Interactive Sonification of Choropleth Maps

Haixia Zhao, Benjamin K. Smith, Kent Norman, Catherine Plaisant, and Ben Shneiderman
University of Maryland

Interactive sonification systems can make georeferenced data accessible to people with vision impairments. The authors compare methods for using sound to encode georeferenced data patterns and for navigating maps.

Auditory information is an important channel for the visually impaired. Effective *sonification* (the use of nonspeech audio to convey information) promotes equal working opportunities for people with vision impairments by helping them explore data collections for problem solving and decision making. For example, the data collected by the US Census are an important source of information for government, industry, and the general public. Georeferenced data are often presented visually using *choropleth* maps—maps with predefined regions colored to show how a variable differs between regions. To access georeferenced data, visually impaired users currently rely on screen readers to linearly speak the geographic region names and data that are presented as table records, often in alphabetical order. Examples include FedStats, the US government statistical data gateway (http://www.fedstats.gov) and Corda, which automatically converts maps and graphs to descriptive text (http://www.corda.com). Such linear textual presentation makes it difficult for visually impaired users to locate specific data and understand data patterns in geographical context. There are many ways to improve visually impaired users' access to such data collections. Ramloll et al.[1] found that using nonspeech sound in 2D numerical tables decreased subjective workload and enhanced data comprehension.

In our effort to solve this problem, we worked with a blind design partner and developed several sonifications. The sonifications have synchronized visual and auditory presentations and follow two design guidelines:

- conform to an auditory information-seeking principle (AISP)[2] and
- have minimum requirements for special software and hardware, making them easily accessible to the general public and helping to improve universal accessibility.[3]

Georeferenced data analysis often involves geographical context information. In the visual mode, a picture is often said to be worth a thousand words. A glance at the geographic distribution pattern of data often gives users valuable information. Our goal is to achieve a similar effect in the auditory mode. The geographical distribution pattern of georeferenced data involves three dimensions, two for the geographical location of a data point on a map, and a third for the numerical value. (See the "Related Work" sidebar for a discussion of work in this area.)

In a pilot user study,[2] we compared a preliminary map-based design to an enhanced table design. The study showed that subjects could perceive georeferenced data distribution patterns on a real map with 51 geographic regions using both designs. Observations and user comments indicate that AISP fits users' pattern-recognition strategies. Based on observations from the pilot study, we designed a user study to compare two map-navigation methods and investigated the effect of using sound to encode vertical geographic positions. The study was part of our continuous effort to identify effective sonification designs for georeferenced data and to gain insights into people's ability to perceive patterns in sonified data.

Auditory information-seeking principle

Inspired by Shneiderman's visual-information-seeking mantra,[4] we proposed an AISP[2] consisting of four elements:

- *Gist.* A gist is a short auditory message presenting the overall trend or pattern of a data collection. It guides further explorations and often lets users detect anomalies and outliers. Because humans perceive sounds as transient time-sensitive stimuli, the gist must be short

Published by the IEEE Computer Society

and allow active attention/rehearsal to transform it from short-term memory to working memory. Furthermore, human auditory perception is less synoptic than visual perception. The gist must present multiple data items serially instead of simultaneously (temporization). Because of the short length and low parallelism, the system must use data aggregation when the data collection is large, for example, more than 100 data items.

- *Navigate*. Navigating refers to the user flying through the data collection, selecting and listening to portions of it. Navigation is an iterative process of the user initiating an action and the system giving feedback about the user's current range of interests. Because sound is transient feedback, users must tie sound to virtual objects and construct mental navigation maps to interact with the data items through the auditory interfaces. The input method design must use input devices suitable for visually impaired users. For example, the traditional point-and-click method via a computer mouse works well for normal-sighted users but is difficult for users with vision impairments.

- *Filter*. Filtering out unwanted data items helps trim a large data collection to a manipulable size and lets the user quickly focus on the items of interest.

- *Details on demand*. Users can select an item or group for further details. Although sonification emphasizes the use of nonspeech sound, speech can be an effective presentation at this level.

Related Work

Previous work has shown that users can interpret a quick sonified overview of 2D line graphs containing a single data series,[1] two data series,[2,3] and bivariate scatterplots.[4] Research has also shown that users can recognize 2D graphical shapes by listening to sounds tracing the shapes' borders.[5]

Few observations exist about the ability to recognize data distribution patterns with more than two dimensions in the auditory mode. Meijer[6] aims to let visually impaired users "see" with hearing. This approach's effectiveness still needs to be established. Wang and Ben-Arie[7] found that people can recognize simple shapes on binary images of 9×13 resolution in which the pixels are raster-scanned slowly. Jeong[8] showed that people can locate the minimum/maximum values on a simplified choropleth map with up to nine regions, with the values presented as different sound volumes.

