

Exploratory strategies and procedures to obtain non-visual overviews using TableVis

Johan Kildal, MSc(IT), Ing Sup, BScME and Stephen A Brewster, PhD, MBCS, CENG

Glasgow Interactive Systems Group, Department of Computing Science, University of Glasgow, Glasgow, United Kingdom

Abstract: TableVis was developed to support computer users, who are blind or visually impaired, in tasks that involve obtaining quick overviews of tabular data sets. An iterative design methodology led to a successful solution to this problem, as derived from evaluations of the interface and its associated techniques for interactive data sonification and support of exploratory processes. However, complex systems can be used creatively in a variety of ways, which are often different from those intended by their designers. Thus, it can be difficult to determine the specific reasons that make a complex interface successful. In an attempt to find a way around these limitations in the most commonly used usability test methodologies, the exploratory strategies and procedures employed by users in their evaluations of TableVis were analysed. These analyses provided a better understanding of the challenges in exploratory tasks from the perspective of users with visual impairments, and helped to identify both best practices and common problems, which form the basis for the next steps to be taken in this line of research.

Keywords: Visual impairment, disability, United Kingdom

Correspondence: Johan Kildal, MSc(IT), Ing Sup, BScME, Glasgow Interactive Systems Group, Department of Computing Science, University of Glasgow, 17 Lilybank Gardens, Glasgow, G12 8QQ, UK. E-mail: johank@dcs.gla.ac.uk. Website: <http://www.multivis.org>

Submitted: April 01, 2006. **Revised:** May 05, 2006. **Accepted:** May 07, 2006.

INTRODUCTION

Advances in information technology in the last decades have contributed enormously to making information more accessible for users who are blind or visually impaired. With computers, accessible media can be generated dynamically in several forms and formats that can often be produced without the need for assistance by the target users. Screen reading technology is an example of a particularly successful solution. The convergence of the advent of screen reading technology with the generalisation of personal computers for work, study and leisure, and the democratization of information availability through the internet, has drawn a new scenario in which information is, by and large, far more accessible in non-visual forms. Other computer-based technologies that can generate accessible content dynamically have also appeared, such as refreshable Braille displays, pin array-based tactile displays, and force feedback devices.

Although such technologies provide access to computer based information non-visually, they do so by presenting information in full detail but with little or no support to obtain overview information quickly and easily at the beginning of the exploration of a new data set. This drawback can be a serious limitation when exploring information that is encountered for the first time. As Shneiderman defends (1), any data exploration starts by obtaining an overview, to be followed by zooming and filtering operations. Other authors also concluded that obtaining overview information is the first step in a data browsing process, for different types of data (2-5) Detailed information should be obtained only on demand, once an

object of interest in the whole set of information has been identified, to avoid overloading the user's working memory. The accessibility technologies mentioned above provide fully detailed information directly, with little or no support to obtain initial overview information. Obtaining overview information non-visually is particularly difficult with numerical data sets like spreadsheets. Accessing numerical values sequentially (as happens when using screen readers) soon saturates the user's working memory (6), and not enough comparisons can be performed between numerical values to reveal overview information.

Not being able to obtain overview information from numerical data sets non-visually has serious implications for visually impaired students and professionals. They cannot freely choose to undertake higher education studies or career paths in disciplines that are numerically intensive, for which collaboration regarding complex data sets is required, as well as keeping up with fast-paced data analysis.

Properties of human auditory perception have been successfully exploited to communicate data in sound (7). The concept of graphs in sound, obtained by mapping numerical values to the pitch of sound was introduced by Mansur (8) Following this concept, substantial research work has been conducted in recent years in the field of auditory graphs, mainly with line graphs (9, 10). Extending these data-sonification techniques to a more general data structure like a data table, we have developed TableVis. This audio-haptic interface was designed to browse tabular numerical data sets, with particular focus on providing the means to obtain overview information quickly and easily, while still supporting filtering of information and obtaining




Fig. 1: The data table to be explored is presented on the active area of the tablet, scaled to fill it completely

details on demand at later stages in the exploration.

