

Interactive Auditory Data Exploration: A Framework and Evaluation with Geo-referenced Data Sonification

Haixia Zhao¹, Catherine Plaisant¹, Ben Shneiderman¹, and Jonathan Lazar²

¹Human-Computer Interaction Lab, University of Maryland, College Park, MD USA

²Dept. of Computer and Info Science, Towson University, Towson, MD USA

Contact author: Catherine Plaisant (plaisant@cs.umd.edu)

Abstract

We describe an Action-by-Design-Component (ADC) framework to guide auditory interface designers for exploratory data analysis. The framework characterizes data interaction in the auditory mode as a set of Auditory Information Seeking Actions (AISAs). Contrasting AISAs with actions in visualizations, the framework also discusses design considerations for a set of Design Components to support AISAs. Applying the framework to geo-referenced data, we systematically explored and evaluated its design space. A data exploration tool, iSonic, was built for blind users. In depth case studies with 7 blind users, with over 42 hours of data collection, showed that iSonic enabled them to find facts and discover trends of geo-referenced data, even in unfamiliar geographical contexts, without special devices. The results also showed that blind users dramatically benefited from the rich set of task-oriented actions (AISAs) and the use of multiple highly coordinated data views provided by the ADC framework. Some widely used techniques in visualization, with appropriate adaptation, also work in the auditory mode. The application of the framework to scatterplots shows that the framework can be generalized and lead to the design of a unified auditory workspace for general exploratory data analysis. Readers can view a supplementary video demonstration of iSonic by visiting www.cs.umd.edu/hcil/iSonic/.

Author Keywords

Interactive sonification, auditory user interfaces, information seeking, vision impairment, universal usability

ACM Classification Keywords

H.5.2.a Auditory (non-speech) feedback; H.5.2.e Evaluation/methodology; H.5.2.q. User – centered design

General Terms

Human Factors

1 INTRODUCTION

Information visualization has produced many innovative techniques/interfaces for people with normal vision to use their tremendous visual ability to explore and discover data facts/trends.

When information is presented with visual properties, such as color and spatial location, it is not easily viewable by users with any level of visual impairment. In addition, visual data interaction is typically done by using pointing devices, such as computer mice, to directly manipulate the visual objects displayed on the screen. Such interaction is hard without sustained visual feedback. Although a few visualization tools (e.g., [35]) allow keyboard-only navigation inside some visual graphs, most visualizations are not usable for users with vision impairment.

One example is the current web-based access to government statistical data. Such data is often geography-related, such as population distribution by geographical regions, and often presented as choropleth maps that typically use colors to show the value for each map region. U.S. Federal law (www.section508.gov), requires that all government information on the web be available to users with disabilities. Other governments around the world have similar types of rules related to web-based information (see www.w3.org/wai for government policies around the world). The paradox is that government agencies are presenting more data through choropleth maps, but at the same time, are increasingly required to make that data accessible to people with disabilities. Therefore, this research not only has technical implications, but also government policy implications.

A widely used accommodation for users with vision impairment to access digital information is to rely on screen readers, such as Window-Eyes and JAWS, to speak the textual content [28]. While JAWS and Window-Eyes are the most popular screen readers, new approaches to presenting text have been presented in the literature. For instance, the BrookesTalk screen reader provided more navigational information to users [36] and math markup languages improve the comprehension of math formulae through screen readers [12]. However, the drawback of the new approaches is that they all focus on other ways of presenting text, not methods for presenting graphical data. To make non-textual elements accessible to screen readers, textual equivalents are needed. For static graphs, it is a standard practice to provide textual labels during the system development [30]. For dynamic graphs, tabular data presentations are used instead (e.g., [35]), or textual summaries can be automatically generated from the data set.

Several problems exist in the current approaches. First, while a concise textual description is helpful, the data interaction that is a critical part of data exploration process is lost. Automatic textual summarization techniques require pre-defined summary templates and do not have enough flexibility to support all user needs in exploratory data analysis. Second, a tabular presentation may be good for basic data browsing but is hard for in-depth data comprehension and analysis. Third, speech can accurately describe information but tends to be long in duration and hard to realize complex information.

Sonification seems to hold the most promise for building interfaces to help data exploration for users with various levels of visual impairment. The success of Braille-based interfaces have been limited in the past, both due to the low levels of Braille literacy (some estimates are 20%) and the high cost of refreshable Braille displays [18]. Tactile maps would seem to be the most likely choice for understanding data on maps, but the high costs of tactile displays mean that they are rarely used [33]. Even special printing devices, such as printers that can print raised paper maps, are prohibitively expensive, and rarely available outside of schools [33]. The most commonly used assistive technology for users with any level of visual impairment is the screen reader, which requires no special experience or education, works off of the standard computer hardware (the speakers), and only requires a software application. Furthermore, users with visual impairment who have been relying on their hearing for a number of years, are able to discern

more detail than the typical user with vision, who would not notice the intricate details being produced by the computer speaker.

Data interaction has been extensively investigated in visualization systems. But little was done regarding whether techniques in visualizations can be translated for use in auditory data exploration without visual aids, and what design implications are involved. Some research used musical sounds to present sonified “overviews” of simple graphs (e.g., [7]) but support for other task-oriented data interactions is typically missing.

We believe it is important to investigate whether an analogue to standard techniques in visualizations can be established for the auditory mode. In this paper, we first describe an Action-by-Design-Component (ADC) framework for designing auditory interfaces for analytical data exploration. We use a set of Auditory Information Seeking Actions (AISA) to characterize task-oriented data interaction without visual aids, identify Design Components for supporting AISAs, and discuss their general design considerations by contrasting AISAs with actions in visualization. This framework has been used to investigate the design space of geo-referenced data sonification. In our earlier work, we gave a partial preliminary version of the framework and reported on some initial sonification designs for geo-referenced data [37][38]. The focus of the earlier work was limited to actions for conveying data distribution patterns on maps. Little evaluation was done regarding the support for general data exploration tasks or the validity of the partial framework. In addition, the initial studies were conducted with blind-folded sighted users, not actual blind users. The ADC framework presented in this paper substantially extended the earlier version, and the current work includes case studies of blind users.

Guided by the ADC framework, we now developed a general exploratory data analysis tool for users with vision impairment, called iSonic. iSonic contains three highly coordinated views and supports AISAs within and across the views. Figure 1 shows the table view and map view. The integration of the scatterplot view is described in Section 7. We will describe iSonic features and discuss the design rationale to illustrate the framework.

Afterwards, we report an empirical evaluation of the keyboard-only version of iSonic with 7 users with complete vision impairment (42 hours of in-depth observation and interview data.) which enabled us to examine the effectiveness of iSonic design choices, to draw general design implications, and to validate the benefits of the ADC framework. By applying the ADC framework to scatterplots and integrating a sonified scatterplot into iSonic as the third highly coordinated data view, we demonstrate the generalizability of the framework which can lead to the design of a unified auditory workspace for general exploratory data analysis. Note: in this paper, when use the term “user with visual impairment,” we refer to users with partial vision, and when we use the term “blind,” we refer to users with no residual vision.

2 PREVIOUS WORK

Sonification, the use of non-speech sound, has been used in various interface designs (e.g., non-visual GUI presentations [3][17]), as well as data presentations [13]. Using the highly structured nature of musical sounds to convey information works even when no everyday auditory equivalence exists, is less tiring and generally more appropriate than everyday sounds [27]. Research has shown that musical sounds enhance numeric data comprehension (e.g. [21]) and humans can interpret a quick sonified overview of simple data graphs (e.g., [2][7][8]). Some guidelines were extracted (e.g., [2][31]) and toolkits were developed to help researchers try different data-to-sound attribute mappings (e.g. [20][32]). While some allow basic user

movements in the graph (e.g., [2]), previous data sonification typically lacks supports for task-oriented data interactions.

In visual data exploration, the information seeking mantra “overview first, zoom, filter, then details-on-demand” [25] characterizes the general visual information seeking process and was an effective visualization design guideline. Several visualization interfaces (e.g., Visage [23], Snap-together [19]) were designed that allow users to construct multiple graphical data views and perform data exploration through unified interaction methods within and across the views.