References

1. J.H. Flowers and T.A. Hauer, "Musical versus Visual Graphs: Cross-Modal Equivalence in Perception of Time Series Data," *Human Factors*, vol. 37, no. 3, 1995, pp. 553-569.
2. T.L. Bonebright et al., "Testing the Effectiveness of Sonified Graphs for Education: A Programmatic Research Project," *Proc. Int'l Conf. Auditory Display*, Helsinki Univ. of Technology, 2001, pp. 62-66.
3. L. Brown and S.A. Brewster, "Drawing by Ear: Interpreting Sonified Line Graphs," *Proc. Int'l Conf. Auditory Display*, Boston Univ., 2003, pp. 152-156.
4. J.H. Flowers, D.C. Buhman, and K.D. Turnage, "Cross-Modal Equivalence of Visual and Auditory Scatterplots for Exploring Bivariate Data Samples," *Human Factors*, vol. 39, no. 3, 1997, pp. 341-351.
5. J.L. Alty and D. Rigas, "Communicating Graphical Information to Blind Users Using Music: The Role of Context," *Proc. ACM SIGCHI Conf. Human Factors in Computing Systems*, ACM Press, 1998, pp. 574-581.
6. P.B.L. Meijer, "An Experimental System for Auditory Image Representations," *IEEE Trans. Biomedical Eng.*, vol. 39, no. 2, Feb. 1992, pp. 112-121; also available as the Voice system at http://www.seeingwithsound.com.
7. Z. Wang and J. Ben-Arie, "Conveying Visual Information with Spatial Auditory Patterns," *IEEE Trans. Speech and Auditory Processing*, vol. 4, 1996, pp. 446-455.
8. W. Jeong, *Adding Haptic and Auditory Display to Visual Geographic Information*, PhD thesis, Florida State Univ. 2001.

Pilot user study

We used the AISP to guide the design of georeferenced statistical data sonifications and conducted a pilot user study of these sonifications. The pilot study was a within-subjects experiment that investigated whether users could perceive geographic distribution patterns of a five-category data collection on a 51-region US state map.[2] The study had three purposes:

- check the feasibility of using sonification to present data referring to real maps (not grids or simplified maps),
- investigate the AISP's validity, and
- obtain early user observations to guide future designs.

We paid nine sighted subjects to participate in the study. Each subject used both a spatial map design and a table design enhanced with geographic location knowledge.

In the spatial map design, we tied head-related transfer function (HRTF) spatial sounds to a US state map to create the effect of a virtual map surrounding the user at the center, as Figure 1, next page, illustrates. We based the spatial sound on the widely used KEMAR mannequin HRTF from the CIPIC HRTF database.[5] For each US state, we played

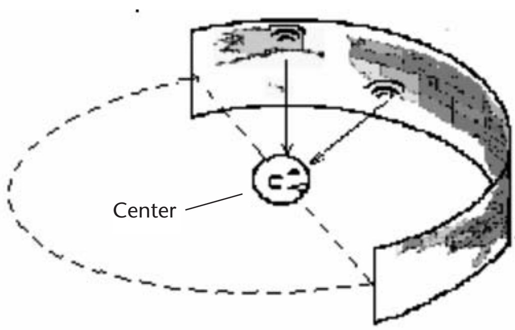

Figure 1. Spatial sound tied to the map creates the effect of a virtual half-cylinder-shaped map surrounding the user, who is located at the center. The illustration doesn't reflect the real spatial parameters.

Figure 2. Sweeping order (in both user studies).

a pitch of a string instrument sound for 200 milliseconds (ms) to indicate its georeferenced datum, then a 100-ms piano pitch to indicate its vertical position. The piano pitch supplemented the low vertical localization accuracy of nonindividual HRTF spatial sounds.[6] Although research has shown that using individual HRTFs can improve the vertical localization accuracy, we chose not to use this method because individual HRTFs are unlikely to be available for general public users.

The five value pitches were from an increasing scale of CEGCE starting from the middle C on a piano keyboard. A higher pitch indicated a higher value. The vertical position pitches ranged from about one octave below middle C to about two octaves above. A higher pitch indicated a state to the north.

Using a keyboard, users could start an automatic spatial sweep from west to east (as Figure 2 illustrates) to listen to a 25-second gist of the data of all the states, navigate the map to explore individual states, and request the spoken detail of individual states. During a sweeping, a bell sound indicated the end of a sweep column, and three consecutive bell sounds indicated the sweeping's end. A bell sound also played when the navigation automatically jumped to a different sweep column.