Section 2 in this article summarizes the main characteristics of this interface and briefly outlines the studies conducted to evaluate the interface and the techniques that it employs, both quantitatively and qualitatively. The limitations of the methodology followed are explained, justifying the need for further analysis of the data collected during those evaluations, with particular attention to the exploratory strategies and procedures employed by the users. Section 3, offers an analysis and discussion of the exploratory strategies and procedures. Observed and recorded data from the explorations performed by blind and visually impaired users are analyzed as a method of learning about the best ways of using the interface, and to identify any common issues that might interfere with the exploratory process. Ultimately, this process of analysis could also help to understand better the mental processes of users with visual impairments, in their attempts to complete data exploration tasks.

PREVIOUS WORK WITH TABLEVIS

A traditional iterative design process was followed to design TableVis. Following initial extensive user-centred requirements capture conducted at various British institutions for the blind (like the RNC in Hereford and the RNIB in Edinburgh), an initial prototype was designed. Then, several iterations were followed in a process in which the prototype (as a whole or specific aspects of it) was evaluated quantitatively and/or qualitatively, leading to the refinement of the requirements and to improvements in the design. In this section, TableVis is described mentioning only the functionality that is relevant to the discussions that follow later in this article. For a detailed list of requirements

captured and a complete relation of functionalities, the reader can refer to previously published work by the authors (11,12).

TableVis is a multimodal interface designed to explore tabular numerical data by generating sonifications of the data interactively, with particular focus on obtaining overview information. Two-dimensional data sets of any size or shape are scaled to fit the whole active area of a WACOM graphics tablet (www.wacom.com), as seen in Figure 1. With the tablet's electronic pen, the user can point at any position on the tablet, which is equivalent to pointing at the cell that is mapped to that position on the tablet. This action produces the whole row or column being pointed at to be sonified. Whether it is the row or the column that gets sonified will depend on the sonification mode that is selected at that time (these modes will be referred to as rows mode and columns mode). Sonifying a complete row or column involves sonifying each and every one of the cells in that row or column by mapping the values contained in them to pitch of sound (with higher values corresponding to higher pitches), and playing them from left to right (rows) or from top to bottom (columns), in very rapid succession. The succession is so fast that sounds are perceived to be simultaneous, as if they were played in a chord rather than in an arpeggio. The reason why sounds are played in an arpeggio rather than in a strictly simultaneous chord is that the user can modify the duration of the sound for each cell, stretching the sonification of a row or column in time, which functions as a detail filtering device, as described in earlier publication by the authors (12). For the purpose of the discussion presented in this article, however, it is sufficient to consider that all the sounds are perceived to be simultaneous. Presenting information in this fashion hides

detail about the structure of the sounds in each row or column, and facilitates making simple comparisons between rows or columns regarding the average pitch of all the sounds contained in them. Thus, a user can tell which of two rows or columns contains sounds that, on average, have higher or lower pitch, or to identify outliers. Comparisons among all the rows or columns reveal features about the whole of the data set, such as trends and patterns along and across the data set. Exploring the same data table in rows mode and in columns mode presents two complementary representations of the same data object (these representations will be referred to as views of the data set, following a visual analogy). Each one of these views contains the same information (the same cells have rendered the same sounds), but grouped in two different ways that complement each other (as the multiple views in the orthographic projection of a three-dimensional object complement each other, to use an example from technical drawing). The interface does not present any information on the visual screen.

Several aspects of the design of TableVis are crucial for the users to be able to conduct non-visual explorations successfully. An absolute positioning pointing device (like a Wacom tablet in the case of TableVis) is used for the users to be able to know about the position of the pointer in the data set. Additionally, the following invariants must be maintained to be able to keep both focus and context information throughout the exploration:

- The complete data set is on display. The user knows that the complete data set completely fills the active area of the tablet and that the physical, tangible borders of this active area correspond to the boundaries of the data set as well.
- Information remains stationary in space. The user knows that each cell in the table always occupies the same physical location on the tablet.
- The active area on the tablet has a fixed size. The user knows the dimensions of the active area of the tablet.

