

SoundBar: Exploiting Multiple Views in Multimodal Graph Browsing

David K. McGookin
The MultiVis Project
Department of Computing Science
University of Glasgow
Glasgow, G12 8QQ
mcgookdk@dcs.gla.ac.uk
www.multivis.org

Stephen A. Brewster
The MultiVis Project
Department of Computing Science
University of Glasgow
Glasgow, G12 8QQ
stephen@dcs.gla.ac.uk
www.multivis.org

ABSTRACT

In this paper we discuss why access to mathematical graphs is problematic for visually impaired people. By a review of graph understanding theory and interviews with visually impaired users, we explain why current non-visual representations are unlikely to provide effective access to graphs. We propose the use of multiple views of the graph, each providing quick access to specific information as a way to improve graph usability. We then introduce a specific multiple view system to improve access to bar graphs called SoundBar which provides an additional quick audio overview of the graph. An evaluation of SoundBar revealed that additional views significantly increased accuracy and reduced time taken in a question answering task.

Author Keywords

Haptics, Non-Speech Audio, Visualisation, Visual Impairment

ACM Classification Keywords

H.5.2. Haptic I/O, Auditory (non-speech) feedback.

INTRODUCTION

Much of the information that we access in our lives is non-textually based. Diagrams, maps, graphs and tables all provide access to information in a faster, more efficient way than textual descriptions. As noted by Kung Fu Tze “*a picture paints a thousand words*”. However for many user groups, accessing such information is challenging. Visually impaired users, of whom there are around 11.4 million in the US [13] and 2 million in the UK (www.nib.org) are one such group. In this paper we shall discuss the problems of access to non-textual graphs, both computer and non-computer based, be-

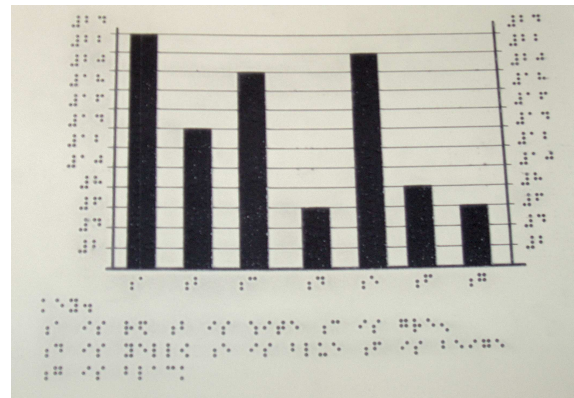


Figure 1. An example raised paper graph used by a visually impaired person. The graph is produced by passing a printout on special paper through a heat printer, creating a raised image on the paper.

fore discussing how these issues may be overcome for graphs using alternate non-visual representations.

GRAPH ACCESS FOR VISUALLY IMPAIRED PEOPLE

Access to graph based information for visually impaired people is a challenging process as graphs must be produced using special materials. The most common technique is to use swell paper, where the graph is printed using a conventional printer on a sheet of special paper. This paper is then passed through a thermal printer, causing the print to rise off of the surface of the paper, creating a tactile relief that can be felt by a visually impaired person. Whilst this does provide access to graphs, it is not without problems. As can be seen from the graph shown in Figure 1, it is not possible using braille to write the name of each bar directly underneath the bar, as would be the case using a visual graph. Each bar therefore is represented by a single digit with a key containing the bar names underneath the graph, creating a level of indirection between each bar and the information that bar represents.

Additionally, raised paper representations are not useful ways to access dynamically changing data graphs. For exam-



Figure 2. A picture of a SensAble Technologies' PHANTOM Omni Haptic Device (www.sensable.com) as used in SoundBar experiment. The user interacts with the "pen" which also contains two buttons for interaction.

ple, a user trying to keep up-to-date with sales information held in a company spreadsheet would need to print the graph on swell paper and pass that printout through a heat printer just to view the graph. Obviously this hard copy would not automatically update and further printouts would be required to inspect the graph at future points. These problems also extend to visually impaired users who are constructing graphs in applications such as Microsoft Excel, who will have no other way to inspect graphs after they have been produced. This can be a particular frustration for those with incomplete knowledge of graphs, such when learning in school. In such situations mistakes are more common, and several iterations may be required to create a correct graph, meaning that a pupil may need to work harder simply to accomplish the same level of performance as a sighted peer.