In the enhanced table, we ordered the states according to their occurrence in the spatial map sweeping. Users could start an automatic sweeping following the table order, navigate the states following the sweep order, and request state details. For each state, we played a pitch of the same string instrument sound for 200 ms to indicate its georeferenced datum, speaking the state name at the same time. All sounds came from the center.

The study showed that subjects could perceive the general pattern type after listening to a 25-second gist just once (the overall accuracy was 56 percent for both the map and the table). After exploring for about 110 seconds, the general pattern-type recognition accuracy increased to 78 percent for the table and 89 percent for the map. Subjects could also grasp the details of the patterns with 67-percent accuracy for the table and 75-percent accuracy for the map. The subjects strongly preferred the map design to the table, although no statistically significant difference was found in terms of performance. The study also indicated that AISP conforms to subjects' information-seeking strategies. Our experience with our blind design partner suggested that we would obtain similar results from visually impaired users. More details about the study are available elsewhere.[2]

Map-based interface designs

Observations and subjects' comments in the pilot study suggested that the map-based design could be significantly improved. Two observations led to the present user study's design:

- The sound indicating the state's vertical position distracted some subjects from the value.
- Irregular state shapes and sizes make defining a good state-by-state navigation matrix difficult. It often causes the actual navigation direction to drift away from the subjects' expectations, possibly causing misinterpretation.

Experimental design

The user study was a $2 \times 2 \times 6$ experiment with two between-subjects factors and a within-subjects factor. We assigned each subject to one of four interfaces, which were defined by the first two factors. The third factor was the six map patterns, each used by all subjects.

The first factor consisted of two treatments: the presence or absence of a vertical position sound (VPS). In the treatment without the VPS, for each state, we played only a 200-ms string pitch to indicate the state's georeferenced data value. In the treatment with the VPS, for each state, a 100-ms piano pitch followed the value pitch to indicate the state's vertical position. The ranges of the value pitches and VPS pitches were the same as those in the first user study.

The second factor was the navigation method. We used two navigation methods: the column interface and the mosaic interface. The column interface used state-by-state navigation (see Figure 3a) similar to the navigation method in the first user study. In this study, however, the navigation didn't automatically jump between adjacent sweep columns when it reached the top or bottom of a column, and we remapped the control keys using only the keys in the number pad and the four arrow keys on a standard keyboard to control the interface. Table 1 lists the available controls.

To generate the map for the mosaic navigation, we lay an even-sized grid on a regular map. We ordered all of the states by size, starting with the smallest state. We then assigned each state the cell closest to the state's center, to ensure that even small states got one cell. Next, we assigned each unassigned cell to the state with the most pixels in the cell. Finally, we manually adjusted the automatic assignment to smooth out the state borders.

Using the number pad, users navigated the cells in eight directions. When they moved from a cell to another cell in a different state, the system played that state's sounds. When the movements were within one state, the system played no sound. This lets users sense the state's size.

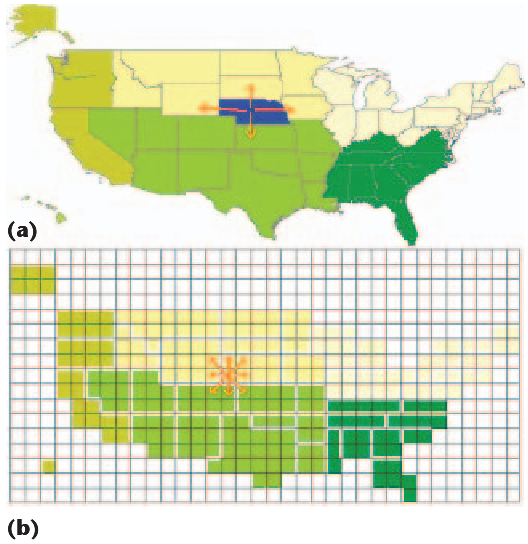

Figure 3. Navigation: (a) state-by-state column navigation and (b) cell-by-cell mosaic navigation.

Table 1: User key controls in the cell-by-cell mosaic navigation and the state-by-state column navigation.