These three invariants are used in combination with the user's senses of proprioception and kinesthesia (the senses of the position and movement of the limbs) to maintain contextual information throughout the exploration and to be able to point directly to different areas on the data set without having to access information sequentially. Because the complete data set is presented on the tablet, there is no need to scroll information; and because the location of the cells remains static, once some area of interest has been identified on the data set, the user is confident that the data will always be on the same physical position. The tablet has a physical, tangible frame around its active area, which provides easily felt reference about the constant size and location of these boundaries (see Figure 2). Feeling these fixed, physical frames, combined with the user's senses of proprioception and kinesthesia, facilitate jumping directly to



Fig. 2: A user explores a data table generating data sonifications interactively with the pen, while the left hand feels the boundaries of the data set to provide contextual information

those areas of interest that have been identified earlier in the data browsing process. Furthermore, because the data set fills up the active area on the tablet completely, the user can also jump directly to specific areas on the data set (the top, bottom left or right edges of the data set, the centre, explore the four corners, etc.). Finally, as users traverse data sets obtaining information through sound, they keep track of the position of the pen (of the data under focus) in relation to the physical boundaries of the tablet (that is, in relation to the boundaries of the data set), maintaining constantly contextual information through the uses of sensory modalities other than hearing, which is being used to analyze data. Maintaining focus and context information is very important to support the functional components of browsing, as described by Kwasnik: orientation, place marking, identification, resolution of anomalies, comparison transitions. (13).

In addition to several iterations of qualitative evaluations of the interface, two quantitative evaluations have been conducted, both with groups of sighted, blindfolded participants and with groups of participants who were blind or visually impaired. The first study showed that obtaining overview information from data tables by exploring them using the sonification technique described above was faster than using a speech synthesiser based interface, whereas the accuracy of the information obtained and the user's subjective mental workload were similar in both cases (11, 14). The second study showed that using TableVis to obtain overview information from data tables is highly insensitive to the size of those tables, at least for the range of table sizes that was tested (sizes from $7 \times 4 = 28$ to $31 \times 24 = 766$ cells) (12).

Although both qualitative and quantitative evaluations showed that the interface was usable and the browsing techniques implemented in it provided efficient access to overview information of data tables, the participants took different approaches to solve the exploration tasks. These differences were detected both from interviews that were conducted with each user after their evaluation of the

system, and in the comprehensive data recorded from their interaction with the system. TableVis offers room for creativity in the use of its functionality, which is not tied to a single possible way of use. This is due to the fact that, while certain theoretically optimum exploratory practises were supported and expected to be used, the inherently interactive character of the interface was preserved as an essential tool to provide freedom for exploration and to let users build their own knowledge of the data. As a consequence, participants in the evaluations of the interface did not always use it exactly in the way intended by the designers, yet they created ways of using the interface that led to successful and efficient data explorations. This richness of ways of use was considered to be a valuable source of information to understand better the various possible approaches to the problem of browsing tabular data that different users may have, and in particular users with visual impairments. These observations motivated the analysis of exploratory strategies and procedures that is presented in the next section.

EXPLORATORY STRATEGIES AND PROCEDURES IN TABLEVIS

Exploring tabular data sets non-visually is a complex task in which understanding the mathematical relations intrinsic to a tabular structure is used to provide meaning to the data contained in the table, in the context of the metadata (information describing the data and giving it a meaning) of the particular table being explored. Similar to the way in which vision or touch are used to explore a physical object, hearing is used in TableVis to explore a tabular data set actively, obtaining different views of the data object via data sonification, and combining them to infer properties of that object. Exploratory strategies and procedures are planned and implemented to carry out those active explorations. An exploratory strategy is the action plan that the user devises to carry out the exploration of a data set in search for information. Exploratory procedures are the ways in which that strategy or action plan is implemented; the ways in which actual interaction with the system takes place, including repeatable subroutines that contribute towards building an understanding of the data set.

As stated in the previous section, the complexity of tabular data exploration tasks and certain degree of flexibility in the interaction with the interface give way to the users' inventiveness to devise their own ways of approaching the problem. Observing these approaches can provide substantial information about which practises should be supported, and it can throw some light on the cognitive processes that take place during the exploration of data tables.

