

PROVIDING A SIZE-INDEPENDENT OVERVIEW OF NON-VISUAL TABLES

Johan Kildal, Stephen A. Brewster

Glasgow Interactive Systems Group
Department of Computing Science
University of Glasgow, Glasgow, G12 8QQ, UK
{johank, stephen}@dcs.gla.ac.uk www.multivis.org

ABSTRACT

Obtaining an overview is an important first step in the analysis of data sets, which cannot be easily done non-visually with current accessibility tools. We present TableVis, a multimodal interface to obtain overview information from numerical data tables non-visually, with the use of an interactive sonification technique controlled from a tangible physical device (a tablet). An experimental study with TableVis is reported, in which it is found that performance is highly insensitive to the size of the data set being explored, for a broad range of data set sizes.

1. INTRODUCTION

Much of the data we receive daily are numerical and they are often arranged in tables. Tables are a very general and foundational structure in which data can be organised. Card et al. [1] say that any set of raw data can be transformed into tabular mathematical relations as a prior stage to a more detailed analysis of the data. Hence tables can be seen as a midway point between raw data and a purpose-built data representation (such as a bar chart or a scatter plot).

A natural way of examining a data table (or any other data structure) that is encountered for the first time always starts with a browsing stage to obtain overview information about it and to be able to judge on the relevance of further analysis, as well as where or how to carry it out. The literature suggests many ways of dividing the task of analyzing a data set into different stages, all of them coinciding in that the first stage has to be one in which a general overview is obtained. Shneiderman [2] has put forward his Visual Information Seeking mantra: ‘Overview first, zoom and filter, then details-on-demand’, where the aim of the first stage is to gain an overview of the entire collection. Based on the same division of stages, Zhao et al. [3] proposed an Auditory Information-Seeking Principle (AISP): ‘gist, navigate, filter, and details-on-demand’, where gist stands for a ‘quick grasp of the overall data trends and patterns from a short auditory message’, and suggesting that if information-seeking in the auditory mode follows the same pattern as in the visual, then the collaboration between users holding data in the visual and in the auditory domains might become easier. Pérez-Quinones et al. [4], also after Schneiderman’s division of stages, propose the following set of information seeking tasks for browsing web pages using speech: Situate, Navigate, Query, Details on demand, where the aim is to help the user identify his/her location within the information space. Other examples include Sauer [5], who makes a two-stage distinction, called ‘orientation’ and ‘analysis’, or Bertin [6], who defines general graph-reading operations, characterised by their level or

scope, as ‘overall’ (looking for overall structure), ‘intermediate’ (identifying trends), and ‘elementary’ (looking for specific, perhaps quantitative, information).

While visually inspecting a data set is quite an efficient way of obtaining overview information quickly, particularly after having constructed appropriate graphical visualisations of the data (which exploit the fact that vision is a highly synoptic sense and easily provides a general view of the whole of a subject), blind and visually impaired (VI) users have far greater difficulties in accomplishing this basic task. Current accessibility tools only provide access to fully-detailed, unitary portions of information sequentially, and do not support mechanisms for summarising information or highlighting its most salient features. Making comparisons between units of data is the most frequent of the functional components of data browsing, which includes also orientation, place marking, identification, transition and resolution of anomalies (Kwasnik [7]). Accessing fully detailed information sequentially imposes great demands in terms of mental workload: the user has to remember all the detailed information, extract meaning from it, compare it against every other unit of information, judge the actual scale of the difference, and try to construct a mental model of the whole data set. All this process leads very quickly to saturation of the user’s working memory (only a few items can be kept in the user’s working memory at a time [8]), as well as being a slow procedure. For example, accessing information in a large spreadsheet is normally done using a screen reader that speaks all the information in the document while the user navigates sequentially across the grid. As Stockman *et al.* describe it [9], obtaining an overview of a large data table in this way is a long and difficult process, and is more difficult the larger the data set is, implying a direct relationship between the size of the data set and the time and effort needed to complete the task. Reading Braille versions of spreadsheet tables presents similar problems and limitations.

In the following sections TableVis is presented, a new interface to explore numerical data tables and extract general overview information from them quickly and easily. A description and brief justification of the design of this interface is given first. Then, the results from a quantitative experiment are reported, in which the effect of the size of the data set over exploration time, proportion of incorrect answers and mental workload using TableVis are investigated. Results from this experiment show that, within a broad range of table sizes, these three parameters are largely independent from the size of the table explored using TableVis.