Computer Based Graph Browsing

In order to overcome these problems, several researchers have investigated ways to present graphs to visually impaired users using computer based haptic and non-speech sound technology. Mansur [8] created a technique for presenting line graphs using non-speech sound, by mapping the value of the y-axis to a musical note. This technique has been further investigated [3, 16], and found to be effective at communicating information.

van Scoy and Kawai [15] developed a tool which allowed users to feel the graph produced by a mathematical equation, by "carving out" the function in a virtual haptic surface, which could be felt by the participant using a PHANTOM haptic device (see Figure 2). Unfortunately, they do not report an evaluation of their system. Yu and Brewster [19] have carried out extensive work developing applications, again using a PHANTOM device, to allow visually impaired people access to graphs. Their work has covered several graph types (such as line and bar graphs and pie charts), as well as

creating guidelines for producing virtual haptic graphs [19]. Their evaluations, which asked users to answer questions about the graphs, showed that whilst virtual haptic graphs were more accurately interpreted, the time taken to extract information was significantly longer than with a raised paper version. This increase in time was attributed to the single point of contact with the graph that the PHANTOM allows, in comparison to contact with all of the fingers that can be used with a standard tactile diagram.

Attitudes to Graphs

Whilst the work discussed above shows that visually impaired people can extract information from graphs, it does not consider whether visually impaired people feel that graphs are a useful means of communicating information. To identify how visually impaired users feel and use graphs, we carried out a group discussion at the Royal National College for the Blind (RNCB) Hereford, UK. The group discussion contained six participants who were all registered blind. The participants were between the ages of 18-30, and were a mixture of both students and staff. Two individual interviews were also carried out with students in the same age range. Participants were asked to discuss several topics around their experiences of graphs in both education and wider life. All participants could see the point of graphs, and could extract information from them although participants did not find graphs to be effective ways to read information, and did not give them the advantages that a graphical representation would give a sighted person. One participant saying, "*When somebody fully sighted is looking at a screen, they can see it as a whole instantly, but somebody visually impaired has got to navigate actually bit by bit all around the screen. Which is why its so hard to get a complete picture of the screen*". Many participants felt that the graph provided another level of indirection between themselves and the information that they wished to get to. Rather than the representation of the graph assisting with their task, participants felt that the representation "got in the way". Many participants felt so strongly about this that they would actively avoid graphs, with one noting that "*I do avoid them as much as I can*".

Whilst the result of the group discussion shows that visually impaired people do not find graphs a useful way to access information, much of the information that is required for mathematical or scientific study is presented via graphs. Whilst screen readers can be used to read out the values, a graph provides more information than just the values of data points, allowing multiple data to be compared quickly and easily, something that current representations and speech do not allow. Ineffective ways to access this information therefore limits access to careers in these disciplines for visually impaired people [5]. Before discussing how graphs may be made more effective and useful for visually impaired users, we discuss reasons why current graphs for the visually impaired, and graphs accessed via a non-visual modality cause the problems identified in the group discussion.

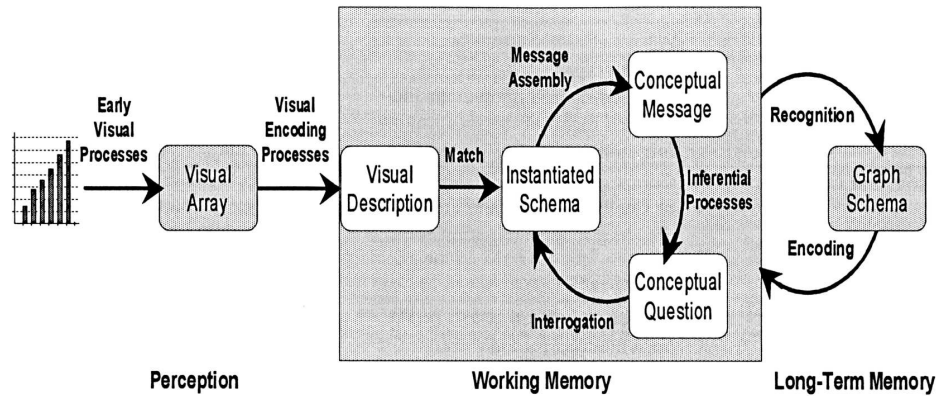


Figure 3. A diagram illustrating the main components of Pinker's theory on graph interpretation. Taken from Lohse [7].

GRAPH BROWSING THEORY

To properly explain the problems with current haptic graph approaches, we must consider how graphs are interpreted. Figure 3 shows a diagrammatic overview of a theory of visual graph perception by Pinker [11]. Pinker's theory considers that the image on the page is perceived and manipulated by the visual system into what he calls a visual description, which describes the primitives of the graph such as rectangle, line etc., the properties of these primitives (e.g. colour, texture etc.) and how they are related to each other (e.g. above, below etc.). This visual description is produced in a few milliseconds, which is important given that the description is held in working (short term) memory, which has both a limited capacity for information, and information contained within it degrades over time [9]. The visual description is then used to fill in a blank "graph schema" which has been selected from long term memory as the best "fit" for the visual description. The graph schema is like a blank form with spaces to fill in information about the graph, such as the bar values etc. The "layout" of the schema makes it easy to perform certain comparisons between data in the graph. For example, the schema may hold the height of each bar in a bar chart and provide easy mental links between the bars to allow a user to quickly identify trend information in the graph. The user interrogates the instantiated graph schema and seeks to answer a question which is likely to rely on only part of the schema. For example, if the user is trying to identify the trend of a graph they will only need to consult relevant parts, such as the relative heights of the bars from left to right. If some data are not available in the graph schema the user will need to trace back to the visual description in order to retrieve the information, with higher cost searches the further back the user must go, ultimately leading to an active visual search of the graph in order to locate the information required. Obviously, whilst such searches are progressing, information that was required and available in the graph schema will degrade. This means that even when the required information is found further searches may be required. As the length of time taken to encode information into the graph schema increases, such as with raised paper

or PHANTOM based solutions, the number of re-searches required increases.