This analysis was performed on data of two kinds: semi-structured interviews conducted with every participant after the evaluation, and rich information about the actual interaction with the interface, that was logged automatically for every participant and data exploration. This information

covered every event in the interaction with the interface, including the complete succession of all the actual positions of the pen on the tablet, time stamped. This collection of data permitted reconstructing and analysing aspects like speed of exploration and trajectories, which could be played back at actual speed, faster or more slowly, for better observation of the events.

During the evaluations from which the data analyzed here were recorded, the participants had to interact with the system and obtain overview information about numerical data tables. All the data tables had 7 columns (days of the week) and 24 rows (hours of the day), and the data contained in each cell was the hourly number of visitors that a hypothetical website received at a particular day and time. Thus, a complete table contained information about hourly number of visitors to that website, in a whole week. Two types of questions could be asked about a data table:

- Find which day of the week (or which time of the day) the website is busiest (or least busy).
- Determine which quadrant of the table contains a cluster of cells with higher or lower values than in the rest of the table

Data from eight participants is considered for the analysis presented here. All eight participants were visually impaired (congenitally or adventitiously impaired, with total or partial blindness). Of the 12 participants who volunteered for the study from which the data originated, only 8 completed it. Three of the four participants who did not complete the study did not develop an understanding of the concepts of tabular data structure and the mathematical relations between rows and columns involved in them, nor did they understand the data sonification metaphors in the rows and columns modes well enough to be able to complete the session in one hour. They were mature participants (between 35 and 60 years old) who had not studied or used mathematical concepts involving tabular structures in recent years. The fourth participant who could not complete the study suffered additionally from a hearing impairment that was severe enough not to let him/her discriminate sound pitches. Each participant received training on how to use the interface, the sonification metaphors, and the functionality available to browse data in tables. Then, each participant was asked to explore 12 different data sets and to obtain overview information to answer to questions of the types i and ii, as shown above (six data sets for each type of question). The order in which these types of questions were asked was counterbalanced. At the end of the evaluation, all the participants who had volunteered were paid for their participation in the study.

Stages of exploration

In the observation of the bigger structure of the exploration process of a single table performed by the participants, three different stages could be distinguished. The first stage

required mental processing of information by the user, and no action took place. The second stage included all the active interaction with the system, where information from the data set was retrieved interactively. In the last stage, mainly or exclusively mental information processing took place again. In the context of the tasks performed by the participants, and in light of the additional information collected from some of the participants through interview and think aloud evaluation techniques, it is proposed that these three stages can be rationalised as presented in the list below.

- Interpret the question in terms of tabular data structure, and devise strategy for exploration. The first step requires the user to understand the question and select the best strategy to explore the data to find an answer to that question. First, a connection between the question and the metadata of the data set has to be established to understand what actions are required to be performed next. For example, if the metadata already known by the user establishes that the columns of the table are days of the week and the rows are hours of the day, then a question that starts with “at which time of the day...” will be asking about “which row...”, and will require a precise row as the answer. Questions about days of the week or about times of the day required users to explore data tables in rows or in columns modes only, depending on whether it is a row or a column that is being searched for. Questions about quadrants in the table, instead, required users to scan the data tables both by rows and by columns, and to combine the information obtained from both views to infer the answer to the question. The translation of the question in terms of the meaning of the data into the strategy to explore the data table requires a good understanding of the distribution of data in a table and the relationships between rows and columns. Some of the users that volunteered to take part in the evaluations of TableVis did not complete the evaluation because they could not complete this first stage because they did not have a good enough understanding of the mathematical relations between the cells in a table. This is the stage that arguably makes the highest mental processing demands on the user.
- Scan the data in one or two directions. This stage includes practically all the interaction with the interface. Interactively, all the information is retrieved in sound. The user performs the actions planned in the previous stage, repeating and refining searches until all the information required in terms of tabular data structure and features in the data has been collected.
- Combine all the information retrieved through interaction with the system, and obtain the answer to the question. Having obtained general information in one or two navigation modalities (depending on the

task), the user still needs to obtain the answer for the question in terms of the meta-data. If the answer is a particular row or column, then the user still has to interact with the system to request details in speech at the position that had been identified to be containing the correct features in the data. On the contrary, if the task was to identify a region in the table that required exploring the data set in both navigation modes, no further interaction with the interface is required, and the user has to perform the mental exercise of inferring the answer by superimposing both complementary views of the data (obtained in different directions of exploration) and identify a particular area in the table as containing features in the data that are in accordance with the information in each separate view.

