes did not differ from chance, $ts < 1.912$, $ps > 0.07$. These results suggest that it was very difficult to discriminate the shapes using auditory information provided by The vOICe.

A direct comparison of the overall discrimination performance using SoundView and The vOICe is shown in Figure 6. Inspection of the figure revealed that the overall percent correct discrimination in the SoundView condition (mean = 66.2%) was more than double the percent correct discrimination in The vOICe condition (mean = 31%). An independent sample t-test confirmed that the

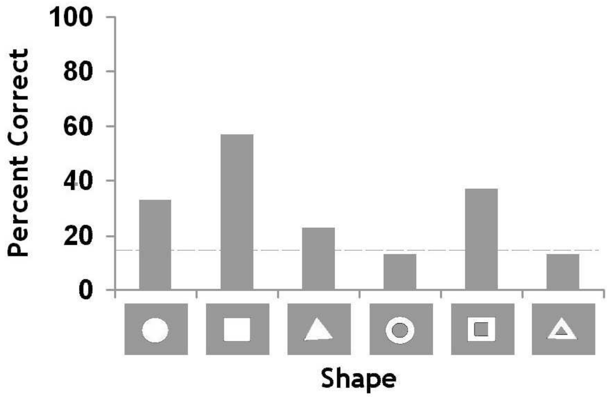

Figure 5: The mean percent correct discrimination performance for each of the six shapes in The vOICe condition of Experiment 3. The broken line represents chance performance.

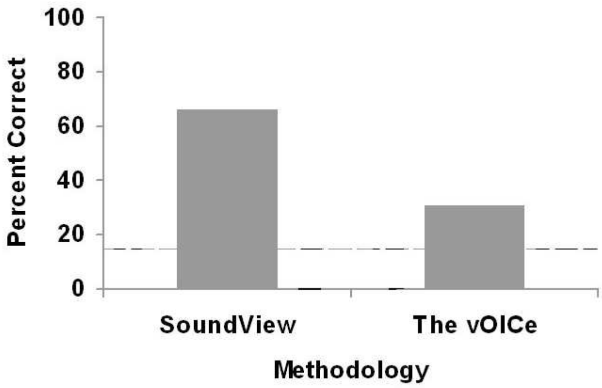

Figure 6: The overall mean percent correct discrimination performance for participants using SoundView and The vOICe in Experiment 3.

discrimination performance was much greater for the group of participants who used SoundView than those who used The vOICe, t(58) = 4.921, $p < 0.001$. A comparison of the two conditions for each shape separately revealed that participants performed better using SoundView than The vOICe on all of the shapes ($ts > 2.953$, $ps < 0.006$) except for the square, for which performance levels did not differ significantly across conditions, t(55) = 1.121, p = 0.267. In general, therefore, the results lead us to conclude that relatively novice users are able to discriminate simple shapes more effectively using SoundView than using The vOICe.

4. COMPARING AUDITORY AND HAPTIC FEEDBACK

Various attempts have been made to make visual information available through haptic devices. In [7] a haptic device for the display of 3D objects and textures was described and user studies on blind and sighted people were performed to determine their ability to determine object properties such as size, angles, and roughness. Complex object recognition was also investigated. User studies investigating the ability of blind people to use a haptic device to perform various task were presented in [8]. The TACTICS system described in [9] allows the printing of complex images as tactile maps on microcapsule paper. It was found that preprocessing the images by edge detection and enhancement resulted in greatly improved performance in recognition tasks. Attempts to augment the haptic display with auditory information are described in [10], where line graphs are displayed through a combination of haptics and sound. Multimodal perception of roughness textures through sound and haptics is described in [11]. Roughness is displayed aurally by piano tones of various frequencies.

SoundView uses a scraping metaphor, yet only renders the audio associated with this action. However if we scrape an object in reality we hear a sound and we feel the surface texture. It therefore seems reasonable to assume that the addition of haptic feedback should improve the performance of the system. On the other hand, perhaps the haptic and auditory channels redundantly encode information in this task, in which case the performance should not change.

Another experiment was performed in an effort to answer these questions. The goals were to identify which of these types of feedback is most useful and/or preferable to participants in recognizing geometric shapes. The extension to SoundView developed for this experiment will be referred to herein as SHView for ”Sound and Haptic View”.