In comparing a number of studies which investigated a number of different visual graph representations Lohse [7] identified that different types of graph more effectively support different types of question. Those representations that minimised the demand on short term memory, by making the information required to answer the question explicit, tended to outperform other representations that did not make such information as obviously available. This has also been noted by Peebles and Cheng [10] who say that although the same set of questions may be answered using two different graph representations (the representations are "informationally equivalent"), it may be easier to extract information from one representation than the other (the representations are not "computationally equivalent"). So if we can "prioritise" the encoding of information relevant to the user's task in non-visual graph browsing by making it easier to get to, the graph will be more useful.

Other Factors Influencing Graph Browsing Efficacy

There are various aspects of both the user and the task which impact how effectively information can be extracted from a graph. Obviously one of the main factors is the user's prior experience of the graph type. The greater the exposure to a particular graph type, the more the schema used to answer questions about the graph can be refined to allow more effective access [7].

Another feature that affects users' ability to extract information from a graph is the type of question that they wish to answer from it. Curcio [4] determined that there are three main types of question that a person would wish to answer using a graph. Each type requiring more information to be contained in the graph schema, or post processing of that information to answer the question.

Literal questions simply require extraction of information contained within the graph. For example a literal question




Figure 4. An example graph as used in the experiment showing different types of food and the number of people who said a particular food was their favourite.

for the graph shown in Figure 4 would be “*How many people said custard was their favourite food?*”. The other two types of question require that users compare the data and manipulate it in some way. These question types are nominally “harder” than the literal type of question since they require information to be retained and processed in short term memory, which is of course shared with the graph schema. One question type, termed “*read between the data*”, occurs when all of the data required to answer the question are explicitly available in the graph but must be manipulated in some way to yield the correct answer. Such a question for the graph in Figure 4 would be “*Which food has a higher popularity, burger and chips or fish and chips?*”. The final question type, “*read beyond the data*” occurs when the data are not explicitly available in the graph and must be constructed from the data available, again using short term memory, to answer the question. An example question of this type for the graph contained in Figure 4 would be “*if the popularity of jelly halved, which would be more popular, jelly or fish and chips?*”. Harder questions require more data and processing of that data, again placing demands on users’ short term memory.

Non-Visual Graph Browsing

Whilst there is a body of work that allows us to consider visual graph understanding, there is no work that considers non-visual browsing of graphs at the same level. Brown, Stevens and Pettifer [2] have considered how node and edge style graphs may be improved however mathematical graphs have not yet been studied. Much of the work that investigates providing non-visual access to graphs simply converts the visual representation of the graph into a raised paper diagram or other haptic representation [15, 19]. However, visual graph representations are optimised to exploit our visual processing capabilities [6], something that, as discussed, is lost when graphs are non-visually presented.

When converting a visual graph representation directly into another modality the resulting graph may have lost much of its ease of use, and be sub-optimal for the modality in which it is presented. Whilst a sighted user can very quickly (in a few milliseconds) move between elements of the graph, it will take several seconds longer for a user to move their finger to inspect each element of a haptic graph. From Pinker’s theory we can consider that haptic graph visualisation has two main problems that must be addressed. Firstly the speed at which information is loaded into a “visual description” needs to be increased. This is a problem given the relatively lower bandwidth of the human haptic system compared to the visual system via all current technologies including raised paper. Since the visual description (and the instantiated schema) are held in working memory, the information already processed by the user will degrade over time, meaning that less can be held at any one time and more reference back to the original source will be required in order for users to answer questions. These problems are only likely to be compounded by increasing the difficulty of questions that a user will wish to answer [4].

We argue that when a graph is converted to another modality, its ability to efficiently communicate information is impaired, making it harder for a visually impaired person to gain the same benefit as would a sighted user. Alternative representations or “multiple views”, each tailored to quickly finding particular information and “fastracking” it into the graph schema, can provide more effective graph access. In the following section we will outline classical visualisation work on multiple views, and discuss why this is a practical and relevant solution to the problem of non-visual graph access, before discussing our first implemented solution: SoundBar.