Key	Mosaic Navigation	Column Navigation
0	Play subgist starting from current state	
1, 3, 7, 9	Go to closest valid cell (that is, a cell that is part of a state) diagonally adjacent to current cell and play sound*	N/A
2	Go to closest valid cell below current cell and play sound	Go to and play next state in the current sweep column
4	Go to closest valid cell to the left of current cell and play sound	Go to and play state in the previous sweep column that's nearest to current state
5	Request state name and value pitch of current state	
6	Go to closest valid cell to the right of current cell and play sound	Go to and play state in the next sweep column that's nearest to current state
8	Go to closest valid cell above current cell and play sound	Go to and play previous state in the current sweep column
Up, down arrow	Go to north- (south-)most valid cell and play sound	Go to and play north- (south-) most state in current sweep column
Left, right arrow	Go to west- (east-) most valid cell and play sound	Go to and play state in first (last) sweep column that's nearest to current state
Enter	Play gist of all states	
+	Request spoken value of current state	
Any key	Stop playing the gist, set current state to be the state just played in the gist	

**Play sound* means to play a percussion sound for each background cell crossed, then play the target state if it differs from the current state.

Figure 4. Pattern types: (a) vertical strip, (b) horizontal strip, (c) cluster, (d) diagonal strips, and (e) no pattern.

When a movement crossed the background (for example, an ocean, which had no data values) to get to a state, the system played a series of percussion sounds before playing the new state's sounds. The number of percussion sounds played equaled the number of background cells crossed. Again, the key controls were the number pad keys plus the four arrow keys (see Table 1).

Whenever users reached a border and could go no further in that direction, a synthesized female voice reminded them that they were at the boundary. During sweeping and navigation, the system synchronized the visual display with the auditory presentation to help subjects better understand the interfaces during the interface explanation. During the training task and experimental tasks, the display was hidden from the users and was visible only to the experimenters. In column navigation, the current state is highlighted in blue. In mosaic navigation, the current cell is always marked by a yellow dot. The dot also moves across white space when users cross a national border or ocean.

All four interfaces use the same sweeping order as in the first user study (see Figure 2). Because all subjects from the pilot user study reported that they couldn't tell the vertical positions of the generic HRTF spatial sounds used, we simply used the stereo panning (0 ~ 127) in the second study to indicate left-to-right sound positions. We changed the sound indicating the end of each sweep column and the end of a sweeping from a bell sound in the pilot study to a percussion sound in the second study.

Second user study

Forty-eight sighted subjects from introductory psychology courses participated in the study to earn extra credit. Their ages ranged from 18 to 51 with a median age of 20. Thirty-seven participants were female and 11 were male. All subjects reported using computers at least one hour per week, and 44 reported using them at least five hours per week. Fourteen subjects reported having had professional music training for a year or more. None of these factors were significantly correlated with performance on any outcome measure at the .05 level. We randomly assigned subjects to one of four interface conditions, and to one of six task-order conditions.

We first tested each subject on his or her ability to recognize instruments, pitches, and stereo panning, and on geographic knowledge. We then taught subjects how to use their assigned interface. (Two experimenters, neither of whom knew the focal hypothesis, ran the experiment. There were no significant differences in performance for subjects run by each experimenter.) The subjects learned the sound design and interface controls while viewing the display. We then hid the display from the subjects, and they practiced by performing a training task with a monotonically horizontal-strip pattern, following the same procedure as in the real test.

Each subject performed six pattern-matching tasks. Figures 4 and 5 illustrate the pattern types and patterns used. Two tasks used vertical-strip maps (one monotonic and another interleaving), two used diagonal-strip maps (one monotonic and another interleaving), and two used cluster maps. We counterbalanced the task order using a Latin square and set up the orders so no subject ever had the same general map type twice in a row. We notified subjects that we would measure both accuracy and speed, but that accuracy was more important.

The task procedure was similar to the first user study but with a few changes. Each task consisted of three steps. First, subjects listened to the gist of the data once and were asked whether they perceived any pattern in the data by choosing from the five pattern types shown in Figure 4.

Subjects also chose their confidence level about the answer based on a 10-percent-break scale. Second, subjects explored the map using the key controls in Table 1 for as long as they needed or up to three minutes. Subjects then again chose the pattern type and their confidence level. Third, experimenters presented four map patterns to the subjects, each with the same pattern type. One of the four map patterns was the actual pattern the users had been exploring, and thus the general pattern type was not necessarily the same as the pattern type they had chosen. Subjects then chose the matching pattern from the four visual patterns, as well as their confidence level.

Figure 6 shows a sample of four such visual pattern choices. We recorded all the subjects' keystrokes as well as the time they took to explore the maps. At the end of the six tasks, the experimenters gave the subjects a post-test questionnaire. The entire process took less than an hour per subject.