4.1. SHView

Some changes were made to SoundView to enable a proper comparison between the alternate feedback modes. A tiny vibrotactile display, or buzzer, was used to provide haptic feedback. Initially, the buzzer was intended to be coupled to the stylus, but it could only be operated at audible frequencies, and made internal parts of the stylus audibly resonate. Instead the buzzer was worn by participants on their dominant wrist, attached via an elastic wristband, as shown in Figure 7. The sensation of the buzzer could be likened to a mobile phone vibrator. The buzzer was interfaced with SoundView via Java Native Interface (JNI) through a PCI I/O board and an external power source. Since the buzzer provided only on/off feedback, SoundView’s original source code was altered such that the system produced audio feedback with a constant frequency composition when the stylus was held inside the shape, and no feedback when it was outside. This differed from SoundView in that no scraping motion was required to produce sound. As a result of these modifications, the system provided comparable binary haptic and audio feedback.

Figure 7: Experimental Setup. Participant is wearing the buzzer on a wristband attached to his wrist. Note the presence of the sound blocking headphones and the occlusion of the participant’s view of his hand by the box.

As a first step, a pilot study aimed at determining if haptic feedback could be effective in a basic geometric shape recognition task was completed. Six participants were asked to recognize a random set of 4 out of a possible 6 basic geometric shapes (a circle, a hollow circle, a square, a hollow square, a triangle, and a hollow triangle). During our pilot sessions, participants were not able to view their hand movements, which ensured that no visual feedback could be aiding the participant in the task.

Shape recognition times varied from 40 seconds to 219 seconds, with an overall error rate of 7/24. However, this large error rate was primarily due to one outlier, a participant who committed 3 errors on 4 trials. In informal interviews, participants found the system generally reliable and pleasant to use. Overall, the pilot study results showed that haptic feedback can enable a participant to perform a basic geometric shape recognition task.

The pilot study revealed that there was an occasional audible difference between the on/off states of the buzzer. Participants in the final study were asked to wear ambient sound blocking headphones during the experiment to ensure that a participant would not be relying on any ambient audible cues from the buzzer. To compensate for the effect of the headphones, the volume of the audio output for the sound condition was adjusted so that it was clearly audible to the participant.

A cardboard box was used to occlude the participant’s dominant hand from view while performing the task, thus eliminating visual feedback from the experiment. The box was constructed such that the range of motion of the dominant hand was not restricted by the box’s walls. The box featured a side opening hidden from participants’ view, which allowed the investigators to monitor the proper functioning of the haptic buzzer, and to observe how participants performed the experimental task (as shown in Figure 7).

The outer bounds of the tablet were marked using tape to indicate the active area of the tablet. This was necessary as, without a physical boundary, participants could find themselves exploring the inactive outer area of the tablet while looking for the shape.

4.2. Methods

A controlled experiment was run using the SHView system to examine the effects of audio, haptic, and combined feedback on a person's ability to recognize basic geometric shapes. Participants were shown a shape sheet that contained a picture of each of the four shapes that they would be asked to identify during the experiment: a circle, a square, a triangle, and a rectangle. The first three corresponded to shapes used in the pilot experiment. Based on the results of the pilot, shapes with holes were found to be recognizable by all participants, and therefore were omitted from the experiment. The rectangle was added to compensate for this omission. Participants were instructed on the use of both the audio and haptic feedback devices. Participants were also told that feedback would occur when the stylus was placed inside a shape, and that no feedback would occur when the stylus was placed outside a shape.

Prior to the experiment, participants were given up to three to explore a training shape, at first while being able to observe their hand movements as well as cursor movements on the monitor, and subsequently under the conditions of the experiment. These included addition of sound blocking headphones, occlusion of the participant's hand, and blocking of the monitor to prevent the participant from observing cursor movements.

After the training stage each participant was instructed to execute three sets of shape recognition tasks, one with each feedback mode, and that the tasks would involve the shapes previously seen on the shape sheet. The shape sheet was removed for the duration of the experiment.

During the experiment participants were instructed that they would have up to 90 seconds to explore each shape, but that they should verbally provide an answer as soon as they were certain. If participants were unable to discern the shape after 90 seconds, they were instructed to give the investigators their best guess. Two investigators were present in the room while participants conducted the experimental task. One operated the system and loaded shapes for the task, while the other interacted with the participant and recorded results.

Following the experiment, each participant was asked to complete two questionnaires, the results of which are summarized in Section 4.4.

4.3. Experimental Design

Twelve graduate students at the University of British Columbia participated in this study. None of the participants took part in either of the previous experiments. Each participant was involved for approximately forty-five minutes. All participants reported normal hearing and normal or corrected to normal vision. No compensation was offered to participants for their time.