However, such a framework or interface is absent for data exploration in the auditory mode without visual aids. Some recent models (e.g. [10][24]) tried to describe interactive data sonification, but they emphasize spatial immersion effects in a physical world modeling of the data set hence may not be suitable for abstract data. More importantly, none has characterized task-oriented information seeking needs in the auditory mode, or addressed design considerations for interaction without visual aids, such as “can users with complete vision impairment operate multiple coordinated auditory views”.




Figure 1: highly coordinated table and map views of the counties of the state of Maryland. Superimposed on the color coded (choropleth) map is a representation of the recursive 3x3 keyboard exploration grid.

3 AN ACTION-BY-DESIGN-COMPONENT FRAMEWORK

The Action-by-Design-Component framework (Figure 2) contains a set of Auditory Information Seeking Actions (AISA), and a set of Design Components to support the actions. In this section, we first describe AISAs, explaining their connection and difference from visual actions. Then we briefly mention some general design considerations for the Design Components to support AISAs. Those considerations will be reviewed in more detail when we discuss the design of iSonic.

3.1. Auditory Information Seeking Actions (AISA)

We believe that an exploratory data analysis task in the auditory mode can be accomplished by a series of Auditory Information Seeking Actions (AISA). Many of the actions resemble those in visual information seeking [25] but involve different cognitive processes and present special design challenges due to the highly transient nature of sound. In our earlier work, we gave a

partial preliminary version of the framework including “gist, navigate, filter, and details-on-demand” [37]. The framework is substantially extended here.

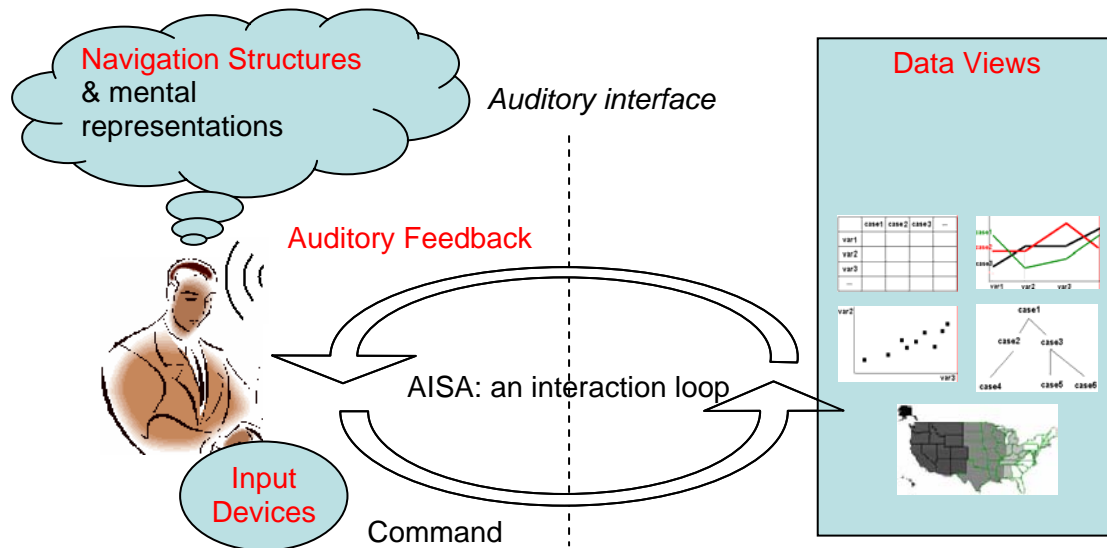


Figure 2: An Action-by-Design-Component framework. An Auditory Information Seeking Action (AISA) is an interaction loop involving four Design Components (in red).

Obtaining a *gist* is to experience the overall data trend via a short auditory message. It guides further exploration and may allow the detection of anomalies and outliers. A gist is an auditory “overview” but has special design and cognition challenges (see Section 3.2) because human auditory perception is much less synoptic than visual perception. Visual overview has tremendous value in visual data exploration. Previous work (e.g., [2][7]) showed that sighted users can interpret data trends from a gist of simple graphs, such as scatterplots and line graphs, yet there has not been enough understanding of how valuable a gist is in the auditory mode and how blind users would actually use gists in exploratory data analysis.

Navigation is “moving around” to examine portions of the data set by listening to a sub-gist of that portion. It needs to follow paths that are natural to the data relations. A visual interface provides a sustained display for users to directly manipulate. In auditory interfaces, users need to construct a mental representation of the display space and virtual navigation structures in order to efficiently move in the data set. Without a persistent display, they can easily get lost. To regain the orientation, users need to *situate* themselves by requesting their status. While navigation is an exploratory action, *searching* is a more fixed-goal action that directly lands on the data items by specifying search criteria. Searching breaks the process of mental representation construction, so situating may be needed to regain orientation after the search is completed.

Filtering out unwanted data items according to some query criteria helps to trim a large data set to a manipulable size, and allows users to quickly focus on items of interest. In visualization, dynamic query coupled with rapid (less than 100 milliseconds) display update is the goal [24]. In the auditory mode, different goals need to be established because such a short time is usually not enough to present a gist of changes. Auditory feedback about the filtering results needs to be given after filtering is done instead of continuous display updates during the filtering process.

When the number of items is small, users can listen to the details. While speech is often too lengthy for obtaining an overall gist, it can be an effective presentation at the *details-on-demand* level. It is hard to understand a data element without the appropriate context. On the other hand, too much detail slows down the sequential presentation and can be overwhelming. Multiple information detail levels that incorporate different combinations of speech and non-speech sounds are needed.

By *selecting*, users specify special interest in particular data items. Those data items are marked and can be revisited later or examined in other contexts.

In visualization, linked brushing allows users to manipulate the data in one view while seeing the results in other views. It requires users to construct and maintain multiple mental representations of the data views simultaneously which can be mentally intensive in the auditory mode. Additionally, auditory feedbacks from multiple views need to be clearly distinguished to avoid confusion and overloading. In the auditory mode, *brushing* can be done in a sequential style by selecting data items in one view, then explicitly switching to another view to examine them in a different data relation. We also need to understand whether blind users can handle and benefit from brushing among multiple views.

3.2. Design Components

Each AISA consists of one or multiple interaction loops in which the user uses an *input device* to issue a command and listen to the *auditory feedback*. The center of the loop is the *data view* that governs the *navigation structure*, allowing the user to build a mental representation of the data space and correctly interpret the auditory feedback.

A data view is a form of presenting the data items and their relations, such as a table, map, scatterplot, or line graph. Research has shown that users with vision impairment were able to learn, interpret, and benefit from non-tabular data presentations. There is also evidence [1][37] that choosing the right data view for a given task dramatically influences performance.

Navigation structures should reflect the data relations in the data view. In some previous work, users used a mouse or other input devices to move in the 2-D or 3-D data space to activate sounds of the data items within a certain distance from the cursor position. Such a “torch metaphor” [6] navigation could be useful for some data views, e.g., a scatterplot, but may be inefficient for others, e.g. a node- link diagram.

The choice of input device needs to consider both effectiveness and universal availability. Speech as input can be tempting but lacks the kinesthetic feedback users can get from operating physical input devices. Sensory feedback can help with users’ orientation and mental representation in the interaction. Card et al. [4] categorized physical input devices by their physical manipulation properties and defined several choice factors such as the cost. We can maximize users’ situation awareness by matching an input device’s properties with those of the navigation structures. However, it is important to keep the system device-independent by providing good alternatives in the absence of the desired device. For success by users with vision impairment, a system should provide interactions optimized for keyboard-only operations, as a keyboard is standard equipment for nearly all computers, and it is frequently used by users with any level of visual impairment (including full blindness), as they have typically memorized the keyboard layout.

As a general principle, the auditory feedback should have a low latency. It should be generally short to fit the short-term memory (STM) or allow pauses for midpoint STM processing. Short and responsive feedback increases user engagement and allows users to quickly refine their control activities in the exploration process. It should synchronize with other display modalities to allow perceptual combinations. While humans are good at selective listening, attending to multiple simultaneous sounds is difficult and the amount of accurate information that can be extracted from simultaneous sound streams is limited [9]. The sounds of multiple items often need to be sequenced along the time dimension instead of being played all at once. This imposes special design challenges when no natural mapping exists from the data relation to the time dimension. When the number of data items is large, data aggregation may be necessary to design short feedback.