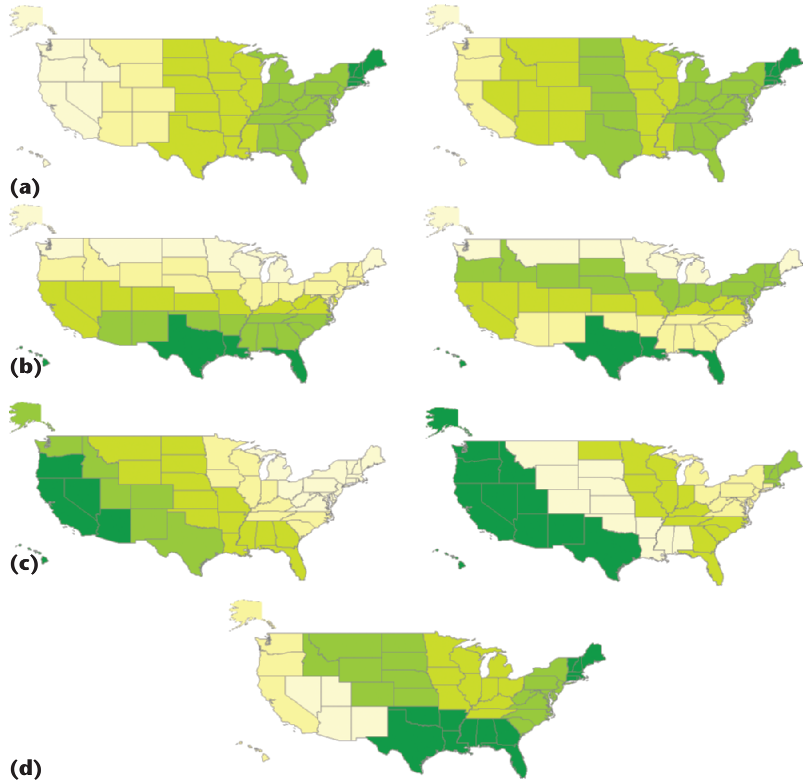

Figure 5. Sample patterns for each pattern type: (a) vertical, (b) horizontal, (c) diagonal, and (d) cluster.

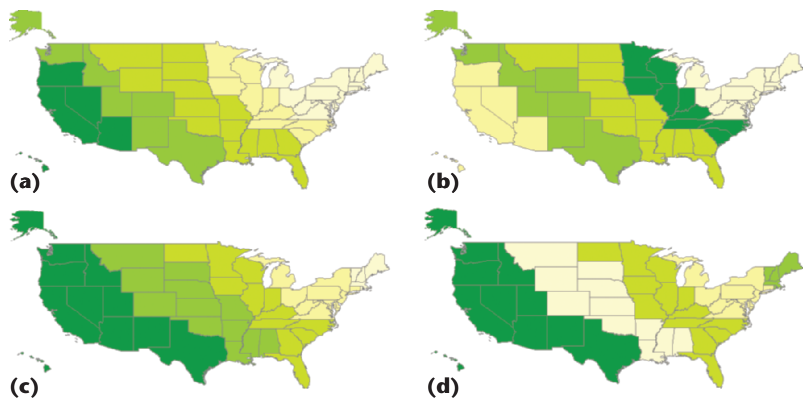

Figure 6. Sample of four visual pattern choices: (a) and (c) diagonal monotonic, (b) and (d) diagonal interleaving. A greener color indicates a higher statistical value. We categorized values into five ranges.

Results

Although overall the tasks were difficult, subjects were able to perceive five category-value distribution patterns on a 51-state real map. The average pattern-type recognition accuracy was 50.7 percent after the gist but before exploration (chance accuracy would be 20 percent because there were five choices). After exploration, pattern type recognition accuracy increased slightly to 55.2 percent, and the specific pattern recognition accuracy was 48.7 percent (with chance accuracy at 25 percent).

At least two likely explanations exist for the low accuracy in this study compared to that in the pilot study. First, the tasks in this study were more difficult. After exploration, subjects had to choose the general pattern type explicitly, rather than having their choice inferred from their selection of a specific pattern. Similarly, subjects had to choose a specific map from four similar choices of the same pattern type, rather than from a set of choices of different pattern types. Second, the differences between subject populations possibly accounted for some of the difference. In this study, we didn't pay subjects or give them performance incentives, as we did in the pilot study. Because the tasks were difficult, the subjects' motivation played an important role.

We found no statistically significant difference in performance across the four interfaces. However, we found that three correlational factors significantly affected subjects' performance:

- *Geographic knowledge*. Subjects' knowledge of US geography was positively correlated with performance on identifying both the general ($r = .31$, $p < .05$) and specific ($r = .36$, $p < .05$) patterns after the exploration period. (The letter r refers to the Pearson correlation coeffi-

Table 2. Average accuracy for identifying different pattern types.