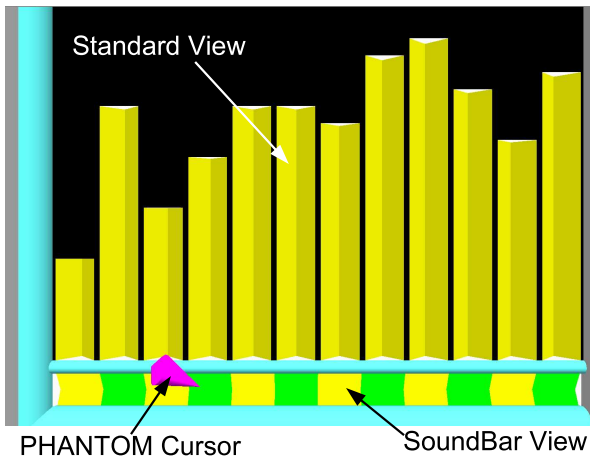


Figure 5. A screenshot of the SoundBar System. Bars are represented as recessed grooves, the SoundBar is located below the bars. When a segment of the SoundBar is touched with the PHANTOM (represented by the cone shaped object), a note proportional to the height of the bar immediately above is played.

MULTIPLE VIEWS

Multiple views have existed in information visualisation research for several years and can be defined as the use of “two or more distinct views to support the investigation of a single conceptual entity” [1]. Such multiple views can offer several advantages over single views, allowing a richer understanding of the data to be gained [12]. However, there are issues that must be taken into account or the advantages of multiple views may be easily lost. Baldonado *et al.* [1] note that to be effective, views must complement each other, should be able to be related to each other but should emphasise different aspects of the same information otherwise any improvements will be lost as users must “context switch” between the different views. Roberts [12] discusses how multiple views can be used to overcome the “data availability paradox” of Woods *et al.* [18], where the amount of information that we have access to is consistently increasing, yet our cognitive abilities to process such information are constant. Based on the discussion of the previous section, we can consider non-visual graph browsing in terms of the “data availability paradox”, where the information available to come to a decision on a task is held constant, but due to the data representation and data integration time required our cognitive abilities are effectively reduced. Whilst we cannot alter the cognitive abilities of the user, we can change the data representation to allow more effective use of those abilities, making multiple views an appropriate technique to consider for non-visual graph browsing, allowing different representations of the data to be presented, each optimising specific types of graph browsing tasks the user may wish.

SOUNDBAR

To investigate if multiple views of a graph can provide effective non-visual graph access, we developed a system to present bar graphs incorporating two distinct views of

the data. One view was a “standard” bar graph view (constructed using the guidelines developed by Yu and Brewster [19]), programmed in SensAble Technologies’ OpenHaptics API and presented via a PHANTOM Omni haptic device. In this view the bars are represented as recessed grooves, as this stops the user from slipping off the bars as is the case with raised lines. The axes are represented as raised cylinders to differentiate them from the bars. All features of the graph can be interrogated by touching them with the PHANTOM pen and pressing one of the buttons on the pen. Doing so presents speech information about that feature using the Microsoft Speech API (www.microsoft.com/speech/download/sdk51/). In the case of the bars and axes, the name is provided. For the graph background, navigation information to find the bars is provided. The entire graph is contained within a box, so the user does not get lost, and all colours used in the visual display are of high contrast to assist users who have any residual sight. The second view of the graph allows users to gain a quick overview via non-speech sound. This second, SoundBar view, is located below the x-axis and was segmented into squares. When the PHANTOM comes into contact with a square, a musical tone is played by a MIDI piano timbre (General MIDI patch No. 000). The pitch of this note being proportional to the height of the bar immediately above that segment (see Figure 5). The value of this MIDI note is calculated based on the formula derived by Brown and Brewster’s work on the SoundVis system [3], leading to a MIDI note value between 35 and 100. As with the standard view bars, touching a square and pressing the PHANTOM button causes the name of that bar to be read out using synthetic speech.

There are several reasons why starting to investigate multiple view presentation of graphs using these two views is appropriate. Whilst the standard graph view is sub-optimal it does present the data in the same way as a sighted colleague would view the data. As Winberg and Bowers [17] identified in a collaborative task between a sighted and visually impaired individual, communication problems can arise if coherent representations are not used for both the sighted and visually impaired user, as a common vocabulary to describe the representation would not exist. We therefore retain the spatial graphical layout of the bars to allow visually impaired and sighted colleagues to discuss the graph with each other. Whilst these views may not support all of the tasks that a user may wish to undertake, if a user can move between the views and integrate the information derived from each view, more views may be identified to further improve understanding of the data contained in graphs.

The SoundBar should provide access to information that visually impaired participants in the group discussion did not feel they could get access to via current graph representations namely a quick overview of the graph. The ability to move the PHANTOM pointer from left to right across the SoundBar is similar to running a finger across a piano keyboard, allowing the user to quickly focus in on areas of interest, excluding immediately bars that do not interest them, thus optimising the transfer of relevant information to short

term memory to accomplish the user's task. This is not something that would be possible if the audio were simply integrated into the existing "standard" view as users would get caught in the recessed grooved bars.