4 ISONIC AND THE DESIGN SPACE FOR GEO-REFERENCED DATA

Guided by the ADC framework, we have systematically explored the design space for geo-referenced statistical data and designed iSonic (Figure 1). Two users without residual vision were involved in the iterative design process. Many iSonic design decisions were based on their suggestions, as well as results from earlier evaluations of some potential design choices.

4.1. iSonic Design Components

Data Views

iSonic provides three highly coordinated data views – a region-by-variable table, a choropleth map, and a bivariate scatterplot. The scatterplot view was integrated into iSonic as an extension after the user evaluation and is described in section 7.

The table shows multiple statistical variables simultaneously. Each row corresponds to a geographical region and columns to variables. Table rows can be sorted by pressing ‘o’ while at the desired column, allowing quick location of low or high values. While geographical coordinates and adjacencies could be added as table columns, such information is better displayed on a map. Subjects in our previous study [37] strongly preferred a map over a table for discovering geographical value trends and performed better on pattern recognition tasks with a map than with a geographical knowledge enhanced table. Other views, such as line graphs or scatterplots, can be helpful for some analytical tasks. They were not included in the evaluation because we wanted to first examine how users could operate multiple coordinated auditory views. Auditory and visual displays are synchronized to allow communication between sighted users and users with any level of visual impairment.

Input Devices and Navigation Structures

When choosing input devices, we considered both device availability and how effectively their physical properties match the navigational properties of the two data views.

In iSonic, the table navigation follows the row and column table structure. It is discrete and relative because what matters is the relative row/column order, not the exact spatial location or size of each table cell. On the other hand, the map navigation follows the regions’ positions and adjacencies. Both the relative region layout and the absolute region locations and sizes are useful.

iSonic works with a keyboard alone. A keyboard is available on most computers and users with any level of visual impairment are very comfortable using the keyboard. We use the arrow keys

as natural means for relative movements in the left, right, up, and down directions. The numerical keypad potentially allows relative movements in 8 directions. Furthermore, blind users tend to have limited or no experience with using mice [33]. The keyboard can also be transformed into a low resolution 2-D absolute pointing device, e.g., by mapping the whole keyboard layout to 2-D screen positions. In iSonic, we map the 3x3 layout of the numeric keypad.

iSonic also works with a touchpad. Touchpads are relatively common. A 14" touchpad costs less than \$150. A touchpad provides high resolution 2-D absolute pointing and allows continuous movements by fingers. The kinesthetic feedback associated with arm and finger movements, combined with the touchpad frame as the position reference, may help with users' position awareness on maps. Tactile maps placed on the touchpad can be helpful [14], but we chose not to rely on them because they need to be changed when the map changes and tactile printers are expensive and rarely available. When resources are available, a generic grid with subtle tactile dots may be used instead as a position and direction aid.

Auditory Feedback

iSonic integrates the use of speech and musical sounds. Values are categorized into 5 ranges, as in many choropleth maps, and mapped to five violin pitches. The same mapping is used in the table view. Various musical instruments are used to indicate when users are outside the map or crossing a region border in the touchpad interface, or crossing a water body to reach a neighboring region in the keyboard interface. Stereo panning effects are used to indicate a region's azimuth position on the virtual auditory map. It is also used in the table to indicate the column order. Using the plus and minus keys, users can switch among four information levels for each region: region name only, musical sound only, name and sound, name and sound plus reading of the numerical value.

There are many alternatives. Sound duration can present the value but would significantly prolong the feedback and is not appropriate when values of many regions need to be presented. Region locations could be mapped to sound locations using virtual spatial sound synthesized with Head Related Transfer Functions (HRTF) [34]. Spatial sound provides high perceptual resolution in the azimuth plane, but is not satisfactory in the elevation plane, especially when a generic HRTF is used. Using individualized HRTF could improve the elevation perception but its measurement is a long process requiring special equipment and careful calibration. Additionally, HRTF spatial sound is computing intensive. While we have connected iSonic to a virtual spatial sound server and plan to investigate the use of individualized HRTF spatial sound, we currently focus on MIDI stereo sound. We also tried to play a piano pitch after each violin value pitch to indicate the region's elevation position. Unfortunately, such extra sound was not found to increase performance [38].

4.2. iSonic Auditory Information Seeking Actions

iSonic supports AISAs in both the table and the map views, including sequential brushing between the two views. Each interface function can be activated from a menu system that also gives the hotkey and a brief explanatory message.

Gist

In the table view, a gist is produced by automatically playing all values in a column or a row. The sequencing follows the values' order in the table, from top to bottom, or left to right. In the

map view, there is no natural mapping from the geographical relation to the time relation. Research has shown that sequencing that preserves spatial relations helps users to construct a mental image of the 2-D representation. Sequencing is done by spatially sweeping the map horizontally from left to right then vertically, like in a typewriter. When the end of sweep row is reached, a tick mark sound is played and the stereo effect reinforces the change. A bell indicates the end of the sweep of the whole map. The same sweep order holds for sub-gists of parts of the map. For both views, the current information level controls the amount of detail in the gist, thus controlling its duration. For example, when the information level is set to “musical sound only”, a sweep of the entire US state map containing 51 regions lasts for 9 seconds.

Navigate

Table navigation is done by using arrow keys to move up, down, left, right, and to top, bottom, left and right edges. Users can press ‘u’ to switch between two modes. In the cell mode, the current cell is played. In the row/column mode, a sub-gist of the whole row or column is played. While it is easy to navigate the table, using a keyboard to navigate maps with irregularly shaped and sized regions brings special design challenges. Relative movements between neighboring regions reveal region adjacency but do not convey region shapes, sizes, or absolute locations. Subjects in our previous studies reported that they only had weak location awareness by using this navigation method. Furthermore, it is a challenge to define a good adjacency navigation path for a map that is not a perfect grid. A movement may deviate from the direction users expect. Reversibility of movements can also be a problem in which a reversed keystroke may fail to take the user back to the original region. To tackle some of the problems, we tested cell-by-cell movements on a mosaic version of the map [38]. However, it did not improve users’ location awareness, and was much less preferred because it required more keystrokes to move around.

We expect that navigations based on absolute pointing may help. Kamel and Landay first used a 3x3 grid recursion method via the keypad in a drawing tool [11]. In iSonic, the map is divided into 3x3 ranges (Figure 1) and users use a 3x3 numerical keypad to activate a spatial sweep of the regions in each of the nine map ranges. For example, hitting ‘1’ plays all regions in the lower left of the map, using the same sweep scheme as the overall gist. Users can use Ctrl+[number] to zoom into any of the ranges, within which they can recursively explore using the 3x3 pattern or use arrow keys to move around. Pressing ‘0’ sweeps the current zoomed map range or the whole map.

With the touchpad, users drag their fingers or press spots on the smooth surface touchpad to activate the sound of the region at the finger position. Stereo sounds provide some complementary direction cues. The sound feedback stops when the finger lifts off. The touchpad is calibrated so that the current map range is mapped to its entire surface. Preliminary observations suggest that both the keyboard and touchpad navigations allow users to gain geographical knowledge. A controlled experiment is planned to compare them in details.

Details-on-demand and Adjust information level

Pressing ‘space’ plays the details of the current region. Another way to get the details is to increase the information level to the maximum level in which all details of a region are given by default when users navigate to that region.

Situate

When users press ‘I’ (as for ‘Information’), iSonic speaks the current interface operational status. In the table, it includes the row/column counts, headings of the current table position, navigation mode, sorting status, regions selected, and so on. In the map, it includes the name of the variable displayed, navigation position, regions selected, and so on.

Select

In both views, users can press ‘L’ (as for ‘Lock’) to select/unselect the current region and press ‘A’ to switch between “all regions” and “selected regions only”. In ‘selected regions only’, AISAs only activate sounds of the selected regions.

Brush

Brushing is done by users switching back and forth between the two views. The views are tightly coupled so that action results in one view are always reflected in the other. For example, users can select a region in the table view and show “selected regions only”. When users switch to explore the map view, only the selected region will be played. By sweeping each of the 9 map ranges, users can roughly but rapidly locate the region on the map.