Pattern Type	General Pattern Type After Gist	General Pattern Type After Exploration	Specific Map
Chance level	0.20	0.20	0.25
Horizontal*	0.44	0.42	0.42
Cluster 1	0.44	0.50	0.52
Cluster 2	0.54	0.50	0.52
Diagonal interleaving	0.50	0.44	0.29
Diagonal monotonic	0.50	0.52	0.46
Vertical interleaving	0.63	0.67	0.63
Vertical monotonic	0.44	0.69	0.50

*The horizontal pattern was the training task and was always first. The other three patterns (cluster, diagonal, and vertical) are averages of two tasks and were counterbalanced throughout the rest of the experiment.

cient, a measure of how related two variables are; p refers to the observed significance level.)

- *Pitch differentiation ability*. All subjects could distinguish between three pitches in our pretest, but seven subjects needed a second try to get the answer correct. These subjects generally did worse on all outcome measures and significantly worse on the general pattern type after exploration ($r = .45$, $p < .05$).

- *Task strategy*. The post-test questionnaire asked subjects to describe their strategies. We compared these strategies by identifying common words and phrases in the descriptions. Some strategies appear to have been more effective than others. Subjects who reported listening for changes did particularly well on the two outcome measures taken after the exploration period. (For the general pattern type after exploration, $r = .36$, $p < .05$; for the specific pattern type, $r = .31$, $p < .05$.) Subjects who reported trying to visualize the map did particularly well identifying the general pattern type after the gist ($r = .49$, $p < .05$). Subjects who reported paying attention to the piano sound (which indicated the states' vertical positions) did significantly worse identifying specific maps ($r = .30$, $p < .05$). The most common strategies were moving around to find particular states and visualizing the map.

The tasks varied in difficulty, and each had only a single pattern type, so we can't separate the difficulty of identifying patterns or maps from the difficulty of specific tasks. However, we can report that the easiest tasks were the vertical patterned maps and the most difficult map was diagonal. Surprisingly, the only task for which identification of the general pattern type improved significantly after the exploration period was the vertical monotonic map ($t(47) = 2.72$, $p < .05$). Table 2 lists the details. This suggests that first impressions were important.

The post-test questionnaire asked subjects to rate the difficulty of both the tasks and the interface. Subjects in all conditions found the task hard (2.8 on a seven-point Likert scale, with 1 as very difficult and 7 as very easy). On the same scale, column navigation users rated the interface 6.0, and mosaic users rated that interface 5.2. This difference was significant ($F(1, 44) = 5.23$, $p < .05$). (F refers to the Fisher test, similar to an analysis of variance.)

Two questions on the post-test questionnaire asked about the stereo sound. The first asked whether the stereo sound helped subjects locate states. Responses were just above neutral (4.5 on a seven-point Likert scale, where 1 was very distractive and 7 was very helpful). The second question asked whether the stereo sound helped subjects picture the data distribution. Subjects who didn't have the vertical position sound gave higher ratings (5.2) to this question than subjects who had the vertical position sound (4.4) ($F(1, 44) = 5.21$, $p < .05$).

Two questions asked whether the vertical position sound helped subjects locate states or picture the data distribution (we only asked subjects in the two conditions with the vertical position sound). The average responses were close to neutral: 3.8 for the first and 4.2 for the second, using the same scale as that for the stereo sound.

One question asked subjects how good a sense they had of where they were on the map. On a seven-point scale, subjects gave an average response of 4.6, which was just better than "some sense" (4), but considerably lower than "good sense" (7). Conditions for this question didn't differ significantly.

Finally, the questionnaire asked subjects about the sounds' tempo. The average response was 4.5 on a seven-point scale, on the "too fast" side of the "right tempo" (4).

Possible improvements

Because performance with the four interface conditions didn't differ significantly, we must draw conclusions cautiously. We can say, however, that the vertical position sound seems to have been unhelpful at best, and subjects who reportedly paid attention to it did worse than those who didn't. It also seems to have taken away from the stereo sound's utilization. Using a vertical position sound doesn't appear to have any advantage, and because such a sound also adds complexity, we don't intend to use it in future sonifications.

Subjects ranked the column navigation interface significantly easier than the mosaic navigation interface, although the interfaces didn't differ significantly in terms of performance. We don't know exactly what users preferred about the column navigation, but it was somewhat simpler to use, required fewer keystrokes, and gave feedback after every keystroke. This suggests a possible improvement to the mosaic navigation interface: providing sound feedback after every keystroke. The mosaic navigation interface used in the study played a state's sound only when the subject entered the state. It gave no feedback when the subject moved within the state unless the subject pressed the 5 or + keys. The mosaic navigation design was intended to give the subject some sense of state size and shape. We can improve the design by providing state value feedback on every keystroke and playing a special sound before the state value sound to indicate whether the subject has entered a new state.