2. TABLEVIS

In this section, the process of design and implementation of the interface is described and briefly justified. A complete description of the current state and functionality of TableVis is also provided.

2.1. Requirements Capture

Extensive requirements capture was conducted in several British Institutions for the blind, (such as The Royal National College for the Blind in Hereford and The Royal Blind School in Edinburgh), to find out about current accessibility tools and techniques being used by blind and VI users to access information in general and numerical data sets in particular. This process led to the confirmation that obtaining an overview of any collection of information or document non-visually is difficult, particularly when trying to explore numerical data sets with more than only a few values.

In terms of design requirements, the system to be designed would have to avoid overloading short term memory, facilitate comparisons between subsets of data, support a focus and context paradigm, provide support for collaboration between sighted and VI users working on the same data sets, and offer a high degree of control in the interaction. An additional requirement was that the interface should not require expensive or difficult-to-get components, but rather off-the-shelf equipment that could be made to function together easily. From these requirements, an interactive table browsing system, called TableVis, was designed.

Most of these requirements were implemented in a multimodal interface in which data sonifications are interactively triggered while context information is gathered through the sense of touch.

2.2. Interface Design

2.2.1. Sonification Technique

Kildal & Brewster's table sonification technique [10] was taken as a starting point for the design of TableVis. With this technique, instead of sonifying each cell in a two-dimensional table, a whole row or column is sonified as a single piece of complex information. Thus, a table can be explored by rows (*rows* mode) or by columns (*columns* mode), and two simplified and complementary views of the same data table are obtained (an array of rows and an array of columns), each one of them containing the complete overview information. The total number of units of information a user can access is very much reduced in number (from mxn , where m and n are the number of rows and columns, to only m or n , depending on the direction of navigation), thus becoming a lot more manageable for the user's working memory to cover the complete data set. Additionally, since each one of the data units (a row or a column) is more complex, only the main characteristics of that large unit of information are picked up by the user, with the finer detail (which is irrelevant as overview information) filtered out by the user's auditory perception.

To construct the sonification of each row or column, all the values in it are mapped to pitch, distributing the complete range of values in the table over a range of 66 MIDI values

(piano sound), where higher numeric values are mapped to higher pitches. This is the same mapping strategy that Mansur used to sonify graphs [11], incorporating guidelines from Brown *et al.* [12]. These sounds are played in rapid succession (left to right or top to bottom, depending on the case). Kildal & Brewster [10] reported that an overview of the data in a table is obtained faster using this sonification technique than using speech-based screen reading software, without a reduction in accuracy. Stockman *et al.* [9] arrived at a similar conclusion in an implementation of data sonification functionality for one of the most widely used screen readers.

During the piloting stage, it was observed that when sonifying a row or column in this fashion, the shorter the duration of each sound (i.e. the faster the sonification), the less detailed information was perceived by a user. In the most extreme case, when sounds were played so close together that they were perceived to be simultaneous, like a musical chord, the overall pitch perceived was a blend of all the pitches in that dissonant chord. In those cases, although little could be told about the precise values forming these chords, carrying out comparisons between different chords and judging relative average pitches was easy, as was spotting outliers. In fact, studies like the one reported later in this paper show that the perceived overall pitch for each one of these chords is a good indication of the relative statistical mean of all the values in the chord. As the duration of each sound was increased (slower sonification), more detail about the structure of each row or column was revealed, but comparing rows and columns became more difficult.

It has to be pointed out that TableVis does not perform any computations with the data in the table, like averaging or selecting extreme values only. The information is presented complete and "as is", rendered according to the mapping convention described above. It is the user's perception of the resulting sonification that does the rest.

Speech is available on-demand to obtain text information about a particular cell, row or column being sonified. The user can also obtain in speech a complete description of the metadata of the table, i.e. description of the variables constituting each of the axes.