Participants were asked to complete the shape recognition task under auditory, haptic, and combined auditory and haptic feedback conditions. The dependent variables recognition time, defined as the time to complete the recognition task for a shape (up to 90 seconds), and error rate were used to assess participants' efficiency and accuracy.

The experiment utilized a fully-crossed, within-subjects design, with four shapes rendered in a random sequence for each of the three conditions. To prevent participants from guessing the last shape in a condition by process of elimination, one of the shapes was randomly inserted twice, for a total of five trials per condition. One of the two occurrences of the shape was then randomly removed from experimental results.

4.4. Experimental Results

Two 3x4 repeated-measures analyses of variance (ANOVA) were used to evaluate the statistical significance of the main effects of feedback mode and shape and any interaction between these independent variables with respect to recognition time and error rate. The mean scores and standard errors for each of these, averaged across participants, are given in Figure 8. Due to the exploratory nature of this work, results with a significance level of $.05 \leq \alpha \leq .10$ are reported below as borderline significant, as indication of possible trends. Mauchly's test of sphericity was non-significant for all four main effects and both interaction effects, indicating normally distributed data. Results for recognition time indicate a non-significant main effect of feedback mode (p=.783) and a non-significant interaction between feedback mode and shape (p=.514). However, there was a significant main effect of shape (F(3,123)=5.154, p=.005). Similarly, results for error rate indicate a non-significant main effect of feedback mode (p=.734) and a non-significant interaction between feedback mode and shape (p=.460), but a significant main effect of shape (F(3,123)=4.333, p=.011). Posthoc pair wise comparisons using the Bonferroni error correction method showed no significant pair wise differences between shapes with respect to recognition time. With respect to error rate, a significant pair wise difference was detected between the rectangle and the circle (p=.006), and a borderline significant one between the rectangle and the square (p=.077).

To determine the presence of learning or fatigue effects, 3x4 repeated-measures ANOVA's were also performed, with block sequence (1, 2, or 3) rather than feedback type coded as an independent variable, for each dependent variable (recognition time and error rate). For both recognition time and error rate, both main effects of block sequence (p=.359 and p=.480, respectively) and interaction effects of block sequence and shape (p=.974 and p=.137, respectively) were not significant, indicating no clear evidence of learning or fatigue effects.

Inspection of Figure 8 suggests that the difficulty of the recognition task varies by shape, even for simple, well-known geometric shapes. This was most evident in results for error rate, where mean values by shape varied considerably. However, there was no evidence that any of the feedback modes lead to consistently better performance for all shapes. These results suggest that feedback mode is less important than inherent shape properties and individual ability in determining performance in a non-visual shape recognition task.

Participants were asked to rank the conditions with respect to usability and preference after completing the experiment. All conditions were found to be useable by at least nine of the 12 participants. However, in terms of preference haptic feedback was ranked last by all but three participants. Audio feedback was ranked first and combined feedback second by the majority of participants. Participants' comments regarding their perception of the feedback conditions were also collected and are summarized in Table 1.

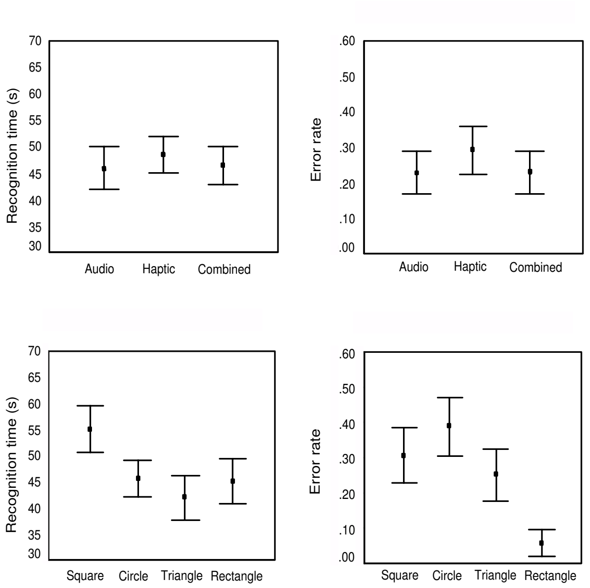

Figure 8: Mean and standard error results: Top-left) Recognition time by feedback mode; Top-right) Error rate by feedback mode; Bottom-left) Recognition time by shape; Bottom-right) Error rate by shape.

User Request	**Support**
Haptic feedback is unpleasant in current form	5 (41.7%)
Audio feedback is unpleasant in current form	2 (16.7%)
Would prefer haptic feedback in the stylus	5 (41.7%)
Would prefer haptic feedback at fingertips	3 (25.0%)
Would prefer haptic feedback in tablet	2 (25.0%)

Table 1: Tabulated questionnaire responses.