Some of the difficulty subjects had in identifying patterns and maps might have been because of the experimental conditions, not the interface itself. Subjects in our experiments had a few minutes to learn the interface, and might have benefitted from more learning time. Visually impaired users would likely spend a considerable amount of time learning and using such systems, so more training time is reasonable. It might also be useful to suggest certain strategies to users during training, such as visualizing the map and listening for changes, as these strategies helped other users.

Observations of visually impaired users

Although we used sighted subjects in our two controlled experiments, we've been working closely with one congenitally blind student during the design process. We have also presented the designs to another blind user. Our experience with them shows that the pattern recognition tasks were difficult but still possible.

Our blind design partner was only able to identify vertical-strip patterns at first. After using a tactile map to learn the definition of the sweeping columns, he could describe in detail most of the patterns, including horizontal-strip patterns, monotonically diagonal patterns, and cluster patterns.

The other blind user works with statistical data and is familiar with US geography. In a brief session conducted at her workplace, she learned the interface and explored a monotonically horizontal-strip pattern using the column navigation without the vertical position sound. She was able to recognize the pattern on the first try.

These observations have led us to some reflections on the training process. Sighted subjects learned the geography and sweeping column definition during both the studies' training stage and their time spent choosing the matching patterns. Visually impaired users also need such learning, but in nonvisual ways. Depending on resource availability, nonvisual training can take place in several ways, such as using a tactile map or speaking the states' sweeping order.

Furthermore, some of the experiment's findings conform to the blind student's comments. For example, the blind student commented that he liked the column navigation better because it required fewer keystrokes than the mosaic navigation. We expect that the automatic sweeping's tempo can increase after prolonged use; however, the blind student thought the current tempo was appropriate and should not be faster. Sighted subjects reported the tempo was close to the "right tempo" but a little "too fast."

Future work

The user study revealed strengths and weaknesses of our sonification interfaces. The observations will help us improve both user training and interface design. Our future work takes two main directions.

First, we plan to replicate the studies with visually impaired users to compare with the observations obtained with sighted users. To do this, we must develop new outcome measures that don't depend on visual displays. One possible way is to let subjects choose from a set of tactile maps with tactile textures indicating the distribution patterns. However, special care must be taken because this approach introduces an extra factor—the tactile perception ability—into the measurements.

Second, we can improve the interface in many ways. For data-to-sound mapping, we can improve the temporization (sweeping order). The user study showed that the easiest patterns to recognize were those whose gradient changes conformed to the sweeping order. We'll investigate this relation further and provide multiple sweeping orders for exploring patterns. On the other hand, we expect that the irregular state shapes and sizes would cause considerable difficulty in defining a suitable sweeping order, on which we also based the column navigation. We already encountered this difficulty when we defined the sweeping columns for the US state map, especially when considering the eastern part of the map. For other maps with even more irregular region shapes and sizes, we might need to use other sweeping methods, such as a mosaic sweeping, instead of state-by-state sweeping.

We could improve the interaction design by using absolute pointing methods. To avoid needing special external devices, we based both current map navigation methods on a standard keyboard. Every keystroke causes a relative, incremental change to the exploration position on the map. The questionnaire showed that subjects had a fairly weak sense of where they were on the map during navigation. Many studies have shown that in both the real world and virtual environments, users combine motor (for example, vestibular) information with sensory (for example, visual) information to construct a mental spatial representation.[7] Zheng et al.[8] investigated how navigation devices and modes of operation affect users' ability to develop an accurate mental spatial representation of a virtual environment. They found that the absolute pointing mode was better than relative mode.

We can provide absolute pointing methods for map navigation in various ways, using either a standard keyboard or a special external device. For example, we can map the keyboard layout to a map, with each key representing a position on the map. Of course, the pointing resolution is limited. A touchpad calibrated to the full map range can provide continuous movement on the map. A tactile grid or map laid on top of the touchpad could further enhance it. Using a tactile map requires users to have access to special devices, such as tactile embossers, and to restrict the flexibility of switching maps (for example, zooming into a state on a state map brings up the state's county details, requiring a different tactile map). A generic tactile grid can be a middle-point solution that provides some tactile positioning cue while reducing the requirement of special devices. More details about our continuous work in map navigations by absolute pointing methods using a keyboard or touchpad are available elsewhere.[9]

Making choropleth maps accessible to visually impaired users is a difficult challenge. We've developed initial designs and put in place a testing procedure that lets us refine our designs, which we hope will lead to improved universal access to georeferenced statistical data. **MM**

Acknowledgments

This material is based on work supported in part by the US National Science Foundation under grants EIA 0129978 (see also http://ils.unc.edu/govstat/) and ITR/AITS 0205271, the US Census Bureau, and the US National Center for Health Statistics. We thank our blind design partner Ahmad Zaghal, and thank Dmitry N. Zotkin for providing the spatial sound technique. We thank Michael Kelehan and Cristina Fraschetti for helping to run the experiment, and Jeffrey K. Smith for helping to design it. We also thank Barbara Shinn-Cunningham and John Flowers for their comments and suggestions.

