JVIB

January 2005 • Volume 99 • Number 1

A Report on a Novice User's Interaction with the Internet through a Self-Voicing Application

Keith S. Jones, J. Shawn Farris, Peter D. Elgin, Brent A. Anders, and Brian R. Johnson

---

Abstract: This study used verbal protocol analysis to examine the behavior of an individual with visual impairment using a self-voicing application to find information on the World Wide Web. The results indicated that executing actions (such as typing or pressing keys) and interpreting the computer system's state (data gathering) were the most frequent and time-consuming tasks. Furthermore, the individual had difficulty determining the effects of her actions on the system and whether relevant information was present on a page. These results suggest that there may be problems in interfacing the user with the software and the way textual information is aurally displayed to the user.

The authors thank Rickey Thomas for his suggestion to develop the transition matrix, as well as the anonymous reviewers who provided valuable critiques. This research was funded by a University Small Research Grant from Kansas State University.

Forty-two percent of U.S. households have access to the Internet (Newburger, 2001), and this percentage is expected to increase in the near future. As access to the Internet increases, so will use of the Internet in day-to-day activities. For example, it is becoming increasingly common for companies to provide only Internet-based product documentation or for universities to provide Internet-based registration systems. Therefore, it is becoming increasingly advantageous to be able to find information on the Internet efficiently.

Background

Accessibility for individuals with visual impairments

To obtain access to Internet-based information, many people with visual impairments (that is, those who are blind or have low vision) use software or hardware that presents auditorially information that normally would be displayed graphically on a computer's screen. These technologies, along with other assistive devices, such as screen magnifiers and braille displays, provide access to computer-based resources, including the Internet. This accessibility has a positive impact on the lives of adults with disabilities (Taylor, 2000a, 2000b).

Assistive auditory interfaces sometimes use a standard application, such as Microsoft Internet Explorer or

Excel, in conjunction with a systemwide screen reader such as JAWS, which speaks the elements of the computer's interface aloud (Blenkhorn & Evans, 2002). Alternatively, the auditory interface may use a specialized application that verbalizes only the elements of the particular application's interface (for a discussion of the relative merits of these two approaches, see Blenkhorn & Evans, 2002).

For example, pwWebSpeak is a stand-alone self-voicing application that is designed specifically for accessing the Internet. A Kansas State University student who uses pwWebSpeak to search for an open calculus class at the time of registration must first press the F2 key to access a window that allows a web address to be entered. While the student types, by default, the application speaks the keyed letters to provide feedback. To request this address, the student presses the Enter key. While the page loads, the student hears feedback that describes the loading process (for example, "page loading"). Then, the application speaks the page's contents. To do so, pwWebSpeak reads from the left to the right side of the screen, beginning with the content at the top left of the page. Thus, the application transforms the organized layout of the web page into a serial flow of information.

While listening to the spoken content, the student hears "Fall 2004 Link Course Schedule," which indicates that the fall 2004 course schedule can be found by following the link. The student then presses the Enter key or Shift-F5 to select the link if the student did not respond before the system read the next piece of text.

As one may imagine, for the 1.5 million Internet users with visual impairments in the United States (Gerber & Kirchner, 2001), using an auditory interface like pwWebSpeak for the first time can be daunting. As the example highlights, these complex devices have many control features that must be accessed largely through key commands (for instance, pressing F4 to pause the verbalization). The vast number and sometimes nonintuitive nature of these commands can make it difficult to learn how to use these auditory interfaces.

In addition, when one of these auditory interfaces is used to access the Internet, the task may be more difficult because many web sites are not designed to facilitate their use. For example, one common problem is that images are often used to display text. If the web site designer did not provide a text description of the graphic in the HTML (Hypertext Markup Language) code, then this graphic is unreadable by an auditory interface (Laux, 1998; for a discussion of other major impediments to accessing web sites, see Slatin & Rush, 2002).

Accordingly, the World Wide Web Consortium (W3C, 2002) established guidelines that address accessibility for individuals who use auditory interfaces, including screen readers and self-voicing applications, but many web designers do not adhere to these rules. In addition, many widely available web technologies do not work well with auditory interfaces, including Java applets. The result is that many of the amenities that are taken for granted today, such as Internet-based product information, can be relatively inaccessible to individuals who use auditory computer interfaces.