2.2.2. Physical Input Device

A WACOM graphics tablet (www.wacom.com) was chosen to control TableVis. This is an input device consisting of a physical flat rectangular surface on which the user can control the position of the mouse pointer with a wireless electronic pen. An ordinary mouse could not have been used because it is a relative positioning pointing device, and it requires visual feedback about the position of the pointer on the screen, making it not eligible for a non-visual interface. A tablet, however, can be configured to operate as an absolute positioning pointing device. In that mode, the position being pointed at on the tablet (using the tablet's pen) is an absolute reference on the 2D active area, meaning that by pointing at the same physical position on the active area of the tablet, the computer will always receive the same pair of coordinates for the mouse pointer. The tablet was augmented by adding a physical, tangible frame around the active area, delimiting it and making the borders easily identifiable using both hands. A small tangible dot was stuck in the middle of each side of the frame, to indicate the midpoint of each one of the sides for additional orientation and navigation cues (see Figure 1).




Figure 1. Graphics tablet with physical frame delimiting the active area, and a rotary knob with push-button function on its side.

The data table to be explored is presented on the active area of the tablet, scaled to fill it completely, as represented in Figure 2. Each cell remains in the same position throughout the exploration.




Figure 2. The data table to be explored is presented on the active area of the tablet, scaled to fill it completely.

Users can access any region of the table directly, knowing that the tangible edges of the active area on the tablet correspond to the limits of the data table. Using their senses of proprioception (the sense of the position of the limbs) and kinesthesia (the sense of the movement of the limbs), users can tell how far apart one hand is from the other, and judge approximately where on the tablet (in the data table) the pen has been placed (Figure 3). They also allow the user to “jump” from one position on the tablet directly onto another, targeting data cells that are in approximately known locations. It can therefore be noted that the tablet is not only an input device, but also an output device conveying focus+context information, which is necessary for the successful use of the interface.

Details-on-demand, as described in the previous section, can be obtained in speech by clicking on any of the two positions of the push-button built on the tablet’s pen. It was decided that both positions should provide the same functionality because pilot testing showed that it was difficult to differentiate one position from the other for some users. When the rows mode is selected, the label corresponding to that row is read out in speech, while when the selected mode is the columns mode, the label of the currently selected column is spoken. In the cells mode, both row and column labels are spoken, followed by the numeric value in the

selected cell, thus being able to get full detailed information about coordinates and numeric value in a particular position, should that information be of interest at some point during the data exploration.



Figure 3. The data table to be explored is presented on the active area of the tablet, scaled to fill it completely.

An additional input device, a Griffin PowerMate (www.griffintechology.com) USB rotary knob with a push-button function, shown in detail in Figure 4 and also seen in Figure 1, was used to modify the sonification speed during exploration and to shift between exploration modes. Turning the knob to the left reduces sonification speed by increasing the duration of each sound in the succession of sounds. Turning it to the right increases the speed. The system implemented achieved an effective fastest speed of 7ms duration per sound. Shifting between exploratory modes is done by pressing on the knob, which acts as a push-button. These functions can be integrated in the built-in controls of some graphics tablet models, which contain customisable push-buttons and sliders on both side margins.

2.2.3. Interaction

TableVis offers three modes of exploration. These will be referred to as ‘cells’, ‘rows’ and ‘columns’ modes. In the cells mode, the cell being pointed at on the tablet is sonified, facilitating ‘freehand’ exploration of the data table, cell by cell. In the rows and columns modes, a complete row or column is sonified when the pen is pointing at it. These two last modes facilitate obtaining quick overviews of the complete data set. Drawing a line across the whole table (horizontally in columns mode or vertically in rows mode), a sonification of all the data in the table is interactively generated, as the pointer moves over different rows or columns. The user can go back and forth, or jump from one end to another, retrieving complete ‘slices’ of data that can be easily compared against each other. If one row or column is judged to be interesting compared to the rest when heard as a chord (for instance because it sounds higher or lower or more dispersed), the user can choose to repeat that particular sonification by tapping on the same position on the tablet that generated that chord. Additionally, using the rotary knob (Figure 4), the user can adjust the speed of the sonification, stretching in time the succession of sounds in a particular row or column to obtain a sound graph similar to those generated for a line graph by Brown and Brewster [13], at the speed chosen by the user. This permits analysing selected parts of the information in more detail.