Filter and Search

Filtering was done by slider-based queries. It is complex even for sighted novice users and was not evaluated in the current evaluation. Searching is obviously helpful but was not implemented at the time of the evaluation.

5 ISONIC IMPLEMENTATION

iSonic was implemented as a Java application that also runs from the Web through Java Web Start Technology [29]. The GUI part was written in Java JFC/Swing and the musical sounds were typically produced through the Java MIDI Sound technique. Speech was produced by sending command network datagram to an accompanying speech server built on Microsoft Speech SDK 5.1. Although the speech operates in a manner similar to a screen reader, including the speech directly in the application, rather than utilizing a screen reader application, provides more flexibility and also more stability of the speech output. Screen readers frequently conflict with applications, and cause crashes, which are one of the most frequent causes of frustration for blind users [15]. Incorporating the speech output directly into the iSonic application greatly decreases the likelihood of a screen reader-caused crash, and is more likely to lead to accurate results in evaluating the iSonic tool. iSonic can also produce virtual spatial sounds by sending command network datagram to a spatial sound server that simulates real world sounds using Head Related Transfer Functions (HRTF). The spatial sound server was developed at the Perceptive Interfaces and Reality Laboratory, University of Maryland at College Park. Additionally, iSonic can be configured to load and play prerecorded speech or musical sound files. iSonic is a real time application, giving real-time highly synchronized visual and auditory responses to every user action, with delays usually less than 100 milliseconds. Readers can try iSonic or viewing the video demonstration by visiting www.cs.umd.edu/hcil/iSonic.

iSonic implementation was based on the Model-View-Controller (MVC) paradigm [22] commonly used as a design framework for modern GUI interfaces. It extends the MVC paradigm by treating the auditory display and GUI display as two parallel subviews (a visual view and an auditory view) residing inside each data relation view (e.g., map, table, and scatterplot). For each

data relation view, the auditory display is produced by the creation and execution of various sound objects (e.g., MIDI events, speech objects). The two displays are synchronized by monitoring the execution of the sound objects. Each data relation view has related controller classes that define the mapping from input devices to interaction commands. Model classes that contain the input data set and each data item's interaction status (e.g., whether a data item is selected) are shared among all data relation views. Each data relation view may also have its own model classes (e.g., map shape data for the map view).

iSonic has a rich set of configuration parameters that can be used to customize its visual, auditory, and interaction behaviors. This allows iSonic to be adapted for different users and to be used as a research tool to compare some design options.

To generate map gists, iSonic accepts user defined map spatial sweep orders as part of the input data. It can also automatically produce sweep orders in the absence of predefined orders. It is a challenging problem to automatically produce map spatial sweep orders that are congruous to what people may define manually based on their visual impressions. iSonic uses a preliminary algorithm that first recursively split the regions on the map into multiple subsets with as little inter-set intrusion as possible, then connect the regions inside each subset to produce the sweep path. The algorithm uses a greedy approach and does not guarantee globally optimized solutions, but produces good results for several U.S. maps tested. More details on the problem of map spatial sweep order generation are available in the leading author's Ph.D. dissertation [39].

6 EVALUATION

During early design iteration for iSonic, controlled experiments were conducted to compare the effectiveness of design alternatives including the choice of data views, map navigation methods and sound encoding schemes [37][38]. However, an exploratory data analysis task is a complex process that involves many interface components. During the process, many inherent human subject variations can come into play, such as experience and cognitive styles. In order to obtain insights into users' auditory information seeking behaviors, we chose to conduct case studies. Through a combination of direct observation, thinking aloud protocol, and in-depth interview, case studies can reveal the underlying design strengths and weaknesses, and capture common user behaviors as well as individual differences.

During the summer of 2005, we conducted intensive case studies with 7 local users who are blind, producing 42 hours of observation and interview data, with an average of 6 hours per user. Using cross-case analysis, we were able to extract common user behaviors and feedback that allowed us to (1) evaluate the effectiveness of iSonic design choices; (2) identify features helpful to each data exploration task category and examine the utility of the ADC framework; (3) identify task road blocks in order to target training and modifications to the interface and the framework.

6.1. Subject Profiles

All seven subjects possessed basic computer skill and relied on screen readers to access computer information. They were all comfortable with maps and tables, had experience with numerical data sets, and used government statistical data at work. All subjects were in the age range of 23 to 55. Three of them were born blind (P2, P3, P4) and the others became legally blind after 15 (P1, P5, P6, P7). None of them had residual vision, and none of them were newly-blind. Among the born blind, 2 were males, one with a college degree (P2) and the other with a

doctorate degree in law (P3). The remaining female (P4) had a masters degree in English. Among the subjects who became blind after 15, one was a male (P7) with a college degree in business and commerce. The other male (P1) was about to finish college in science and technology. For the two females, one had a college degree (P5) and the other had a master degree (P6), both in social science. All subjects volunteered to participate, and were compensated for their time.

6.2. Material and Tasks

The studies used the basic iSonic configuration that is accessible to most computer users: stereo auditory feedback through a pair of speakers and a standard computer keyboard as the input device.

Three data sets were used, one for training, one for testing, and one for post-test free exploration. The data was 2003 census data on general population information, employment of population with a disability, housing value and vacancy, education levels, and household income. The training data set contained 8 variables and was about the 50 US states plus the District of Columbia. The test data set contained 12 variables and was about the 24 counties in the state of Maryland. The post-test data set was about the 44 counties in the state of Idaho (but subjects were not told what it was). Subjects' geographical knowledge of US states and Maryland counties ranged from excellent to very poor. This allowed us to observe the influence of geographical knowledge on task behaviors and interface usability.

Seven tasks were designed for each data set. The tasks are based on those used in previous research on finding statistical data on government web sites [5]. Three tasks required value comparison in the geographical context (T5, T6, T7), and four did not need any geographical knowledge (T1, T2, T3, T4). Task orders were different between the training and testing sessions, but were consistent for all subjects. The testing tasks are summarized below.

T1: (Find min/max) Name the bottom 5 counties with the lowest housing unit value.

T2: (Find the value for a specific item given the name) What is the population of Dorchester county?

T3: (Correlation) Which of the two factors is more correlated to "Median household income": "percent population with bachelor's degree and above", or "Percent employed population"?

T4: (Close-up item comparison) For what factor(s) does Montgomery county do better than Frederick county: (1) employment rate for population with a disability, (2) percent population with at least college education, (3) household income, and (4) average housing unit value.

T5: (Find items restricted first by value relations then by geographical locations) How many of the bottom 5 counties with the lowest housing unit value are in the western part of the state? Name them?

T6: (Find items restricted first by geographical locations then by value relations) For all three counties that border Frederick, plus Frederick, which one has the highest percent housing unit vacancy?

T7: (Value pattern in geographical context) Comparing "population with a disability" and "percent population with a disability", which variable generally increases when you go from east to the west and from the north to the south.

Subjects also performed a similar set of testing tasks in Microsoft Excel 2002 with their usual screen readers (all happened to have experience with the JAWS screen reader), and compared the task experience. It was not our intention to compete with Excel. Rather, we considered Excel as the standard tabular data viewer, and used the comparison as a method to solicit user comments on what interface features were helpful to each task. All subjects had some previous experience with Excel, while some were expert users. We did not provide tactile maps when subjects used Excel, because many blind users do not have access to tactile maps (only P5 uses one, owned by the state government of Maryland).

6.3. Procedure

Each case study was carried out in two sessions on consecutive days, at the subject's home or office. In the first session, the subject listened to a self-paced auditory step- by-step tutorial, tried out all iSonic features and practiced seven sample tasks with the training data set. For each training task, a sample solution and the correct answer were given. Subjects could either first try to solve the task on their own, or directly follow the sample solution.