Usability of auditory interfaces

The literature on the usability of auditory interfaces has focused primarily on the development of tools to assess usability, surveys of web users who use auditory interfaces, and the development and testing of novel interface approaches that may be integrated into next-generation interfaces. In addition, there is a growing literature related to usability testing with individuals with visual impairments. Each of these literatures is discussed briefly.

Usability tools

There is a growing list of software solutions that aid in diagnosing accessibility issues (for a list of currently available tools, see W3C, 2004). One of the most popular of these tools is Bobby (2002), which examines a web site for violations of the W3C's accessibility guidelines and recommends changes that should be made to the site's HTML. This tool can be helpful, especially when web site designers are not knowledgeable about the W3C specifications.

Survey findings

Others have studied the usability of auditory interfaces through surveys (see, for example, Earl & Leventhal, 1999; Gerber, 2002; Leventhal & Earl, 1997). One benefit of these studies is that they provide a way of determining which aspects of usability do and do not vary over time. For example, by comparing their two surveys, Earl and Leventhal reported that training issues related to the use of auditory interfaces were consistently problematic, whereas the choice of operating system and the associated usability issues changed over time. Thus, these surveys, when conducted at regular intervals, are capable of describing variations in usability issues.

Novel interaction techniques

A large proportion of the literature on the usability of auditory interfaces concerns the development of novel interaction techniques that may be incorporated into future versions of the interfaces. For example, James (1998) described several generalizations from his work with auditory HTML interfaces. Specifically, he researched the benefits of using different voices to code different aspects of the display. For instance, links may be spoken in one voice, and the body text may be spoken in a different voice; in this way, the user could be aware when a link is present without the auditory interface being required to say "Link" before each hyperlink.

In addition, several researchers have examined the use of particular sounds or sound characteristics to signify interface components (see, for instance, Alty & Rigas, 1998; Blattner, Sumikawa, & Greenberg, 1989; Gaver, 1989, 1993; Raman, 1997). Gaver (1989) focused on using naturally occurring sounds that would normally be associated with metaphorical aspects of the interface. For instance, Gaver associated emptying the computer's trash can with a clunking sound. In this way, the sounds have a natural relationship with the things that they are associated with and, because of that relationship, can convey information about these events.

Usability testing

There is a growing literature on how to conduct usability testing with people who are visually impaired. Gerber (2002) and Barnicle (2000) explored how existing usability methods may need to be modified for use with people with visual impairments. For instance, Gerber argued that focus groups, a technique that is generally regarded as less than optimal by the usability community (Nielsen, 2001), may be useful for testing with individuals who are visually impaired. Similarly, Coyne and Nielsen (2001b) reported on how best to conduct usability testing with individuals with visual impairments and suggested that usability testers should become familiar with the specific auditory interface that a user will employ (such as pwWebSpeak or JAWS) because, without that knowledge, it may be difficult or impossible to truly understand what the user is doing. In addition, the Usability for Visually Impaired electronic discussion group (archives of which can be found at <http://groups.yahoo.com/group/uvip>) is a vast, freely available knowledge base of usability issues for people with visual impairments.

Focus on users' behavior

The research just discussed, along with the work of others (including Morley, Petrie, O'Neill, & McNally, 1999; Vanderheiden, Boyd, Mendenhall, & Ford, 1991), provide valuable information about the ways in which auditory computer interfaces may be enhanced. However, these studies have not provided information about the behaviors of users of auditory interfaces.

Although there have been some accounts of the mental models formed by experienced auditory interface users (see, for example, Kurniawan, Sutcliffe, & Blenkhorn, 2003; Kurniawan, Sutcliffe, Blenkhorn, & Shin, 2003), there have been no reports of the behavioral processes used by a novice who is trying to find information on the Internet. The study reported here used verbal protocol analysis to examine the behavior of a user with a visual impairment who was searching for information on the Internet.

Method

Participant

The participant (hereafter called Mary, a pseudonym) was a 21-year-old woman who was recruited from Kansas State University's Office of Disabled Student Services. She was paid $8.00 per hour for participating.

Mary had a detached retina in her right eye as the result of an automobile accident. She was able to detect visual information, but had lost all peripheral vision in her right eye and had double vision, frequent perceptions of flashing light, and problems with the extraocular muscles. She said that her left eye usually felt fatigued. Mary reported that looking at a monitor for longer than 20 minutes was painful, and therefore she wanted to learn how to use a purely auditory interface, since it would allow her to access information on the Internet without discomfort.