Figure 4. Rotary knob, used in TableVis to control the speed of the sonification

2.3. Interface Evaluation and Research Question

During several iterations of the design and implementation process, the interface was piloted with blind and VI users as well as with blindfolded sighted users. Once TableVis reached the state of development described above, a complete evaluation test was designed and conducted with both blind and VI users, and with sighted users. The design of that evaluation was based on exploring a collection of tables with a fixed size (24 rows and 7 columns) and extracting overview information from them, while time for exploration, percentage of correct answers and overall subjective workload were measured. Results were compared with those derived from using a speech-based system [14].

As results from that study were being collected, another parallel research question arose: How does the size of the data set affect the use of this interface? Is there a linear relationship between the size of the data sets and the performance parameters being measured? It could be expected that increasing the table size would produce an increase in the time needed for the exploration, a reduction in the percentage of correct answers and an increase in the overall subjective workload. Additionally, it could be expected that this effect would follow some kind of proportionality, where the more the table size was increased the bigger the variation would be in the three parameters measured. Given the characteristics of this interface, including the property that the sonification technique employed in it has to simplify the underlying data set, it could be hypothesised that the effect observed after increasing the table size would not be very strong. If this was found to be the case, then MultiVis could be a way of providing part of the sinopticity that vision offers when looking at visual representations of numerical data, where an overview of the whole is obtained easily with little effect from the complexity of the system or data set being explored. Blind and VI users could perform initial explorations of data sets to obtain overview information with comparable investments in time and effort, and with similar accuracy, for data sets of different sizes.

The rest of this paper will describe how this question was addressed and the results that were obtained.

3. EFFECT OF THE SIZE OF THE DATA SET IN THE USE OF TABLEVIS

An experiment was designed to find out about the influence of the size of the data set in the performance and user experience obtained when browsing them with TableVis to extract overview information.

3.1. The Data Sets

In order to cover a wide range of data sets in this experiment, three different table sizes were chosen: 28 cells (4 rows and 7 columns), 168 cells (24 rows and 7 columns) and 744 cells (24 rows and 31 columns). The table same table size and type as for the general evaluation of TableVIs was chosen as the medium size, and a larger and a smaller size were added to cover a broad enough range of sizes in both directions. The meta-data was common to all three table sizes, containing realistic information about the number of visits to a popular imaginary website, over a period of time. Table 1 summarises the table sizes used in the experiment. Table 2 shows the ratios in table sizes used in this study. A medium-sized table is 6 times the size of a small table, and a large table is 4.429 times the size of a medium. Therefore, the overall range of sizes considered in this experiment covers a range in which the size of the largest table is 26.57 times the size of the smallest. Figure 5 shows a visual representation of the relative sizes of the tables used in this study.

	Rows	Columns	Cells
Small	4 (part of the day: night, morning, afternoon or evening)	7 (days in a week)	28
Medium	24 (hours in a day)	7 (days in a week)	168
Large	24 (hours in a day)	31 (days in a month)	744

Table 1. Summary of the three table sizes used.

	Small	Medium	Large
Small	1	6	26.57
Medium	0.167	1	4.429
Large	0.038	0.226	1

Table 2. Table size ratios used in the experiment. Numbers in the table are calculated dividing the size (number of cells) of the entry at the top by the size of the entry on the left.

In order to create complexity in the data, the numerical values stored in each cell for any of the table sizes were large figures with 5 digits. Data were generated using the same algorithms for every table size. These algorithms created information with enough random noise to make it difficult to observe trends or patterns by merely inspecting a cluster of cells. However, some more general patterns were embedded in the whole of the data set. In some occasions, values in rows or columns would be higher or lower overall than the values in the neighbouring rows or columns. In some other cases, there would be a cluster of cells in which values were higher or lower than in the rest of the cells in the table. Each table contained several rows, columns or cell clusters that were

higher or lower than the rest in the table, although only one of them would be the absolute highest or lowest.




Figure 5. Visual representation of the relative sizes of the three tables selected for this study, created by super-imposing the small table (darkest shade) on the medium-sized table (medium shade), and both on the large table (lightest shade).

3.2. The Experiment

This section describes how the experiment was designed and conducted. The aim of the experiment is to find out whether the size of the data table in a data exploration task to obtain an overview has any effect on the time required to explore the data set, and on the accuracy of the information obtained as well as on the subjective overall mental workload (as a measure of the effort required from the user to complete the task). The null hypotheses being tested are that table size does not have any effect on these parameters. The only independent variable in the experiment is the size of the table. Three conditions are compared: *small* tables (4x7 cells), *medium* tables (24x7 cells) and *large* tables (24x31 cells).