In terms of Pinker, using both views should allow the user to gain a quick overview of the graph, deciding quickly what information is relevant to their task, optimising the transfer of relevant information into an appropriate graph schema. Additionally, the user can quickly decide and safely ignore irrelevant information, meaning less chance of relevant information in the graph schema being lost as short term memory degrades. Additionally, the quick access nature of the SoundBar should allow information to be relocated easily in the graph if it is lost from the graph schema and needs to be retrieved. In this way our approach is consistent with Shneiderman's rule on information browsing "*overview first, zoom and filter, then details on demand*" [14].

EVALUATION OF SOUNDBAR

To determine if SoundBar and multiple views are a useful aid to interrogate non-visual graphs, an evaluation was performed. SoundBar was compared to a haptic bar graph system which was identical to the system previously described, but did not incorporate the SoundBar view. In common with much other work involving systems for visually impaired people it is difficult to get enough visually impaired participants to carry out a statistically meaningful evaluation. We have therefore taken the approach of carrying out an experiment with blindfolded sighted users, before validation of the results with a small group of visually impaired participants.

Procedure

Twelve undergraduates from the University of Glasgow participated in the experiment; each was paid £6 on completion. All participants were aged between 18-30 and comprised of 7 men and 5 women. Participants had to browse graphs, and answer questions about them using both the SoundBar system (SoundBar Condition) and a similar bar graph presentation system which did not contain SoundBar (Standard View Condition). The order in which participants carried out the conditions was counterbalanced to avoid order effects, additionally two sets of graphs and questions were used, and again these were counterbalanced between conditions.

Participants were first asked to read and sign a consent form, before being given a verbal briefing about the experimental procedure. Participants were then familiarised with the PHANTOM Omni device. Participants then were introduced to whichever system they were using in the first condition via a standard sample graph. The experimenter walked through the system with the participant and explained all of the important details. During this time the participant was allowed to look at the computer screen, however during the actual experiment the participant was blindfolded to simulate the effects of blindness as much as possible.

Each graph used in the experiment had 12 bars, and each represented data on common everyday objects which did not require specific domain knowledge to understand. Partici-




Figure 6. A graph illustrating the average scores for correctly answered questions in the standard and SoundBar conditions. Shown with standard deviations.

pants were asked a range of questions based on the different types identified by Curcio [4] described earlier. For example, a typical question for the graph in Figure 4 would be "*The following graph shows the number of people who said that a particular type of food was their favourite. What were the three most popular foods?*". Participants gave their answers verbally which were recorded by the experimenter. Participants were then asked if they were sure of their answers, if participants responded that they were, the time taken was recorded and the next graph and question were presented. On completion of the first condition, informal comments by participants were taken before the second condition was carried out. At the end of the experiment participants were debriefed on the purpose of the experiment and asked for any other comments.

Results

Both the number of correct responses, and time taken for those responses was recorded. Each correct response from the user was allocated a mark, yielding a score out of 12. The average scores for each of the conditions are illustrated in Figure 6. A paired t-test on this data ($t(11)=3.95$, $p=0.002$) showed that there were significantly more correct responses in the SoundBar Condition than in the Standard View Condition.

The time taken by a participant to complete a condition was calculated as a sum of the time taken to answer each question in that condition. The mean time taken for each condition across all participants is shown in Figure 7. A paired t-test on this data again showed significance ($t(11)=3.28$, $p=0.007$), with the SoundBar Condition taking significantly less time to complete than the Standard View Condition.

The results shown clearly indicate that SoundBar is a more accurate and efficient method to retrieve, and compare information in bar graphs than a standard haptic representation. However, these results do not necessarily support the claim of this paper that multiple views of a graph can be of benefit in non-visual graph browsing. Indeed, one possibility is that




Figure 7. A graph illustrating the average time taken by participants to answer questions in the standard and SoundBar conditions. Shown with standard deviations.

participants only used the SoundBar view rather than both together.

Usage Strategies

In order to identify if the use of more than one view was helpful, after carrying out each condition, participants were asked about the strategies they used in answering the questions given. Additionally, a log of the position of the PHANTOM was recorded for each question that the participant answered. This log, sampled at 0.25 second intervals, allows us to reconstruct the exploration of the graph by the participant and uncover the strategies used.

SoundBar Condition Usage

An analysis of all of the logs produced in the SoundBar condition revealed that 54% showed usage of both the views available, 40% showed the use of only one view (largely the SoundBar only), whilst 6% of logs were not usable or had been corrupted. This shows that in most cases there was a benefit in having both of the views available. A further analysis of the cursor logs revealed a number of distinct strategies that participants used.