Thus, before she participated in the experiment, Mary did not have any experience using an auditory computer interface, which was advantageous for our purposes, since we were interested in problems encountered while learning to use an auditory interface. However, she had used computers regularly for five years and the Internet for three years before her accident.

Typically, a usability study includes about five people to balance the cost of data collection with the benefits of finding important usability issues (Nielsen, 1998), and this recommendation has been extended to usability studies involving individuals who are visually impaired (Gerber, 2002). Unfortunately, funding and recruiting-area limits precluded our inclusion of more than one person in the present study. Virzi's (1992) findings, however, indicate that conducting usability testing with a single participant identified 80% of the most severe usability problems, which suggests that valuable information on usability can be garnered from a single participant.

It might also be argued that the research design should have included an individual who was not visually impaired, to determine differences in behavior between users who do and do not have visual impairments. Although such information could be valuable for other purposes, it would not lead to a better understanding of the behavior of users with visual impairments. In addition, because visual and auditory interfaces are two different systems, each with its own usage patterns, the behavior of a sighted individual using a visual interface could not serve as a standard or comparison for the behavior of the visually impaired individual using an auditory interface.

Software and hardware

Access to the Internet was provided by a personal computer using the Windows 98 SE operating system, which had a local-area-network connection with ample bandwidth. The auditory-interface software that was used to access the Internet was pwWebSpeak 32

(1999), which is an application designed for reading web pages. PwWebSpeak was an attractive platform for our purposes because it was relatively inexpensive and has all the major features of an auditory computer interface. Because Mary had some visual capability, it was necessary to ensure that she was not exposed to the visual feedback that was provided by pwWebSpeak. To do so, we completely occluded the computer's graphical user interface (GUI) by using a program that presented the current time of day. All Mary's actions and verbalizations were videotaped with her consent.

Tasks and procedures

All Mary's tasks involved seeking information. Occasionally, Mary sought information for personal interest (like the name of a band's album), but most of her tasks were given to her. All the tasks pertained to specific questions (such as "What causes thunder?"), rather than to general information (see Box 1 for examples of 10 tasks; the complete list of 84 tasks can be found in the online version of the journal at <www.afb.org/jvib990105appendix.asp>). The use of specific search tasks is consistent with the methods used in other usability tests with people who are visually impaired (see, for example, Coyne & Nielsen, 2001a).

The use of specific search tasks, rather than allowing Mary to perform tasks in a free-form manner ("surfing" the Internet) while examining where she had difficulty, may have limited the sources of information in our study. Such freedom, however, would have limited the useful data related to any given set of tasks, such as searching versus browsing. In addition, given that the focus of this study was on understanding how a user finds specific pieces of information on the Internet, the use of specific search tasks seemed appropriate. The nature of Mary's tasks required her to explore a variety of web sites that were each designed and organized differently. Since Mary was required to interact with such a diverse set of web site designs, it is doubtful, therefore, that the web sites' design or organization had a confounding effect on this research.

Mary was instructed to verbalize all her thoughts while searching for information, not just those related to concrete acts such as pressing a key. If the rate of her verbalizations slowed, the experimenter reminded her to continue to think aloud. Before testing, Mary first had to be familiarized with the software and hardware and testing conditions.

The first two sessions (a total of 8 hours) were spent familiarizing Mary with the auditory interface while doing concurrent verbal protocols. Thus, these sessions represent situations in which she was learning two tasks: to use the interface and to do verbal protocols. At the end of the second session, Mary stated that she was comfortable with the testing environment and procedures, so official data collection began at the beginning of the third session. The first two practice sessions were not used for data analysis; only the final 10 sessions (a total of 39 hours) were used.

Thirty-nine hours is longer than the typical usability testing session (Mayhew, 1999). However, because of the serial nature of the tasks, using an auditory interface might take longer than a GUI, for which parallel processing of information is easier. In addition, since Mary was a novice user, it was desirable to ensure that our testing time frame covered any trends in learning that we might encounter.

Data coding

Verbal protocols were broken down into individual units of thought. Each utterance was time stamped. To classify the utterances, we used Norman's (1988) seven-stages-of-action model, a general model of how users interact with a technological system, as a coding scheme. The seven stages, with their translation into a coding scheme for the verbal protocols, are shown in Table 1. An eighth category was added to classify irrelevant comments. Given that there are no models of interaction specific to auditory interfaces, Norman's model was chosen as a coding scheme, as it is general enough to encompass the specifics of auditory interfaces and is generally easier to understand than the alternatives.