Although these three stages are observed to happen in the sequence in which they have been presented, that a certain degree of overlapping and iteration of the whole sequence is present could be expected, generally speaking.

Exploratory strategies and procedures

Having described the three stages of exploration, this section analyzes the data observed and recorded during the evaluations of the interface that corresponds to the second of those stages, i.e. during the active exploration of the data through interaction with the interface. Figures 3 to 10 show the traces of the pen on the tablet, for each one of the eight participants that completed all twelve data explorations. Each Figure in the table includes the traces of all 12 explorations of a participant. These traces were also analyzed, reproducing them as an animation and observing the exploration of each data set separately. Light shade traces (small horizontal lines) were generated in rows mode. Dark shade traces (small vertical lines) were generated in columns mode. Marks closer to each other in a trace line




Fig. 3: Traces of participant (a). Light shade traces (small horizontal lines) were generated in rows mode. Dark shade traces (small vertical lines) were generated in columns mode. Circles indicate requests of details in speech (label of current row or column, depending on navigation mode, represented by the shade of the circle)




Fig. 4: Traces of participant (b). Light shade traces (small horizontal lines) were generated in rows mode. Dark shade traces (small vertical lines) were generated in columns mode. Circles indicate requests of details in speech (label of current row or column, depending on navigation mode, represented by the shade of the circle)




Fig. 7: Traces of participant (e). Light shade traces (small horizontal lines) were generated in rows mode. Dark shade traces (small vertical lines) were generated in columns mode. Circles indicate requests of details in speech (label of current row or column, depending on navigation mode, represented by the shade of the circle)




Fig. 5: Traces of participant (c). Light shade traces (small horizontal lines) were generated in rows mode. Dark shade traces (small vertical lines) were generated in columns mode. Circles indicate requests of details in speech (label of current row or column, depending on navigation mode, represented by the shade of the circle)




Fig. 8: Traces of participant (f). Light shade traces (small horizontal lines) were generated in rows mode. Dark shade traces (small vertical lines) were generated in columns mode. Circles indicate requests of details in speech (label of current row or column, depending on navigation mode, represented by the shade of the circle)




Fig. 6: Traces of participant (d). Light shade traces (small horizontal lines) were generated in rows mode. Dark shade traces (small vertical lines) were generated in columns mode. Circles indicate requests of details in speech (label of current row or column, depending on navigation mode, represented by the shade of the circle)




Fig. 9: Traces of participant (g). Light shade traces (small horizontal lines) were generated in rows mode. Dark shade traces (small vertical lines) were generated in columns mode. Circles indicate request of details in speech (label of current row or column, depending on navigation mode, represented by the shade of the circle)




Fig.10: Traces of participant (h). Light shade traces (small horizontal lines) were generated in rows mode. Dark shade traces (small vertical lines) were generated in columns mode. Circles indicate request of details in speech (label of current row or column, depending on navigation mode, represented by the shade of the circle)

indicate that the pen was being moved more slowly on the tablet. Circles indicate requests of details in speech (label of current row or column, depending on navigation mode, represented by the shade of the circle). Some demographic information about the participants: all participants were in the range of 18 to 25 years old, except for participant d (Figure 6), who was in the range 46 to 55 years old. Participants a, c, e and g (Figures 3, 5, 7, 9) were totally blind, blindness being congenital for participants a and e (Figures 3, 7). The other four participants were partially sighted, but required accessibility aids to use computers. Participants b and d (Figures 4, 6) used screen readers, while participants f and h (Figures 8, 10) used screen magnifiers and high-contrast colour schemes.