In the beginning of the second session, those subjects with limited Excel experience were given time to practice. After adjusting the speech rates to the subjects' satisfaction, they performed seven tasks similar to the training tasks in both Excel and iSonic. For each pair of tasks, the subject first did the Excel task then the iSonic task and finally compared the interface experience for that task. The iSonic task was similar to the Excel task but modified. They used the same testing data set but involved different variables, so data learning between tasks can be ignored. We asked the subject to do the Excel task first because we wanted to minimize the effect on the Excel task resulting from the geography learning in the corresponding iSonic task. While there was a chance of strategy transfer from the Excel task to the iSonic task, the Excel task might also have benefited strategically from the iSonic training task. An interview was conducted after subjects performed all the testing tasks in both interfaces. Finally, subjects were asked to freely explore an unknown map and data (the post-test data set) for 5 minutes and report things they found interesting. This was to observe what users would do when they encountered a new map and data.

6.4. Cross-case Task Completion and Error Analysis

After spending an average of 1 hour 49 minutes going through all the interface features by following the tutorial, subjects successfully completed 67% of the training tasks without referring to the sample solution or any other help. After the training, subjects were able to retain their newly acquired knowledge and successfully completed 90% tasks on the next day in a different context without any help. For 74% of the tasks that subjects used different strategies than the given solution in the training, they adopted the sample strategies in the test session. Details on each subject are available in the leading author's Ph.D. dissertation [39].

For tasks that did not require geographical knowledge, the average testing success rates were similar for iSonic and Excel, both at 86%, although subjects ranked iSonic easier than Excel, at 7.9 vs. 7.0 based on a 10-point scale (a higher number being easier). The explicitly reported reasons, in decreasing order of frequency, included: (1) the pitch was helpful in getting the value pattern and comparing values; (2) it was easier to sort in iSonic because sorting was done by pressing one key in the desired column to toggle the sorting status, instead of handling multiple widgets in the dialog window as in Excel; (3) it was helpful to isolate a few regions from other

interfering information by selecting; (4) It was flexible to adjust the information level during the task; (5) there was more than one way to get the same information.

For geography-related tasks, the average testing success rate was 95% in iSonic. In Excel, the two subjects with excellent knowledge about Maryland geography (P3 and P5) achieved a success rate of 67%. Other subjects either skipped some tasks due to the lack of geographical knowledge or tried to make an educated guess but gave incorrect answers, resulting in an average success rate of 20%. On a 10-point scale, subjects gave iSonic an average of 8 on easiness for all the 3 geography-related tasks, and gave Excel an average of 5.8 for the tasks they performed. The explicitly reported reasons included: (1) the map was easy to use and very helpful (mentioned by all subjects in all 3 tasks); (2) it was great to be able to switch between the map and table, select things in one view then look at them in another; (3) the pitch was helpful in getting the value pattern and comparing values; (4) there was more than one way to get the same information.

Overall it was easy for the subjects to choose an efficient combination of interface features to do the tasks (average 7.4 on a 10-point scale with 10 being easy). Correlation tasks, however, turned out to be challenging. Most subjects understood the concept but did not know how to do it efficiently in iSonic until they viewed the sample solution. Only P7 easily came up with the sample solution. He sorted the main variable ascending in the table view, then in the row/column navigation mode, swept other columns with “pitch only” to check which one has more consistently increasing pitch pattern. Other subjects mostly went across all requested columns to check if the pitches or numbers were consistently small or large for each region. Some also sorted one or all columns. One subject (P6) said she would have the data plotted in a scatterplot or multi-line graph and had her human reader look for the highest correlation. All subjects, except P4, were able to learn from the sample solution and successfully applied it in the test session. The geographical value pattern tasks were easy for most subjects except for P7 who guessed the answer correctly but was very uncomfortable. Instead of “visualizing the map”, he emphasized accuracy by trying to calculate and compare the average value for each of the 9 map ranges. This was consistent with our earlier finding that task strategies affect geographical value pattern recognition [38].

Setting aside the above strategic difficulties, incorrect answers in iSonic were caused by two common errors: (1) subjects sorted the wrong variable (a third of all errors). This might be due to the high similarity of variable names, and that the interface did not confirm the variable being sorted when the sorting key was pressed. (2) Some subjects skipped the 1st region in the table (a third of all errors) because pressing the down arrow key after hearing “already top edge” took the subjects to the 2nd row instead of the 1st.

6.5. Task Strategies and iSonic Usage Patterns

Map vs. table

All subjects used the table for most value comparisons, and used the map when they needed to compare items in the geographical context (e.g., T7) or to acquire/confirm region locations. The table was often used to change the variable to display on the map, but more importantly, the sorting feature was used to find minimum or maximum values, named regions, and values of specific regions. The table was also used to compare the values of multiple regions, and to check correlations. The map was used sometimes by a few subjects to find regions.

Brush

All subjects became proficient in switching between the table and the map views according to the changing needs for data relations during the task. The tight coordination between the map and the table views was considered the most significant strength of iSonic by all subjects. “It is cool to select things in one view and look at them in the other”. “The biggest advantage of this tool is the ability to quickly change between the table view and the map view”. To find items restricted first by value relations then by geographical locations (e.g., T5), most subjects first used the table to find items meeting the value restriction, selected to isolate them, then switched to the map to check their geographical locations. Some subjects skipped the use of the map and used their pre-test geographical knowledge to judge if the selected items satisfied the geographical restriction. A few subjects first used the map to find all items that met the geographical restriction, remembered them, then sorted the table to find items satisfying the value restriction, and reported the intersection of the two sets. The latter two strategies relied on subjects’ memory of the intermediate results and caused some errors. Subjects said they would have used selecting to mark items during view switching if the number of items were larger. To find items restricted first by geographical locations then by value relations (e.g., T6), most subjects first found and selected items meeting the geographical restriction on the map, then either used the pitch and value in speech to check if they meet the value restriction, or switched to table and used sorting to compare their values.

Use of Pitch

Using pitches to present numeric values was considered intuitive, entertaining and very helpful to data comprehension. It took some subjects a few tasks to get used to this idea but they became increasingly inclined to using pitches for both trend analysis and close-up value comparison. “Pitch makes it a lot easier and quicker to compare values”. “Tones are very helpful to find patterns in a series of values. In some extent it helps me to do things I used to do with (visual) graphs”. “All the other applications are boring. iSonic has its personality. It has the map that I really enjoyed. The tones are entertaining and fun”. To use pitches, most subjects either changed to the pitch-only information level (especially for trend analysis), or used the level with both pitches and numbers in speech, but quickly navigated through items, only waited for a number to be spoken for confirmation purpose (in value comparison). Some subjects were able to tell the absolute value category using only one pitch while some needed to use other pitches as references. All subjects, except P4, were comfortable with the simultaneous pitch and speech presentations. P4 reported that pitches and speech interfered with each other, and requested to tone down the pitch volume. However, she declined the suggestion to completely remove pitches, because she used pitches exclusively in trend discovery.

Adjust Information Level and Details-on-demand

All subjects frequently adjusted the information level during a task. Subjects mostly used name plus pitch or name plus pitch together with the value in speech. When the information level with value in speech is used, many subjects cut it short by navigating to another item before the value speech finished, and only waited for it to finish when they wanted to confirm the value. In automatic map sweep searching for a region, spoken values were typically removed. To sweep the map or a table column for value patterns, e.g., for geographical patterns or correlations, most subjects used the pitch-only level because it let them skim through the data the fastest. A few subjects chose to keep the names on to keep track of the meaning of each sound while still being

able to go through the data at a decent pace. To find a named region on the map, P7 often used the “name only” level. Details-on-demand was mostly done by increasing the information level to the maximum level instead of pressing the ‘space’ key.

Gist

Table sweep was very intuitive. To check value patterns, e.g., for the correlation tasks, some subjects used an automatic pitch-only sweep of each column by navigating the table in the row/column mode.

Automatic sweep of the whole map was typically done with pitch only or with the region name spoken along with the pitch. P3 said “automatic sweep will be my first step to get acquainted with a new map to get the big picture” During the post-test free exploration of an unknown map, P3 swept the map several times in pitch-only to obtain a rough idea of where the highly populated regions were before starting to explore. P2 swept the unknown map once and accurately reported that most highly populated regions were in the west, by judging from the pitches and the sound panning positions. P2 was the only one that consciously used stereo panning cues in tasks. Most subjects said it was not difficult to understand how the sweep was done, but they need to know what the map looked like to make sense of it. Once they broke the whole map into nine smaller ranges and swept each range using the keypad, it made more sense. All subjects, except P7, were able to easily tell if a variable has a given geographical distribution pattern, by sweeping the nine ranges in pitch only. Unexpectedly, map sweep was also frequently used by all subjects to locate a region on the map. This was typically done with the region names spoken, and often combined with the arrow key navigation and 9- range sweep. It was also used to check what regions have been selected.