The most common use of the SoundBar as described by participants was to “filter” the bars, by getting a quick overview of the graphs to identify candidate bars and as such excluding the irrelevant bars to the task. This strategy fits with the objectives of the SoundBar, reducing the amount of information that must be retained in short term memory and thus improving user performance. The use of sound allowing a quick overview of the graph to be obtained and the collocation of the views allowing the user to selectively inspect the relevant bars haptically.

Take for example, the trace shown in Figure 8. Here the participant has been tasked to find the top three selling domestic appliances in 2004 (i.e. find the highest three bars). The participant has used the SoundBar to exclude those bars that are too low, but cannot accurately distinguish between the pitches to be confident about which three bars are the highest.




Figure 8. A screenshot showing a cursor trace for a participant answering a question in the SoundBar condition. The question was to find the three top selling domestic appliances, i.e. the three highest bars. Each “dot” represents a sample point of the PHANTOM cursor.

Because of this the participant has searched from the top of each of the three bars that they think are the highest to ensure their decision is correct. This is a strategy that participants also used in the Standard View Condition, but is more optimal with SoundBar as the participant will start with a high bar rather than the first bar, which may be the lowest.

The same strategy is also shown in Figure 9. This trace shows a participant trying to determine what the major ingredient in a breakfast cereal was (i.e. find the highest bar). Again the participant has used the SoundBar to identify what they think is the highest bar before searching from the top of that bar to confirm their decision. Participants who had used this strategy expressed that they found the SoundBar an effective way to communicate graph information, but were not confident that their answer was correct if there were two bars with a similar value, the difference in pitch not being enough to base a conclusive judgement on.

Standard View Condition Usage

As with the SoundBar condition, trace logs of cursor movement were also recorded for the standard view condition. An analysis of these as well as comments by the participants revealed a number of strategies that support the addition of the SoundBar view to bar graph browsing.

All participants commented that they found the standard view condition to be more difficult than the SoundBar condition. Participants noted that they had less confidence that their answers were correct than when using the SoundBar. One reason for this mentioned by participants, and confirmed with the trace logs, is the possibility to easily miss out bars (see Figure 10). Since bars can be at any height in the graph, there are only a few strategies that can be employed which will ensure that all bars are considered.

One strategy is an exhaustive search, where the user moves from left to right in the graph, moving to the top of each




Figure 9. Another screenshot illustrating a participant using the SoundBar as a quick overview mechanism, before using the standard graph view to make a final decision. The graph represents ingredients in a breakfast cereal and the participant is tasked to find the main ingredient (i.e. highest bar).

bar and then back to the base, before moving onto the next bar (see Figure 11). This strategy was the most commonly used, but suffers from needing to inspect all of the bars. Bars cannot be excluded such as when a graph is visually presented. There is no data available, until after detailed inspection of the bar, as to whether it should be excluded or not. This means all information must be slowly transferred to the graph schema, with irrelevant information being held there as well as the information relevant to the user's task. Due to the limitations of short term memory, information that is required to answer the task may be lost or degrade in the time taken to perform an exhaustive graph search. Browsing the graph in this way excludes the point of using a graph at all, with the representation simple making it more cumbersome to extract data (as previously discussed by participants in the group discussion).

A final strategy employed by participants, and only in questions where the objective was to find the highest bar, was to move to the top of the first bar and move right from the top of that bar until another bar was located. An adaptation of this strategy was also used in the SoundBar condition (see Figure 8). This strategy was noted to be useful if the highest bar(s) was being found, however not for the lowest bar, where higher bars would interfere. The ability of this strategy to filter irrelevant bars is also dependant on the relative heights of the bars. If the bars uniformly increase from left to right, then the highest bar will not be found until all the bars have been investigated.

Discussion

The evaluation on blindfolded sighted users has shown that the incorporation of multiple views in non-visual graph browsing can significantly improve user performance. From the discussions with participants and analysis of the cursor traces for both conditions, the conversion of the standard visual bar graph representation to the haptic modality introdu-

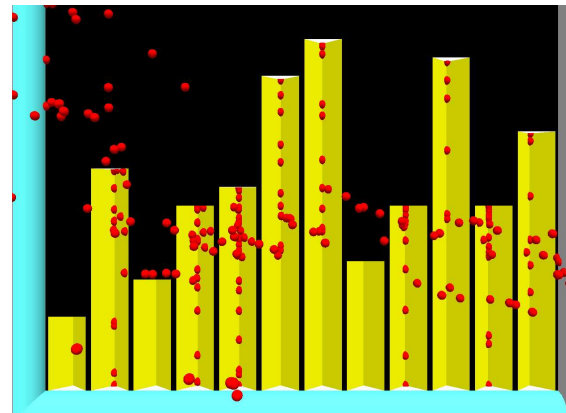


Figure 10. A screenshot showing a cursor trace from the Standard View Condition. The participant was tasked with finding the three lowest bars. The strategy adopted has caused the participant to fail to locate any of the three required bars.