3.2.1. Available functionality

All the TableVis functionality described earlier was available during the experiment, except the cells navigation mode, which was not used as this was a test of the overview functionality of our system:

- Rows and columns navigation modes, in which a complete row or column is sonified when a point on the tablet is selected. Every data exploration started, by default, in the rows mode. Shifting from one mode to another was done by pressing on the rotary knob. The speech synthesiser would announce the new mode by saying ‘rows’ or ‘columns’.
- Sonification speed control, by means of turning the rotary knob (left to decrease speed and right to increase it again). By default, every new data exploration would reset to the maximum speed.
- Details on demand, in speech, about the selected row or column, by clicking on any of the two positions of the built-in button on the tablet’s pen.
- Tapping on the same position on the tablet would repeat the same sonification.

3.2.2. Data Sets and Tasks

Three sets of twelve tables were constructed, one set per size of table. Participants in the study were asked to explore each one of the thirty-six tables using TableVis, and to answer a general question about the data which required obtaining an overview. Each table had one of two kinds of questions, with six tables for each type of question in each set:

1. Find which row / column has the overall highest / lowest values. Examples of this kind of question are: ‘Which day of the week / month gets more / less visits?’; ‘Which part / time of the day is normally the website busiest / least busy?’ As explained earlier (Table 1), weeks and months are the columns, depending on the table size. Similarly, part of the day and time of the day are the rows.
2. In which quadrant of the table are the highest / lowest values?

The first kind of question asks the participant to compare all rows or columns and identify the one with overall highest or lowest values in it. Participants were expected to select the appropriate navigation mode (rows or columns) and make comparisons until the correct answer was found. The overview information was obtained by observing variations of the data in one of the axes, and an exact answer was required. An answer was considered to be correct if it matched the row or column with the highest or lowest statistical mean, depending on the question in particular.

The second kind of question asked participants to identify in which of the four quadrants of a table were the cells with the highest or lowest values in the whole data set. The four quadrants in a table (top-left, top-right, bottom-left and bottom-right) are the four areas resulting from drawing a horizontal and a vertical line in the middle. Data sets with this kind of question contained peaks and valleys in the data, and the highest or lowest peak or valley was the one that needed to be identified. This was normally formed by a cluster of cells, and clearly located in one of the four quadrants, although there could be several candidates in different quadrants of the same table. The answer was judged to be correct if it corresponded to the quadrant containing the highest or lowest single value (depending on the question), which was always part of a cluster of cells with high or low values, in the whole data set. This kind of question required that the participant used both rows and columns exploration modes to obtain overview information about the evolution of the data in both axes. The accuracy required for the answer was lower in this case, as only the quadrant containing the distinctive cluster of values had to be identified.

3.2.3. Procedure

The experiment was run with eighteen sighted, blindfolded participants. Each session took about one hour to complete.

After an introduction which described the aims of the experiment, training about the use of all the functionality made available in TableVis for the experiment was provided. Participants explored some data tables without blindfolds, until all the concepts were well understood, including the meaning of the questions and the kinds of answers that were expected. Participants were informed that they would have up to two minutes to explore each table and give an answer to the question. Each question would be given at the time of presenting each new table. Participants were also informed

that they were not expected to consume all 120 seconds of exploration time, and that they should give the answer as soon as they knew it. The procedure of giving the answer was to say it loudly, to be noted by the experimenter. In the case of the first kind of question, this would mean saying the day of the week or day of the month in some cases, and saying the time of the day or the part of the day (night, morning, afternoon or evening) in others. Participants could obtain this specific information by requesting speech details-on-demand, as it has been explained earlier. When the participant said the answer, the experimenter would stop the timer (the computer recording the total time of exploration) and write down the answer on a data-collection form. Then, a new table would be loaded for exploration.

The order in which conditions (sets of tables with a particular size) were presented was counterbalanced to avoid learning effects in the results

At the end of each condition (of each set of tables for a particular table size), participants filled-in a NASA-TLX form [15] that provides feedback on the subjective workload experienced in that condition. The same form was used for all three conditions, so that participants placed three marks for each scale, providing ratings for each condition relative to each other. In the NASA-TLX forms, participants had to rate the following categories in a scale from 1 to 20: mental demand, physical demand, time pressure, effort expended, performance level achieved, frustration experienced. We added two additional factors: annoyance and overall preference to gain some further information.