The differences in the approaches taken by the participants are analyzed in the following list:

Use of speed of scan. The speed at which the pen is moved on the tablet, determines the number of data sonifications per unit of time that are generated. Thus, faster speeds present more information per unit of time and are more appropriate to obtain very general information (like trends, patterns, or location of outliers), but worse to extract richer information about each one of the sonified rows or columns. The opposite is true for slower speeds. Except for participant b (Figure 4), all the other participants developed a procedure that involved scanning the table by moving the pen at an approximately constant speed, to get an initial overview of the data. Although the speed was constant, it varied substantially between participants. Participant a (Figure 3) made use, in average, of some of the fastest hand movement speeds, whereas participant f (Figure 8) consistently used a slow, uniform speed. Participant e made use of the broadest range of uniform speeds between scans, making some of the traverses at fast speed and some others much more slowly. Participant b (Figure 4) showed a different approach to the use of speed. This participant

scanned the tables slowly, stopping for a short time after hearing each one of the sonifications to analyze it mentally. In Figure 4, these speed variations can be observed in the frequent variations in the density of the trace-plots.

Number of traverses. In most cases, traversing the table only once was enough for all the participants to know the approximate location of the data with interesting features. After the first traverse, they normally went back directly and quickly to the area in which they had identified interested data, simply to confirm the finding. If two or more interesting locations had been identified, they were quickly compared, trying to ignore all the rest of the sonifications that still were being generated as the pen moved. Most participants performed approximately the same number of traverses in all the tables they explored. Participant b (Figure 4) rarely required more than one, while participant h (Figure 10) rarely performed less than three full traverses. All the other participants showed the behaviour described above, i.e. perform a full scan and then confirm the findings locally. As an exception, participant e (Figure 7) showed great variability in the number of full traverses performed in each navigation mode, ranging from one to seven. It was also observed that, in general, scans performed at faster speeds of pen movement required more full traverses of the data set, but the whole process did not necessarily take longer time to complete.

Direction of scan. Both in rows and in columns navigation modes, most scans started near the top-left corner of the tables, for the majority of the participants. Clear examples of this approach are participants d and h (Figures 6 and 10), but most of the other participants show this tendency too. An explanation for this choice may lie in the meaning or meta-data of the particular sets being explored. In them, the top-left corner is the origin for both axes. This would suggest that context is maintained during the exploration also regarding the meaning of the data being explored, what would make the final stage of translating the information gathered into a meaningful answer easier. Further investigation is required to confirm this hypothesis. In some other cases where explorations required both rows and columns views, some participants scanned the data very efficiently by drawing lines that followed the areas near the borders, starting a new traverse where the previous one ended, thus not having to lift the pen and relocate it somewhere else to start a new scan. Participants a and c (Figures 3, 5) are good examples of this.

Use of tangible, physical borders. By design, the rulers are intended to make the working area tangible to support obtaining context information, and in that sense, it is necessary to use them to explore data sets successfully, although this is normally not done consciously. By observing the way in which the tables were scanned and how participants could go directly back to areas of interest, it can be inferred that the rulers were used for this purpose

by all the participants. Some participants also chose to use the rulers as guides to traverse the tables horizontally and vertically. Participant a (Figure 3) made heavy use of the top and left physical borders, making virtually every exploration in this way. Participant c (Figure 5) also used the tangible borders as guides, mainly the left and bottom border, but also the top border.

No other participants used the physical borders as guides, except very occasionally. Two considerations can be made here. The first one is that it could be expected that participants with total blindness, particularly those who have a congenital impairment, might find more benefit in using the rulers as guides to traverse the surface of the tablet. In fact, both participants a and c were totally blind. Furthermore, participant c, who used the borders as guides more consistently, was congenitally blind. However, participants e and g (Figures 7 and 8, adventitiously totally blind and congenitally totally blind, respectively), show that explorations could be performed without any additional support from external physical guides. The second consideration refers again to participant a (Figure 3), who made use of the highest speeds for scanning the data, and also used the tangible frame as guide for the pen most consistently, suggesting the possibility that using the frame in this way can be a technique that may help obtain higher efficiency during exploration.