ces a number of problems. It is easy for users to miss bars, and as such miss potentially important data relevant to their task. Additionally, as the time to extract information from the graph using the PHANTOM is longer than when doing so visually, the linear search through the graph places demands on short term memory, and as such by the discussion of Pinker earlier, means that information will be lost from the graph schema. Whilst participants identified strategies to improve performance above looking at the height of every bar (e.g. moving from the top of one bar to the next highest bar to the right), these strategies are not always applicable, and may not yield improvements in performance. The incorporation of SoundBar allowed participants to get a quick overview of the graph. Whilst in many cases participants were unable to answer questions directly from the SoundBar, it did provide a filtering mechanism to quickly exclude bars. These particular views allowing the problems identified through the analysis of Pinker's graph theory to be overcome. SoundBar allows information that is and is not relevant to the user's task to be quickly filtered, thereby optimising the limited available bandwidth. As only the relevant bars are encoded into the graph schema, the time taken from instantiating the schema to answering the question is reduced. This reduces the likelihood that relevant information already in the schema will have degraded, and will not require first class searches of the graph to be retrieved. The two views available in the SoundBar Condition provided different representations of the data, each suited to different uses, but when used together providing a better insight into the data allowing tasks to be more effectively carried out.

Whilst these results show multiple views of non-visual graphs can be effective, they do not show improvement with the target user group of this work: visually impaired users. In the remainder of this paper we will discuss work which has evaluated the described multiple views with visually impaired participants to determine if there are improvements in graph browsing with this user group.




Figure 11. A screenshot illustrating the most common technique used in the Standard View Condition. The participant moves from left to right, inspecting the height of each bar in turn.




Figure 12. A graph illustrating the average scores for correctly answered questions in the standard and SoundBar conditions. Shown with standard deviations.

EVALUATION OF SOUNDBAR WITH VISUALLY IMPAIRED PEOPLE

As already stated, the evaluation reported on SoundBar was performed on blindfolded sighted participants. In order to determine the value to visually impaired people of multiple views, the experiment was re-run using four participants at the RNCB Hereford. Participants comprised of three males and one female between the ages of 18 and 30. All were registered blind and were paid £15 for participating in the experiment. Since it is inappropriate to perform statistical analysis on such a small population, we do not report a formal statistical analysis here. However graphs of correct answers and time taken (both calculated in the same way as the blindfolded experiments) are shown in Figures 12 and 13 respectively.

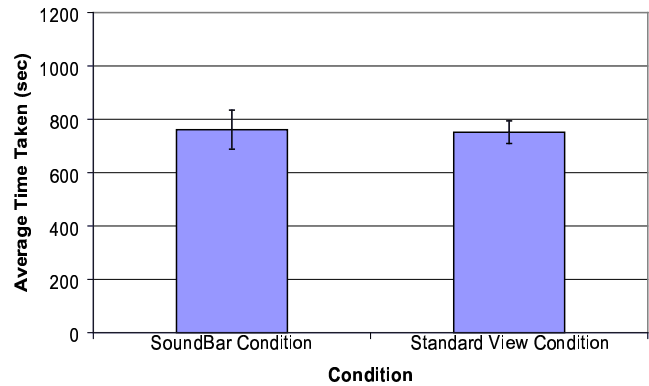


Figure 13. A graph illustrating the average time taken by participants to answer questions in the standard and SoundBar conditions. Shown with standard deviations.

The results indicate that SoundBar provides improvements in accuracy for answering questions, which is consistent with the results of the blindfolded study. The time taken for both conditions is similar. The proportion of cursor traces that showed the use of both views is lower than in the blindfolded experiment. 65% of trials showed use of only SoundBar, 26% showing use of both views, with 9% of logs malfunctioning.

In the Standard View Condition, as with the blindfolded study, there were several instances of participants missing out bars completely in the graph. This would seem to confirm the comments by the blindfolded sighted participants that this is a problem which the SoundBar overcomes.

Improving SoundBar

Overall participants were positive about the inclusion of the SoundBar and felt that it was a useful addition to graph browsing. Participants also mentioned several problems that they had, which could benefit from the inclusion of other views of the graph. Notably participants mentioned that the ability to compare bars which were not consecutively presented was problematic with intermediate notes or bars getting in the way. They discussed how the SoundBar view could be filtered, either presenting a threshold tone, which any bar under a certain value would be given. Or switching off tones for certain bars, so only those bars which the user was interested in would be heard for comparison. These views are consistent with Shneiderman’s mantra on information visualisation “*overview first, zoom and filter, then details on demand*” [14]. The SoundBar providing the overview, whilst the potential additions described by participants providing the filtering.