5. CONCLUSIONS

The results of the user studies show that the SoundView system does allow users to detect simple black and white shapes with much better than chance performance. [Quote 6-6 and 6-18 recognition rates]. Objects with curved boundaries are more difficult to detect than polygonal objects, possibly due to the linear scraping motion which is the most common exploration motion used. The difference in performance between the eighteen alternative forced choice experiment and the six alternative is undoubtedly due to the confusion of similar shapes.

The comparison with The vOICe shows that SoundView performs better on untrained, sighted subjects. Because The vOICe by nature requires more training than SoundView this result does not necessarily indicate the superiority of SoundView for trained users.

To determine if the addition of haptic feedback has the potential of improving the performance of the system, we developed SHView, which adds haptic feedback in the form of a buzzer worn on the wrist. We found no significant difference between performance using auditory feedback alone, haptic feedback alone, or both combined. Although most participants preferred to use audio feedback, this can be attributed to a perceived need for improvement of the haptic feedback mechanism. Further studies are required to assess the usability of different implementations of audio and haptic feedback modes. Perhaps if the haptic feedback and auditory feedback are designed to provide complementary rather than redundant information the performance does improve. One possible approach would be to use haptic feedback to detect edges in the images. This would correspond more closely to real exploration of shapes by touch.

We believe the overall results of these studies are encouraging. Future studies are clearly required before the system can be considered a practical vision substitute. We are currently attempting to understand the exploration strategies adapted by participants to explore the images and how they relate to the performance by capturing and analyzing the motion. Clinical studies on blind subjects are planned in he near future.

Acknowledgements

This works was supported by IRIS/PRECARN and under NSERC Research Grant 84306.

6. REFERENCES

[1] Richard E. Cytowic, *SYNESTHESIA, A Union of the Senses*, The MIT Press, Cambridge, Massachusetts, 2002.

[2] Kees van den Doel, "SoundView: Sensing Color Images by Kinesthetic Audio," in *Proceedings of the International Conference on Auditory Display*, Boston, 2003.

[3] P. B. L. Meijer, "An experimental system for auditory image representations," *IEEE Transactions on Biomedical Engineering*, vol. 39, no. 2, pp. 112 – 121, January-March 1992.

[4] Peter Meijer, "The vOICe - Seeing with Sound, http://www.visualprosthesis.com," 2003.

[5] P. B. L. Meijer, "Seeing with Sound for the Blind: Is it Vision?," in *invited presentation at the Tucson 2002 conference on Consciousness, April 8-12, 2002. Abstract no. 187 in "Toward a Science of Consciousness" , in Consciousness Research Abstracts (a service from the Journal of Consciousness Studies)*, Tucson, Arizona, USA, 2002, p. 83, Oxford University Press.

[6] K. C. Armel and V.S. Ramachandran, "Acquired synesthesia in retinitis pigmentosa," *Neurocase*, vol. 5, pp. 2993 – 296, 1999.

[7] C. Colwell, H. Petrie, and D. Kornbrot, "Use of a haptic device by blind and sighted people: perception of virtual textures and objects," in *I. Placencia-Porrero and E. Ballabio (Eds.), Improving the quality of life for the European citizen: technology for inclusive design and equality: http://phoenix.herts.ac.uk/SDRU/pubs/TIDE/colwell2.html*, IOS Press, Amsterdam, The Netherlands, 1998.

[8] Calle Sjöström, "Designing Haptic Computer Interfaces for Blind People," in *Sixth International Symposium on Signal Processing and its Applications*, Kuala-Lumpur, Malaysia, 2001.

[9] J. P. Fritz, T. P. Way, and K. E. Barner, "Haptic Representation of Scientific Data for Visually Impaired or Blind Persons," in *Proceedings of the Eleventh Annual Technology*

and Persons with Disabilities Conference, California State University, Northridge, Los Angeles, 1996.

[10] R. Ramloll, W. Yu, S. Brewster, B. Riedel, M. Burton, and G. Dimigen, "Constructing sonified haptic line graphs for the blind student: first steps," in *The fourth international ACM conference on Assistive technologies*, Arlington, VA, 2000.

[11] M.R. McGee, P.D. Gray, and S.A. Brewster, "Feeling Rough: Multimodal Perception of Virtual Roughness," in *Proceedings of Eurohaptics*, Birmingham, UK, 2001.