Two judges independently coded all the utterances. They were allowed to select more than one code, if necessary, but only the codes that both judges agreed on were analyzed. Cohen's kappa was used to assess interrater reliability. This analysis revealed a reliability of .66 (N = 6,748) with the miscellaneous category included and .73 (N = 5,653) when the miscellaneous category was treated as missing data, indicating substantial agreement higher than chance (Fleiss, 1981; Landis & Koch, 1977). Although these values are inflated slightly by the way in which differences in judging were reconciled, the values are acceptable.

Finally, to analyze the protocols, we condensed the codes into a single data set by applying three rules: (1) If only one of the judges coded an utterance as miscellaneous, then the nonmiscellaneous code was used; (2) if both judges coded an utterance as miscellaneous, then the utterance was removed from the data set, and the time associated with that utterance was added to the previous utterance (that is, the previous thought or behavior was assumed still to be operating); and (3) if both judges assigned different codes, and neither was coded as miscellaneous, then the utterance and the time were eliminated from the data set. This procedure resulted in 5,283 coded and time-stamped utterances that were analyzed. The majority of the excluded data points were coded as miscellaneous by both judges. The coded and time-stamped utterances that were included in the study accounted for roughly 32 hours' worth of data collection.

Results and discussion

Overview

The data were analyzed in three different ways. First, the total number of utterances for each stage and the total amount of time spent in each stage were analyzed, which identified the stages of action that were the most labor intensive. Second, the amount of time spent per utterance in each of the stages was analyzed, which provided a view of the amount of time spent processing different kinds of information. Last, the frequencies of various transitions among the stages were analyzed, which provided information about the relationship among the stages, something that could not be captured by the other analyses.

Frequency of utterances for each stage

Figure 1 shows the percentages of utterances across all the sessions for each of Norman's stages of action. The two most salient categories are Executing the Action and Interpreting the System State. The high occurrence of Executing the Action utterances (31.04%) indicates that Mary frequently expended effort performing physical actions (for example, going back and typing in the address). In addition, the high occurrence of Interpreting the System State utterances (44.69%) indicates that she frequently expended cognitive or attentional effort in collecting information that was relevant to her goal, such as finding information or links on a page.

Amount of time spent in each stage

Figure 1 also shows the percentage of time that was spent at each stage of action during the 32 hours of data collection. Again, the two most notable categories of action were Executing the Action and Interpreting the System State. The high proportion of time spent executing actions (26.77%) shows that Mary spent a large amount of time performing physical actions, but that she spent even more time interpreting the system state (53.31%). Apparently, collecting information with an auditory interface requires a lot of time. Although Forming the Goal took little time, this was a consequence of the methodology, since typically the goal was given to her.

Amount of time spent per utterance

Figure 2 shows the mean time spent performing the various categories of action for each utterance. As can be seen by the 95% confidence intervals for the means, Interpreting the System State (*M* = 26.27 seconds) took a lot of time whenever Mary arrived at this stage. Although the middle five stages all require a lot of time, Interpreting the System State consistently took substantially longer than did all the other stages, which suggests that collecting information was the most labor-intensive stage for these tasks.

One may argue that when Mary was listening for goal-relevant information, she was simply listening to the output of the speech synthesizer for the information that she sought and that doing so is not labor intensive. The literature on vigilance (sustained attention) suggests otherwise, however. Specifically, monitoring tasks, such as listening for a specified signal, are cognitively demanding (Parasuraman, 1986). Thus, it is likely that a large amount of time spent Interpreting the System State reflects a true labor-intensive component of using an auditory interface.

Implications

The finding that Executing the Action and Interpreting the System State accounted for the highest percentages of both utterances and time spent in each stage does not, by itself, indicate a problem with the auditory interface. That is, the high frequencies associated with these stages of action may be normal for operating any auditory computer interface. The finding does suggest, however, that if a designer wants to simplify the task of searching for information on the Internet via an auditory interface, then he or she needs to concentrate on these two stages of action.

Design implications

These findings suggest that designers should evaluate existing feedback mechanisms to determine whether they facilitate the search for information. Specifically, the data suggest that Mary expended a great deal of time and effort (16 of the 32 hours that were analyzed) determining whether pertinent information was present on a given web page. Thus, facilitating this important process should greatly simplify the use of auditory interfaces for information-seeking tasks.