Navigate

Navigating the table was easy. All subjects mostly used cell mode because “it allows finer control of what to play”. The row and column mode was used by some subjects to sweep a column for the correlation and close-up comparison tasks.

All subjects reported that overall it was very easy to navigate the map. The 3x3 exploration was frequently used by all subjects except P2 who mainly used arrow keys to navigate and used sound panning to judge region locations. All subjects understood the mapping between map locations and the 3x3 layout of the keypad. They were able to use the 3x3 exploration to find the map location of a specific region, and to find what regions are in each map part. While subjects mostly looked for a region by navigating the table (typically by first ordering it alphabetically), sometimes they used the map. They often first used the 9 numeric keys to find out which range contains that region, then used arrow keys to move to that region. The 3x3 exploration also allowed some subjects to acquire knowledge about the overall map shape and the region layouts. During the study, P3, P5, P6, and P7 reported the overall map shape and region density distribution. P7 also used two-level recursive 3x3 exploration to find the county layout in the central and eastern parts of northern Maryland.

Subjects seemed to be able to zoom into/out of the 9 map ranges and stay aware of their zooming positions. Many subjects played with zooming extensively in training but did not use it in the test. Their explanation was there was no need from the tasks and the Maryland map only has 24 counties. If there was a need to focus on one area on a much bigger map, zooming could be helpful. Many subjects expressed the concern that zooming may become mentally intensive as

the map scale grows. One observed problem with 3x3 zooming is that some adjacent regions are assigned to different ranges and thus not reachable from each other after zooming. The problem can be remedied by allowing zooming centered on a region of interest.

Arrow key navigation was essential to find a region's geographical neighbors and was used by all subjects in adjacency tasks. It was also used often to explore regions in a small map range, typically identified earlier with the 3x3 exploration. While P2 mostly used arrow keys to navigate the map, most subjects were inclined to use the 3x3 exploration because it gave the absolute region locations.

“Arrow key navigation takes me everywhere on the map. It is not efficient especially when I am not familiar with the map”. “The nine keys tell me what are in the northwest and so on. It narrows me down to a specific range”.

To address the irreversibility problem in arrow key navigation, iSonic supports previous/next navigation to let users go through every region once and only once, following their order in the map sweep. Although a few subjects mentioned the irreversibility problem, they thought it was a natural fact about maps and no one seemed to be bothered. No one used previous/next navigation after the training because “there is no need for it” or “it does not make sense on maps”.

Situate

Subjects used situating to get the table sorting status, the current table position, the current map and map position, and the number of selected regions. Many subjects reset the interface before each task and did not use situating much since they remembered what they had done. However, all subjects considered this function essential so they do not need to redo the work “after a bathroom visit”.

Select, Search and Filter

All subjects were able to use selection and switch their focus between “all regions” and “selected regions only”, even across the two data views. Some subjects requested the ability to select variables besides selecting regions. Subjects also requested first-letter searching of regions. Filtering was not requested since the data sets are small.

6.6. DISCUSSION

It is clear that iSonic enabled subjects to find facts and discover data trends within georeferenced data, even in unfamiliar geographical contexts. The design choices in iSonic were overall easy to use and allowed subjects to effectively explore data in the map and table views without special interaction devices.

The studies do have limitations. The subjects might have made favorable comments because they wanted to please the experimenters. An average of 6 hours' use was not enough to go beyond the novice usage stage. Investigation of the tool's long-term use in real work circumstances will provide further understanding. We only tested users without any residual vision. Further studies with partially sighted users may reveal different usage patterns and visual- auditory interactions that may modify our results and framework.

However, the studies provided clear evidence that the blind users dramatically benefited from the set of task-oriented user actions (AISA) and the use of multiple highly coordinated data views offered by the Action-by-Design-Component framework. Some widely used visualization

techniques, such as the visual information seeking mantra and the use of multiple highly coordinated visualizations, with appropriate adaptation, work in the auditory mode. The key conclusions and design implications were:

- (1) All subjects were capable of choosing and switching between highly coordinated table and map auditory views, in order to complete the tasks. We believe users could also deal with more and different views such as graphs.
- (2) Using musical pitches to present numerical data makes it easier to perceive data trends in data series and enhances close-up value pair comparison. The integrated use of musical sounds and speech allows users to listen to overall trends and to get details.
- (3) A single auditory feedback detail level is not sufficient. Our 4 levels were all used productively. While it is hard to understand a data element without the appropriate context, too much detail slows down the sequential presentation and can be overwhelming for gaining the big picture. Designers need to carefully select multiple information levels and let users adjust it to fit their tasks.
- (4) A rapid auditory gist is valuable in conveying overall data trends and guiding exploration. For maps, perceiving spatial relation from a sequence of sounds can be difficult, but sweeping the map as separate smaller ranges in a consistent order was effective.
- (5) Navigation structures should reflect the data relation presented by the data view. In the map, designers would do well to provide 3x3 exploration using the numeric keypad and adjacency navigation using arrow keys. Users benefited from absolute localization and relative movements. Even a coarse map partitioning mapped to the physical spatial layout of a numeric keypad can provide valuable geographical knowledge. Stereo sound panning can be helpful but seems to be secondary in giving location cues for most subjects.
- (6) Selecting was valuable for all subjects in focused data examination. They were able to operate selection within and across data views and accomplish brushing.

7 EXTENSION TO GRAPHS

To examine the generalizability of the ADC framework, we applied it to other graphical data views. Both a line graph and scatterplot are considered. In the visual mode, a line graph clearly shows the value change for one or more variables over the period defined by another variable (typically time stamps). The table view in iSonic, with one variable sorted and a second variable swept automatically, produces the effect of a single variable line graph sonification in which X presents the sorted variable and Y presents the second variable. Details on interactive line graph sonification are available in the leading author's Ph.D. dissertation [39]. In this paper, we focus on the sonification of a bivariate scatterplot and its integration in iSonic.

Flowers and colleagues [7] found that a sonified overview of a 50 data sample bivariate scatterplot was quite as efficient as the visual equivalence in conveying sign and magnitude of correlation. The sonification was produced by a dot scanning method, in which a vertical scan line moved along the X axis at a constant speed and the Y value of each data sample (dot) was presented by a pitch when encountered by the scan line.

While information on individual dots can be interesting, sonifying every single dot could become overwhelming for scatterplots containing a large number of data items or many overlapping data items. To sonify a scatterplot, we used a dot aggregation method, in which dots are aggregated

into cells in a 2-D heat map. In Figure 3, the 2-D space of the scatterplot is equally divided into 9x9 grid cells. The choice of 9x9 resolution is guided by the 3x3 number pad exploration design. For each cell, the number of dots it contains is the value shown in the corresponding heat map. This spatial clustering and binning method is also fundamental in calculating “entropy”, a measurement strongly advocated by MacEachren et al. [16] for analyzing 2-D scatterplots as opposed to correlational coefficient.

Guided by the ADC framework, it was easy to systematically design the scatterplot sonification. Many interface designs for choropleth maps can be directly applied to the scatterplot heat maps, such as the sound-encoding schemes, spatial sweep, 3x3 recursive 2-D space exploration, and arrow key navigation, filtering, and so on. In terms of the actual implementation, much code from the map view can be reused in the scatterplot view without much modification, due to the high similarity of their sonification and interaction methods. One main effort lies in the design and implementation of data structures to generate, store, and maintain dot coordinates and semantically zoomable heat map representations of the scatterplot. The other part involves visually rendering the scatterplot. Below is a brief description of the scatterplot view in iSonic and its integration with the table and map views.

Gist

A gist is produced by spatially sweeping the grid cells. Empty grid cells can be included or omitted. When empty cells are included, users can expect the same gist length, sweep path and speed for all scatterplots using the same grid resolution. Each empty cell can either add a pause or a background sound to the gist.