Complete data about the whole interaction of each participant with the computer was collected, enough to be able to reproduce in detail all the data exploration process for each table. Data collected included traces of the position and speed of the pen on the tablet, mode of exploration, use of the speech and changes in the speed of sonification.

3.3. Analysis of Results

3.3.1. Results from the experiment

Figure 6 shows, for each participant, the average time for the exploration of a table until an answer to the question was given. In the case of some participants (10, 11, 17, 18), the time invested in the exploration increased quite uniformly with the size of the tables, as could have been expected. In the rest of the cases many other combinations occur, and no clear correlation is observed at first sight.

Figure 7 presents results about the percentage of correct answers given by each participant, for different table sizes. Visual inspection of these results does not reveal any clear relationship between both variables (table size and percentage of correct answers). In 15 out of the 648 table explorations performed in total, participants who had identified the correct answer gave a wrong one because the pen slipped to a neighbouring row or column at the time of pressing the button on the pen to obtain details about the exact position. This difference would make the scores in Figure 7 slightly higher, although the difference is very small. It should be investigated, however, how the mechanism to obtain speech details on demand could be improved.




Figure 6. Average time for exploration of tables in the three table size conditions, for each participant in each condition (S: small tables; M: medium tables; L: large tables).




Figure 7. Percentage of correct answers by each participant in each one of the three conditions (S: small tables; M: medium tables; L: large tables).

Finally, Figure 8 shows each participant's overall subjective workload after exploration of the tables in each one of the three conditions. This index was calculated from the ratings provided by the participants in the NASA-TLX forms as the arithmetic mean of mental demand, physical demand, time pressure, effort expended, frustration experienced and performance level achieved, after inverting the scale for the last category mentioned. Also in this case, although a tendency to find higher workload ratings for larger table sizes can be observed in some cases, no conclusions are immediately obvious.

Statistical analysis of variance was carried out to decide if table size could actually have any effect in the parameters considered above (time for exploration, percentage of correct answers and overall mental subjective workload). In every case, the null hypothesis was that table size does not have any effect in the parameters measured. Then, one-factor repeated-measures ANOVA tests were performed for each case.

Results from the analysis of variance showed that table size does have an effect on the time taken to explore the data tables ($F_{17}=8.17; P<0.01$). At a level of significance $p=0.01$, a Tukey HSD *post hoc* test showed that the average time for exploration was longer for large sized tables than for small sized ones, confidence interval $CI=(2.67,21)$. No significant differences were found, however, between small and medium sized tables, with $CI=(-4.42,13.92)$, nor between medium and large sized tables, with $CI=(-2.08,16.25)$.




Figure 8. Overall subjective workload in the NASA-TLX scales, by each participants and in each one of the three conditions (S: small tables; M: medium tables; L: large tables).

In the case of the percentage of correct answers, the analysis of variance showed that table size does also have certain effect on it ($F_{17}=6.88$; $P<0.01$). Performing a Tukey HSD test at a level of significance of $p=0.01$, a significant decrease in the percentage of correct answers was found between medium and large tables, $CI=(-1.70,-20.52)$. However, no significant differences were found, between small and medium sized tables, with $CI=(5.25,-13.58)$, nor between small and large sized tables, $CI=(2.47,-16.36)$.

Similar results were found for the overall subjective workload. The analysis of variance revealed some variation in the overall mental workload due to the change in table size ($F_{17}=6.24$; $P<0.01$). Tukey HSD tests revealed that the increase in overall mental workload was significant ($p=0.01$) between small and large sized tables, $CI=(0.23,3.79)$. Variations in the overall subjective workload were not significant between small and medium sized tables, $CI=(-1.02,2.55)$, nor between medium and large sized tables, $CI=(-0.54,3.03)$.