CONCLUSIONS

Graph Browsing is about more than simply reading out the values of the graph. Yet for many visually impaired people this is how they must access computer based graphs via screen reading software. Whilst it is possible to print off special raised paper copies of the graph, this, as confirmed by visually impaired people, is a difficult way to extract infor-

mation which does not provide all of the advantages of the visual representation. Non-speech audio and haptic interaction provides an opportunity to improve access to non-visual based graphs. From the results of both the evaluations we can conclude that such access to graphs can be improved by the use of multiple different representations, each optimising one particular aspect of the user's task. Our initial investigation has revealed that incorporating the SoundBar, providing a quick overview of the graph, may improve both the number of correct answers to questions, and reduces the time taken to answer those questions. The strategies used by participants are consistent with those expected from the visual graph browsing work of Pinker [11]. Gaining a rapid overview allowed participants to decide quickly what information in the graph was, and was not, relevant to their task. Information that had to be retained in the graph schema to answer the set questions could then be prioritised over the limited bandwidth available. Therefore less information had to be retained in the schema for a shorter period of time. This being one of the main issues identified from Pinker's work. As such, more effective use of short term memory was being made, with less information being lost, necessitating less re-searching of the graph which is extremely costly.

Whilst we have shown the utility of multiple views for non-visual graph browsing, the work is currently limited to bar graphs. Future work will investigate improving access to other types of graph such as line and pie. In doing so access to non-visual visualisations can be significantly improved for visually impaired people.

ACKNOWLEDGMENTS

We would like to thank all the RNCB and RBS. This work is supported by EPSRC grant number GR/S86150/01.

REFERENCES

- Baldonado, M. Q. W., Woodruff, A., and Kuchinsky, A. Guidelines for using multiple views in information visualization. In *AVI 2000* (Palermo, Italy, 2000), ACM Press, pp. 110–119.
- Brown, A., Stevens, R., and Pettifer, S. Issues in the non-visual presentation of graph based diagrams. In *International Conference in Information Visualisation 2004* (2004), IEEE.
- Brown, L., and Brewster, S. A. Drawing by ear: Interpreting sonified line graphs. In *Proceedings of ICAD 2003* (Boston, Massachusetts, 2003), ICAD, pp. 152–156.
- Curcio, F. R. Comprehension of mathematical relationships expressed in graphs. *Journal of Research in Mathematics Education* 18, 5 (1987), 382–393.
- Dimigen, G., Scott, F., Thackeray, F., Pimm, M., and Roy, A. W. N. Career expectations of british visually impaired students who are of school-leaving age. *Journal of Visual Impairment and Blindness* 87 (1993), 209–210.
- Kosslyn, S. M. Understanding charts and graphs. *Applied Cognitive Psychology* 3, 3 (1989), 185–226.
- Lohse, G. L. Models of graphical perception. In *Handbook of Human-Computer Interaction*, M. Helander, T. Landauer, and P. Prabu, Eds., 2 ed. Elsevier, Amsterdam, 1997, pp. 107–135.
- Mansur, D. L. *Graphs in Sound: A Numerical Data Analysis Method for the Blind*. M.sc., University of California, 1985.
- Miller, G. A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review* 63, 1 (1956), 81–97.
- Peebles, D., and Cheng, P. C. H. Modeling the effect of task and graphical representation on response latency in a graph reading task. *Human Factors* 45, 1 (2003), 28–46.
- Pinker, S. A theory of graph comprehension. In *Artificial Intelligence and the Future of Testing*, R. Freedle, Ed. Laurence Erlbaum Associates, Hillsdale, New Jersey, 1990, pp. 73–126.
- Roberts, J. C. On encouraging multiple views for visualization. In *Information Visualization IV'98* (London, 1998), IEEE, pp. 8–14.
- Ross, D. A., and Blasch, B. B. Development of a wearable computer orientation system. *Personal and Ubiquitous Computing* 6, 1 (2002), 49–63.
- Shneiderman, B. *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, 2nd ed. Addison-Wesley, Reading, 1998.
- van Scoy, F. L., Kawai, T., Darrach, M., and Rash, C. Haptic display of mathematical functions for teaching mathematics to students with vision disabilities. *Haptic Human-Computer Interaction LNCS 2058* (2001), 31–40.
- Walker, B. N., and Cothran, J. T. Sonification sandbox: A graphical toolkit for auditory graphs. In *ICAD 2003* (Boston, MA, 2003), ICAD, pp. 161–163.
- Winberg, F., and Bowers, J. Assembling the senses: Towards the design of cooperative interfaces for visually impaired users. In *Proc CSCW 2004* (Chicago, Illinois, 2004), vol. 1, ACM Press, pp. 332–341.
- Woods, D. D., Patterson, E. S., and Roth, E. M. Can we ever escape from data overload? a cognitive systems diagnosis. *Cognition, Technology and Work* 4 (2002), 22–36.
- Yu, W., and Brewster, S. A. Evaluation of multimodal graphs for blind people. *Journal of Universal Access in the Information Society* 2, 2 (2003), 105–124.