Shape of lines. Despite the pen not needing to be dragged in straight lines to scan a table as long as it travels between left-right and top-bottom, straight lines were the most common shape drawn on the tablet during the data explorations. In addition to the two participants that used the physical borders of the tablet as guides (participants a and c, Figures 3, 5), the rest of the participants, with the exception of e (Figure 7), drew lines that were not only remarkably straight but also very close to being horizontal and vertical. Partially sighted participants f and h (Figures 8, 10) are particularly good examples, as it is participant g among the blind participants who did not use the physical borders as guides. Although it could not have been expected that a user kept the pen within a single row or column confidently without some additional help (and indeed there is no need for it to obtain an overview in TableVis), the traces presented here show good discrimination between horizontal and vertical, and the ability to stay approximately in a straight line, for participants with a variety of visual impairments. Participant e (Figure 7) used a particularly free style of exploration, in which very fast scans were nervously drawn, vaguely following horizontal and vertical lines.

Use of details on demand. With very few exceptions, this functionality was used as it had been intended, that is, only at the end of the exploration and only if detailed information was required in the answer. A good use of this functionality was also made when a participant detected a position during the initial scan that could well be the answer

that was being searched. Then, the participant would request details in speech (the text of the label of that row or column) before continuing traversing the table. If at the end of the traversal no other part of the table had been found that could be a plausible candidate for the answer, the answer could already be given without having to go back to that position to get detailed information, because that information had already been collected when it was first identified, and the user could still remember it. The best example of this strategy was followed by participant b (Figure 4), as part of the strategy that this participant followed.

Some common issues and possible solutions

The observation of the exploratory procedures followed by the participants revealed also some common situations that appeared to be confusing or at least in need for further support from the interface.

Modality confusion. Probably the most common source of temporary confusion emerged when a participant thought that one navigation mode was active while it was in fact the other one that was active, despite the confirmation of the selection received through synthesised speech when the modality was selected. In those cases (for instance, when a participant tried to explore columns with the rows mode active), no sounds emerged from the scanning of the data, except occasionally and erratically (when the pen crossed “by accident” the border between two rows). The severity of this problem is minor because not obtaining the expected response from the interaction alerted the participants about the problem, which could be rectified immediately. Interestingly, however, it was common for some participants to finish a horizontal scan and start with a vertical one forgetting to change the selection on the navigation mode. Some property in the sonification that was different between both modes could provide constant confirmation about the currently selected modality. Alternatively, having to be engaged actively in the selection of one of the modes could avoid these brief episodes of confusion. Further investigation must be conducted about the best ways to solve this problem.

Comparison of non-adjacent rows of columns. A different kind of issue was observed in situations in which a participant had identified two or more possible candidate rows or columns for the answer, which were not adjacent to each other, and comparisons had to be performed. Some participants used the fingers in one hand to mark the locations of those points of interest and then tried to jump directly between fingers to compare only the sounds that had been pre-selected. Because those positions were not adjacent on the table, moving the pen between them would also produce sonifications of the rows or columns that were in between. This was difficult to avoid even by lifting the pen at one position and putting it down at the next, because the technology used in the wireless graphics tablets (like the

WACOM tablet used in TableVis) allows the pen to communicate with the tablet at several millimeters of distance, before they get in contact. For this reason, performing “clean” jumps is difficult. Even if this behavior of the physical interface was not an issue (pressure sensitivity could potentially be used to make sure that the pen is in contact with the tablet as a requisite to produce sound), the accuracy of hitting specific target locations using proprioception but not having visual feedback is limited, and when trying to jump from one location to another the pen would often land off-target, obtaining undesired sonifications also in this case, which would interfere with the performance of distant comparisons. Additionally, marking several points on the tablet with the fingers from one hand is inaccurate and often the reference got lost by having moved the hand accidentally. This situation would benefit from some form of external memory aid that could be used to mark interesting positions that could be easily found again, and which also permitted filtering out the information that had not been selected to perform comparisons between the selected positions without interferences from other sonifications that got triggered in between. The authors are currently investigating the use of EMA-Tactons, external memory aid Tactons or structured vibrotactile messages (15).