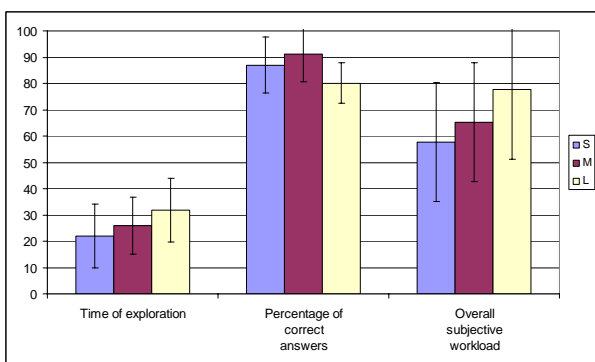


Figure 9. Summary of results, showing overall means and standard deviations for the three measured parameters in the three conditions (S: small tables; M: medium tables; L: large tables).

Results are summarised in Figure 9, in which means of the scores all the participants are graphed are represented for each condition. All three parameters have been standardised to be represented in a common 0-100 scale. Standard deviations are also represented in this figure.

3.3.2. Discussion

Results presented above support the hypothesis that table size has very small effect on the parameters measured in the experiment. For the time for exploration and overall subjective workload in particular, some significant increases were found for the largest increase in table size (from smallest to largest) whereas differences are not significant for intermediate steps. This also suggests that the effect of variations in table size are very small, requiring a very large increase in the amount data for the effects to be significant. Notice that the confidence intervals showing some significance are very close to including zero.

The analysis of the results obtained for the percentage of correct answers given are slightly different. The only significant reduction in the percentage of correct answers was found between the medium and the large sized tables. However, a larger increase in size (from the smallest to the largest sized tables) was not found to be significant. Furthermore, although not a significant difference either, the percentage of correct answers seemed to increase rather than reduce when the size of the tables was increased from small to medium, against what was expected. The authors tend to think that this apparently contradictory behaviour is due to random variations in the results and to the fact that the effect of the table size on the percentage of correct answers, if at all present, is very small. Caution has to be exercised, however, since there could be some other factor that introduces systematic differences of the kind that have been observed. One possibility could be the fact that a table with very few cells mapped to the whole surface of the tablet produces large spaces between cell boundaries. Some participants reported this as being a source of confusion in the case of the smallest tables. Cell boundaries were very far apart and the pen had to travel longer distances to compare adjacent cells, which led to longer periods of silence that made comparing sounds more difficult. Further research is required to clarify this.

These results are similar to those observed when the interface was evaluated with blind and VI users [14] (see Figure 10). Although those results must be taken with care due to the small population that was used (6 participants), they suggest that a similar (slightly longer) length of time is needed to explore a table, with a percentage of correct answers higher than 80%.

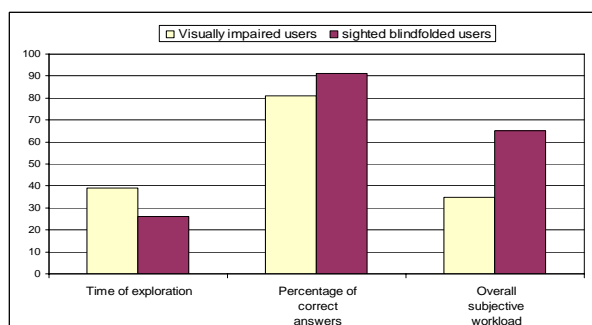


Figure 10. Comparison of results summary for sighted blindfolded and for blind and VI participants.

Sighted participants reported almost twice as much overall subjective workload as blind and VI participants did. An explanation of this difference could be the fact that blind and VI participants rated this interface in comparison to a speech-base interface, and TableVis was rated relative to that one, which produces high levels of subjective workload in

tasks like those described in this paper. It could be expected (although it needs to be investigated) that performance by blind and VI users exploring larger and smaller data tables would be also similar to the performance by sighted blindfolded users reported here.

Finally, it should be observed that the sonification technique employed here seems to be efficient to represent relative values for statistical means of all the sounds in each row or column. As mentioned before, an answer was judged to be correct if it coincided with the highest (or lowest, depending on the case) statistical mean of all the values in each row or column, and results show that answers were correct in well above 80% of the tables explored.

4. FUTURE WORK

As an extension of the experiment presented here, further research could be conducted to find the minimum table size in which the use of TableVis would be beneficial for the users to obtain overview information, in comparison to using more traditional data exploration techniques using speech. An upper limit for the use of this technique could also be investigated.

Finally, ways of extending the technique presented here to support a longer cycle in data analysis, beyond extraction of overview information, can also be investigated. One research direction can be to investigate ways of conducting basic statistical analysis as part of initial data browsing.