Navigate

A heat map can be navigated in a relative or an absolute style, like in a choropleth map. Figure 3 (b) illustrates a 3x3 partition and the absolute navigation by activating a sub -gist of each of the 9 ranges. Users can zoom into a range or a grid cell for in-depth examination. Upon zooming, the heat map disaggregates. Figure 4 (a) shows zooming into range 6 of the scatterplot in Figure 3, and Figure 4 (b) shows the disaggregated heat map. The relative and absolute navigation methods can be recursively applied to the zoomed range.

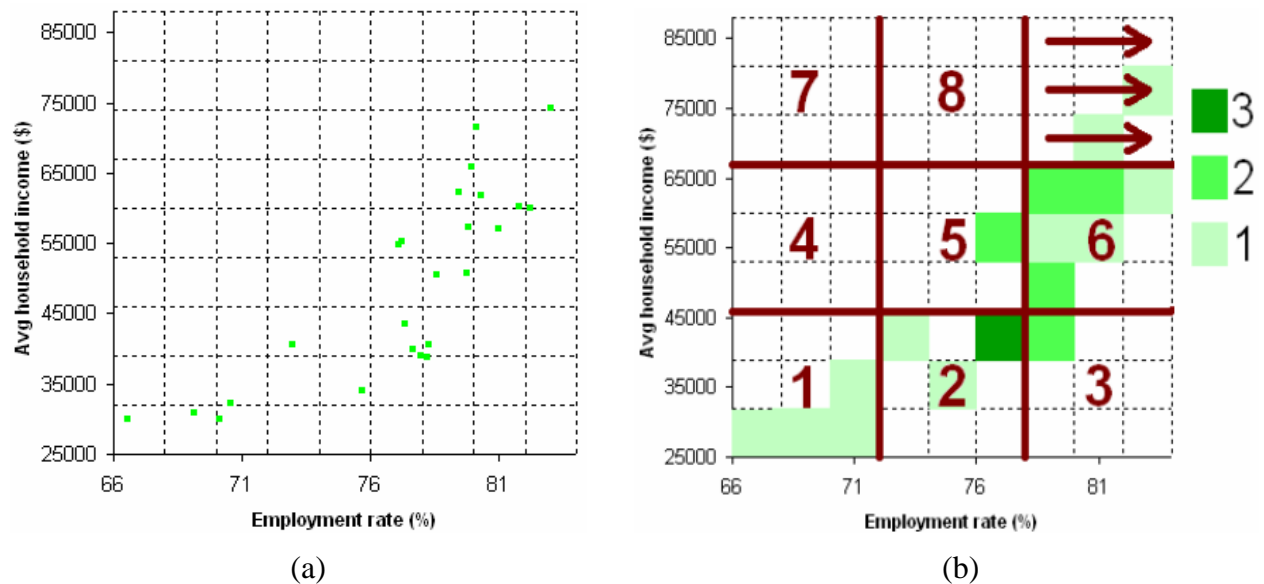


Figure 3: A bivariate scatterplot and its heat map at the resolution of 9x9 grid. The color indicates the number of dots in each grid cell. A darker green indicates a higher value.

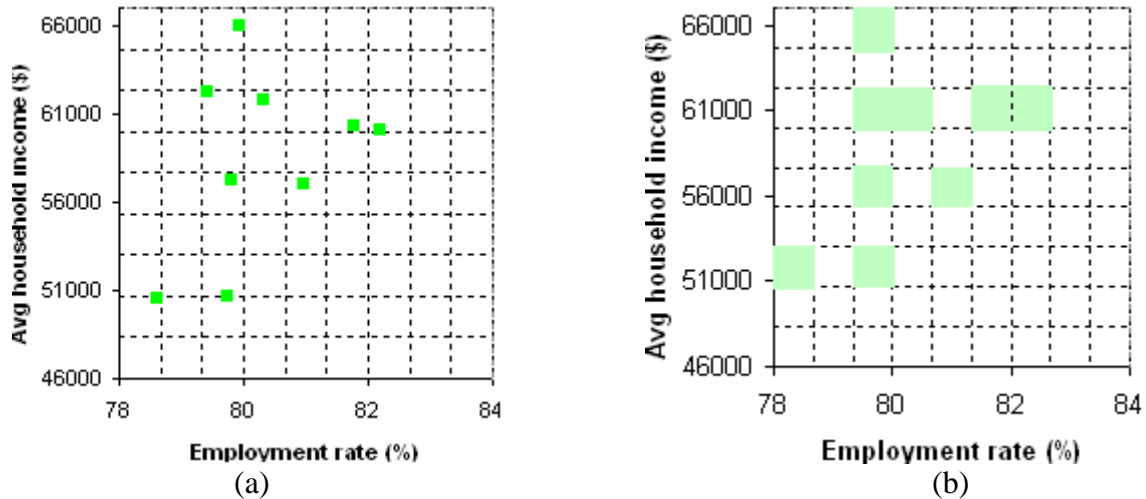


Figure 4: 3x3 zooming: (a) zoom into range 6 of the scatterplot in **Figure 3** (a) the 9x9 heat map for the partial scatterplot in (a).

Information level and details-on-demand

Four information levels are provided. At level 0, the grid cell position is spoken. At level 1, the heat pitch is played. At level 2, the heat value is also spoken. At level 3, the IDs (the region names) of the dots in the cell are spoken as well. When users press ‘space’ to request details, the X and Y value ranges of the current cell are also given besides the information given at level 3.

Select

Selection is applied to individual grid cells in the heat map, and affects all dots inside the operated cell. Users can switch among viewing “all regions” or “selected regions only”. In the later status, unselected dots fade into gray, and a new heat map is calculated to count only selected dots. This is useful especially when users select some regions in the table or map views and want to examine only those regions in the scatterplot view (see *Brush and coordination with other data views* below). It is also usually possible to select a single dot in a multi-dot cell, by zooming into that cell to get a more scattered view. The dot aggregation sonification method could be combined with the dot scanning method as two sonification modes.

Situate and filter

When users press ‘I’, iSonic speaks the name of the X and Y variables, the X and Y range of the current view, the number of dots in the view, the current navigation position, and so on. Filter is done in a similar way as for the map and table views.

Brush and coordination with other data views

The scatterplot view is highly coordinated with the table and map views during all AISAs, such as selection, navigation, filtering, and adjusting information level. Users can change the X and Y variables of the scatterplot view by selecting the desired variables in the table view. Regions selected in the table or map view are automatically selected in the scatterplot view, and vice versa.

8 CONCLUSION AND FUTURE WORK

We described an Action-by-Design-Component framework for designing auditory interfaces for analytical data exploration. We applied the framework to geo-referenced data and built a data exploration tool iSonic. Evaluation of iSonic with 7 blind users showed that the rich set of task-oriented actions (AISAs) and the use of multiple coordinated data views offered by the ADC framework were effective and beneficial to blind users, to accomplish complex tasks. The application to scatterplots demonstrates the generalizability of the framework which can lead to an auditory workspace with multiple highly coordinated graphical data views and a set of unified user actions for general analytical data exploration. We hope to encourage designers and researchers to apply the ADC framework in their auditory graph investigation, and in turn, to further validate and refine the framework. The similarity among the sonification and interaction designs of choropleth map and scatterplot indicate that some design methods may be effective across various graphs. Identifying what methods are commonly effective and what methods only work for specific graphs will help to extract rules to be used in the automation of interactive sonification for any data sets.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under Grant No. EIA 0129978 and ITR/AITS 0205271. We thank our subjects for participating and thank Alex Silver for helping to run the studies.

REFERENCES

- [1] Bennett, D.J., Edwards, A.D.N., Exploration of non- seen diagrams, Proc. Intl. Conf. on Auditory Display (ICAD), (1998).
- [2] Brown, L., Brewster, S.A., Ramloll, R., Burton, M., Riedel, B., Design guidelines for audio presentation of graphs and tables, Proc. ICAD (2003).
- [3] Brewster, S.A., Using nonspeech sounds to provide navigation cues, ACM Trans. on CHI, 5, 3 (1998), 224-259.
- [4] Card, S., Mackinlay, J.D., Robertson, G., The design space of input devices, Proc. ACM CHI (1990)
- [5] Ceaparu. I., and Shneiderman, B. Finding Governmental Statistical Data on the Web: Three Empirical Studies of the FedStats Topics Page, Journal of the American Society for Information Science and Technology, 55, 11 (2004) 1008-1015.
- [6] Donker, H., Klante, P., Gorny, P., The design of auditory user interfaces for blind users, Proc. NordiCHI (2002).
- [7] Flowers, J.H., Buhman, D.C., Turnage, K.D. Cross- modal equivalence of visual and auditory scatterplots for exploring bivariate data samples. Human Factors, 39, 3 (1997), 340-350.
- [8] Franklin, K.M., Roberts, J.C., Pie chart sonification, Proc. IEEE Information Visualization (2003), 4-9.