CONCLUSIONS

Exploring numerical data tables non-visually to obtain overview information that can answer high-level questions about the data is a complex task. TableVis provides support to complete tasks of this kind through interactive sonification of the data. Analyzing the way in which users with visual impairments approach this problem and make use of the functionality available in TableVis provides valuable insight into the practices that should be supported and the issues that must be tackled. This approach also provides a better understanding of the problem of data exploration from the perspective of users with visual impairments.

This paper has focused on observing and examining the exploratory strategies and procedures used by participants with visual impairments who evaluated the interface. A three-stage approach to completing the task has been described. Whereas the general exploratory strategy chosen by all the participants was very similar, the way in which the actual data were gathered varied between participants. A number of exploratory procedures have been described, and further investigation about them is suggested to obtain evidence about which practices are the best and which cognitive aspects of the exploration should be supported better.

Issues regarding the discrimination of navigation modes and the need for external memory aids to select and filter information have been identified as well. Future research will address these questions.

ACKNOWLEDGMENTS

We want to acknowledge the institutions for the blind and visually impaired that have supported this research, especially the RNC in Hereford for their committed participation and support. We also want to acknowledge the EPSRC for funding the MultiVis project (grant GR/S86150/01), part of which is the work presented here (www.multivis.org).

REFERENCES

1. Shneiderman B. The eyes have it: a task by data type taxonomy for information visualizations. In: IEEE Symposium on visual languages, 3-6 Sept 1996. Boulder, CO, USA: IEEE Comput. Soc Press, 1996; 336-43.
2. Zhao H, Plaisant C, Shneiderman B, Duraiswami R. Sonification of geo-referenced data for auditory information seeking: Design principle and pilot study. In: Int Conf Audit Display (ICAD). Sydney, AU, 2004.
3. Pérez-Quñones MA, Capra RG, Shao Z. The ears have it: A task by information structure taxonomy for voice access to web pages. In: Rauterberg GWM, Menozzi M, Wesson J, ed. Human-Computer Interaction-INTERACT'03. Zürich, CH: IOS Press, 2003: 856-9.
4. Saue S. A model for interaction in exploratory sonification displays. In: Cook PR, ed. Int Conf Audit Display (ICAD). Atlanta, GA, USA, 2000.
5. Bertin J. Semiology of graphics: Diagrams, networks, maps. Madison, WI, USA: Univ Wisconsin Press, 1967/1983.
6. Miller GA. The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychol Rev* 1956;63:81-97.
7. Kramer G. An introduction to auditory display. In: Kramer G, ed. Auditory display: Sonification, audification and auditory interfaces. New York, NY, USA: : Addison-Wesley, 1994:1-77.
8. Mansur DL. Graphs in sound: A numerical data analysis method for the blind. Unpublished dissertation. Los Angeles, CA, USA: Univ Calif, 1985.
9. Walker BN, Cothran JT. Sonification sandbox: A graphical toolkit for auditory graphs. In: Int Conf Audit Display (ICAD). Boston, MA, USA, 2003.
10. Brown LM, Brewster SA, Ramloll R, Burton M, Riedel B. Design guidelines for audio presentation of graphs and tables. In: Proc ICAD 2003 Workshop Audit Displays Assist Technol. Boston, MA, USA: 2003:284-7.
11. Kildal J, Brewster SA. Non-visual overviews of complex data sets. In: Extended Abstracts CHI 2006. Montreal, Canada: ACM Press, 2006.
12. Kildal J, Brewster SA. Providing a size-independent overview of non-visual tables. Glasgow, Scotland, UK, 2006.
13. Kwasnik BH. A descriptive study of the functional components of browsing. In: Proc IFIP TC2/WG2.7

- Work Conf Engineering Human-Comput Interaction. North-Holland, 1992:191-203.
14. Kildal J, Brewster SA. Explore the matrix: Browsing numerical data tables using sound. In: ICAD2006. London, UK, 2006.
 15. Kildal J, Brewster SA. VibroTactile external memory aids in non-visual browsing of tabular data. In: McGookin DK, Brewster SA, eds. First Int Workshop Haptic Audio Interaction Design (HAID). Glasgow, Scotland, UK: 2006:40-3.