5. CONCLUSIONS

Obtaining an overview of a data set is a necessary first step in its analysis, which should be performed quickly and easily. Current accessibility techniques do not facilitate this task for blind and VI users.

TableVis is a multimodal interface to explore and obtain overview information from tabular numerical data tables. This interface uses a physical tangible input device (a tablet) to generate interactive sonifications of the data sets. This technique was shown to be easy to use by blind and VI users and by sighted users alike.

In this paper, the design of the interface has been described. An experiment to investigate the effect of table size on performance using TableVis was reported. Results from this experiment show that performance in data table exploration to obtain overview information using TableVis is largely insensitive to the size of the data set, at least within the broad range of data set sizes considered in this study. Increases in the size of the data sets being explored lead to minimum increases in the time needed for the exploration and in the overall mental workload, and it seems to produce only slight reductions in the percentage of correct answers given to overview questions about the data sets.

6. ACKNOWLEDGEMENTS

This research is supported by EPSRC grant GR/S86150/01, the MultiVis project (www.multivis.org).

7. REFERENCES

- [1] Card, S.K., J.D. Mackinlay, and B. Shneiderman, *Information Visualization*, in *Readings in Information Visualization. Using Vision to Think*, S.K. Card, J.D. Mackinlay, and B. Shneiderman, Editors. 1999, Morgan Kaufmann Publishers, Inc. p. 1-34.
- [2] Shneiderman, B. *The eyes have it: a task by data type taxonomy for information visualizations*. in *IEEE Symposium on Visual Languages*. 1996. Boulder, CO, USA: IEEE Comput. Soc. Press.
- [3] Zhao, H., et al. *Sonification of Geo-Referenced Data for Auditory Information Seeking: Design Principle and Pilot Study*. in *International Conference on Auditory Display (ICAD)*. 2004. Sydney, Australia.
- [4] Pérez-Quiñones, M.A., R.G. Capra, and Z. Shao. *The Ears Have It: A Task by Information Structure Taxonomy for Voice Access to Web Pages*. in *Human-Computer Interaction -- INTERACT'03*. 2003. Zürich, Switzerland: Published by IOS Press, (c) IFIP.
- [5] Saue, S. *A model for interaction in exploratory sonification displays*. in *International Conference on Auditory Display (ICAD)*. 2000. Atlanta, Georgia, USA.
- [6] Bertin, J., *Semiology of graphics : diagrams, networks, maps*. 1967/1983: Madison, Wis. : University of Wisconsin Press.
- [7] Kwasnik, B.H., *A Descriptive Study of the Functional Components of Browsing*, in *Proceedings of the IFIP TC2/WG2.7 Working Conference on Engineering for Human-Computer Interaction*. 1992, North-Holland. p. 191-203.
- [8] Miller, G.A., *The Magical Number Seven Plus or Minus Two: Some Limits on Our Capacity for Processing Information*. *The Psychological Review*, 1956. **63**: p. 81-97.
- [9] Stockman, T., G. Hind, and C. Frauenberger. *Interactive Sonification of Spreadsheets*. in *ICAD 05-Eleventh Meeting of the International Conference on Auditory Display*. 2005. Limerick, Ireland.
- [10] Kildal, J. and S.A. Brewster. *Explore the matrix: Browsing numerical data tables using sound*. in *ICAD2005*. 2005. Limerick, Ireland.
- [11] Mansur, D.L., *Graphs in Sound: a Numerical Data Analysis Method for the Blind*. 1975, University of California.
- [12] Brown, L.M., et al. *Design Guidelines for Audio Presentation of Graphs and Tables*. in *Proceedings of ICAD 2003 Workshop on Auditory Displays In Assistive Technologies*. 2003. Boston, MA.
- [13] Brown, L. and S.A. Brewster. *Drawing by Ear: Interpreting Sonified Line Graphs*. in *ICAD*. 2003. Boston, MA.
- [14] Kildal, J. and S. Brewster. *Non-visual Overviews of Complex Data Sets*. in *Extended Abstracts of CHI 2006*. 2006. Montreal, Québec, Canada: ACM Press.
- [15] Hart, S. and L. Staveland, *Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research*, in *Human Mental Workload*, P. Hancock and N. Meshkati, Editors. 1988, Elsevier Science Publishers B.V. p. 139-183.