- [9] Handel, S. Listening: An Introduction to the Perception of Auditory Events. MIT Press (1989).
- [10] Hunt, A., Hermann, T., Pauletto, S., Interacting with sonification systems: closing the loop, Proc. IEEE Information Visualization (2004).
- [11] Kamel, H.M., Landay, J.A., The integrated communication 2 draw (IC2D): a drawing program for the visually impaired. Proc. ACM CHI (1999).
- [12] Karshmer, A. and Gillian, D. Math Readers for Blind Students: Errors, Frustrations, and the Need for a Better Technique. Proceedings of the 2005 International Conference on Human-Computer Interaction (HCII) (2005) [on CD-ROM].
- [13] Kramer, G., Walker, B., Bonebright, T., Cook, P., Flowers, J., Miner, N., Neuhoff, J., Sonification Report: Status of the Field and Research Agenda (1997) <http://www.icad.org/websiteV2.0/References/nsf.html>
- [14] Landau, S., and Gourgey, K., Development of a talking tactile tablet, Information Technology and Disabilities, VII(2), 2001.
- [15] Lazar, J., Allen, A., Kleinman, J., and Malarkey, C. What frustrates screen reader users on the web: A study of 100 blind users. International Journal of Human-Computer Interaction (2006, in press).
- [16] MacEachren, A.M., Dai, X., Hardisty, F., Guo, D., and Lengerich, G., “Exploring High-D Spaces with Multiform Matrices and Small Multiples”, *Proceedings of the IEEE Symposium on Information Visualization*, pp. 31-38, Washington, D.C.: IEEE Computer Society, 2003
- [17] Mynatt, E.D., Weber, G., Nonvisual presentation of graphical user interfaces: Contrasting two approaches, Proc. ACM CHI (1994).
- [18] National Federation of the Blind. Braille Usage: Perspectives of Legally Blind Adults and Policy Implications for School Administrators, available at: <http://www.nfb.org/brusage.htm> (accessed on 5/3/2006).
- [19] North, C., Shneiderman, B., Snap-Together Visualization: A user interface for coordinating visualizations via relational schemata, Proc. Advanced Visual Interfaces (2000).
- [20] Pauletto, S., Hunt, A., A toolkit for interactive sonification, Proc. ICAD (2004).
- [21] Ramloll, R., Yu, W., Riedel, B., Brewster, S.A., Using non-speech sounds to improve access to 2D tabular numerical information for visually impaired users. Proc. BCS IHM-HCI (2001).
- [22] Reenskaug, T., “The Model-View-Controller (MVC): its past and present”, available at http://heim.ifi.uio.no/~trygver/2003/javazone-jaoo/MVC_pattern.pdf, (accessed on 01/12/2006)
- [23] Roth, S. F., Chuah, M. C., Kerpedjiev, S., Kolojejchick, J. A., Lucas, P., Towards an information visualization workspace: combining multiple means of expression, Human-Computer Interaction, 12, 1-2 (1997), 131-185.
- [24] Saue, S., A model for interaction in exploratory sonification displays, Proc. ICAD (2000).

- [25] Shneiderman, B., The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations, Proc. IEEE Symposium on Visual Language (1996).
- [26] Shneiderman, B., Plaisant, C., Designing the User Interface: Strategies for Effective Human-Computer Interaction, 4th Edition, Addison Wesley (2005).
- [27] Sikora, C.A., Roberts, L.A., Murray, L., Musical vs. Real world feedback signals, Proc. ACM CHI (1995).
- [28] Slatin, J. and Rush, S. Maximum Accessibility. New York: Addison-Wesley (2003).
- [29] Sun Microsystems, Java Web Start Technology, <http://java.sun.com/products/javawebstart/>, (accessed on 09/06/2003)
- [30] W3C, Web Content Accessibility Guidelines, <http://www.w3.org/WAI/intro/wcag.php>
- [31] Walker, B.N., Lane, D.M., Psychophysical scaling of sonification mappings: a comparison of visually impaired and sighted listeners, Proc. ICAD (2001).
- [32] Walker, B.N., Cothran, J.T., Sonification sandbox: a graphical toolkit for auditory graphs, Proc. ICAD (2003).
- [33] Wall, S., and Brewster, S. Feeling what you hear: Tactile feedback for navigation of audio graphs. Proceedings of the ACM CHI (Computer-Human Interaction) Conference, (2006), 1123-1132.
- [34] Wenzel, E.M., Arruda, M., Kistler, D.J., Wightman, F.L., Localization using nonindividualized head-related transfer functions, Journal of the Acoustical Society of America, 94, 1 (1993),111-123.
- [35] Willuhn, D., Schulz, C., Knoth-Weber, L., Feger, S., SAILLET, Y., Developing accessible software for data visualization, IBM Systems Journal, 42, 4 (2003). <http://www.research.ibm.com/journal/sj/424/willuhn.html>
- [36] Zajicek, M., Powell, C., and Reeves, C. A web navigation tool for the blind, Proceedings of the ACM ASSETS conference (1998), 204-206.
- [37] Zhao, H., Plaisant, C., Shneiderman, B., and Duraiswami, R., “Sonification of geo-referenced data for auditory information seeking: design principle and pilot study”, Proceedings of the International Conference on Auditory Display (ICAD), Sydney, Australia, International Community for Auditory Display, July 6-10, 2004.
- [38] Zhao, H., Smith, B.K., Norman, K., Plaisant, C., and Shneiderman, B., “Listening to maps: user evaluation of multiple designs of interactive geo-referenced data sonification”, *IEEE Multimedia Special Issue on Interactive Sonification* 12, 2 (2005) 26 -35.
- [39] Zhao, H., “Interactive sonification of abstract data – framework, design space, evaluation, and user tool”, Ph.D. Dissertation, University of Maryland, College Park, 2006. Also available as Human-Computer Interaction Lab. Technical Report HCIL-2006-06.

Relationship to previous work

This paper contains the summary of the doctoral dissertation [1] conducted by Haixia Zhao under the direction of Catherine Plaisant and Ben Shneiderman, with Jonathan Lazar's participation in the final user study with 7 blind users. None of the previous publications covers this substantial final user study.

The 2003 Digital Government Research Conference paper [2] covered several ideas for improving accessibility for choropleth maps, but did not describe the interface, experiments or framework.

A paper presented in the ICAD 2004 conference [3] described the first notions of the framework (40%) for analyzing sonification interfaces and described a very limited prototype of the interface.

The IEEE Multimedia paper described the initial experiments with sighted users based on an early prototype [4].

During the development of the dissertation, Haixia Zhao presented two demonstrations at conferences and participated in the CHI 2005 and ASSETS 2005 Doctoral Consortia which resulted in two-page summaries of the work in progress.

[1] Zhao, H., *Interactive Sonification of Abstract Data - Framework, Design Space, Evaluation, and User Tool*, Ph.D. Dissertation.

[2] Zhao, H., Plaisant, C., Shneiderman, B., Improving Accessibility and Usability of Geo-referenced Statistical Data, *Proc. of the Digital Government Research Conference* (2003) 147-150

[3] Zhao, H, Plaisant, C., Shneiderman, B, Duraiswami, R., Sonification of Geo-Referenced Data for Auditory Information Seeking: Design Principle and Pilot Study , *Proc. of International Conference of Auditory Displays* (2004)

[4] Zhao, H., Smith, B.K., Norman, K., Plaisant, C., Shneiderman, B., Interactive Sonification of Choropleth Maps: Design and Evaluation, *IEEE Multimedia Special Issue on Interactive Sonification* 12, 2 (2005) 26 -35.