References

1. R. Ramloll et al., "Using Nonspeech Sounds to Improve Access to 2D Tabular Numerical Information for Visually Impaired Users," *Proc. 15th Ann. Conf. British Human–Computer Interaction Group*, Springer, 2001, pp. 515-530.
2. H. Zhao et al., "Sonification of Geo-Referenced Data for Auditory Information Seeking: Design Principle and Pilot Study," *Proc. Int'l Conf. Auditory Display,* Int'l Community for Auditory Display, 2004.
3. C. Plaisant, "Information Visualization and the Challenge of Universal Access," *Exploring Geovisualization*, to be published in J. Dykes, A. MacEachren, and M.-J. Kraak, eds., Elsevier, 2005.
4. B. Shneiderman, "The Eyes Have It: A Task by Data Type Taxonomy for Information Visualization," *Proc. IEEE Symp. Visual Languages*, IEEE CS Press, 1996, pp. 336-343.
5. V.R. Algazi et al., "The CIPIC HRTF Database," *Proc. 2001 IEEE Workshop on Applications of Signal Processing to Audio and Electroacoustics* (WASPAA 01), IEEE Press, 2001, pp. 99-102.
6. E.M. Wenzel et al., "Localization Using Nonindividualized Head-Related Transfer Functions," *J. Acoustical Soc. America*, vol. 94, no. 1, 1993, pp. 111-123.
7. D.J. Simons and R.F. Wang, "Perceiving Real-World Viewpoint Changes," *Psychological Science*, vol. 9, no. 4, pp. 315-320.
8. X.S. Zheng, G.W. McConkie, and B. Schaeffer,

"Navigational Control Effect on Representing Virtual Environments," *Proc. Human Factors and Ergonomics Soc. 47th Ann. Meeting*, CD-ROM, Human Factors and Ergonomics Soc., 2003; http://hfes.org/publications/2003proc.html.
9. H. Zhao, C. Plaisant, and B. Shneiderman, "I Hear the Pattern—Interactive Sonification of Geographical Data Patterns," to be published in *ACM SIGCHI Extended Abstracts,* ACM Press, 2005.

Haixia Zhao is a PhD candidate in the Department of Computer Science and a member of the Human–Computer Interaction Laboratory, both at the University of Maryland at Colege Park. Her research focuses on human–computer interaction, with current interests in information visualization, auditory user interfaces, and universal usability. She has an ScM in computer science from the University of Maryland at College Park.

Benjamin K. Smith is a graduate student in the Neuroscience and Cognitive Science program at the University of Maryland at College Park, a project manager at the Laboratory for Automation Psychology, and a member of the Human–Computer Interaction Laboratory. He has an ScB in psychology from Brown University. His research interests include human–computer interaction and cognitive psychology.

Kent Norman is a professor of cognitive psychology at the University of Maryland. He is the director of the Laboratory for Automation Psychology (http://www.lap.umd.edu) and is a founding member of the Human/Computer Interaction Laboratory (http://www.cs.umd.edu/hcil). His research interests include judgment and decision making, human–computer interaction, and usability methods. He has a PhD in experimental psychology from the University of Iowa.

Catherine Plaisant is an associate research scientist at the Human–Computer Interaction Laboratory of the University of Maryland Institute for Advanced Computer Studies. Her research focuses on the design and evaluation of new interface technologies that are usable and useful. Her interests include information visualization and public access information systems. She received a PhD in engineering from the University Pierre and Marie Curie in Paris, France.

Ben Shneiderman is a professor in the Department of Computer Science, founding director of the Human–Computer Interaction Laboratory, and a member of the Institute for Advanced Computer Studies (http://www.umiacs.umd.edu) and the Institute for Systems Research (http://www.isr.umd.edu), all at the University of Maryland at College Park. His research interests include information visualization, human–computer interaction, and universal usability. He has a PhD in computer science from the State University of New York at Stony Brook.

Readers may contact Haixia Zhao at the Dept. of Computer Science, Univ. of Maryland, College Park, MD 20742; haixia@cs.umd.edu.