Several authors have used the spatial qualities of sound to facilitate the search for clues as to whether relevant information is on a given web page (see, for example, Goose & Möller, 1999; Savidis, Stephanidis, Korte, Crispien, & Fellbaum, 1996), especially when several sources of information are presented simultaneously, allowing users to move through the interface more efficiently. Other authors have suggested that annotation strategies may facilitate the identification of instances of goal-relevant information on a web page (see, for instance, Asakawa & Takagi, 2000), and some auditory web interfaces, like BrookesTalk, have incorporated such a feature (see, for example, Zajicek, Powell, Reeves, & Griffiths, 1998). Specifically, a computer system could examine a web page's HTML code and create separate annotations that inform users about how the page is laid out and about the web page's content. For example, BrookesTalk allows users to access lists of headings, links, and keywords, as well as an abridged version of the web page, to name just a few of the options (Zajicek et al., 1998). Given the results of our study, it is likely that the inclusion of these kinds of annotations will help a user determine whether further inspection of the web page is necessary.

Frequency of transitions among the stages of action

The frequency with which Mary transitioned among the stages of action was examined. It was thought that this might identify issues that were not addressed by the other analyses.

Expected transitions

On the basis of Norman's (1988) model, we expected that Mary would transition among the stages of action in a manner consistent with the solid arrows in Figure 3. That is, we expected that she would begin by forming a goal, which would lead to forming an intention about how to move toward that goal, which, in turn, would motivate the specification and subsequent execution of actions, which would cause a change in the state of the system. She would then perceive and subsequently interpret the changes in the state of the system and evaluate whether the desired outcome had been met. Once this evaluation had taken place, we expected that if the goal was not satisfied, she would either specify a different action that was consistent with the initial intention or form a different intention.

Observed transitions

To determine the types of transitions that were characteristic of Mary's behavior, we created a matrix that provided the frequency associated with each transition from one stage to another. As can be seen in Figure 4, data above the diagonal of the matrix (indicated by the gray cells) represent situations in which Mary moved forward through the stages of action. Data along the diagonal of the matrix represent situations in which Mary did not transition to a different stage of action but, rather, stayed in the current stage. Data below the diagonal of the matrix represent situations in which Mary transitioned backward in the stages of action.

Given the goal of identifying bottlenecks in usability, our attention focused on relatively large frequencies along or below the diagonal of the matrix (that is, relatively high frequencies where Mary either did not progress through the stages of action or backtracked through the stages). On the basis of visual inspection, cells that were marked by black circles or crosses warranted closer examination. Figure 3 provides a graphical depiction of these marked transitions (the dotted and dashed arrows) overlaid on the expected transitions based on Norman's (1988) model.

Potentially nonproblematic transitions

On the basis of our inspection of the transcripts of Mary's verbal reports, it appeared that several of the high-frequency transitions were not problematic. These transitions are denoted by black circles in Figure 4 and with dotted arrows in Figure 3. Specifically, it appears that the high frequencies associated with transitioning from Executing the Action to Specifying the Action (194), as well as transitioning from Interpreting the System State to Specifying the Action (142) happened because Mary did not verbalize all her thoughts and actions.

Our examination of the transcripts of Mary's verbal report also indicated that the high frequency associated with transitioning from Executing the Action to Specifying the Action resulted from Mary verbalizing what she did, rather than what she planned to do. In addition, it appeared that the high frequency associated with Executing the Action and then Executing the Action again (476) resulted from Mary making repeated actions of the same type (for instance, repeatedly going back through web pages to find a previous page). Because of the nature of these transitions, these four high-frequency transitions were not considered problematic.

Potentially problematic transitions

The other transitions (marked with black crosses in Figure 4 and with dashed arrows in Figure 3) were deemed potentially problematic after inspection of the verbal reports. Specifically, the relatively high frequency associated with Perceiving the System State, followed by continuing with the same stage (158) occurred because Mary repeatedly verbalized that she had yet to discover what effect her action had on the system. Gerber (2002) also noted that the users in her testing sessions sometimes had difficulty determining whether an action (like clicking a link) resulted in the desired consequence (like going to the selected page). This difficulty was exacerbated by the fact that auditory interfaces present information serially. Because of this difficulty, Mary had to continually attend to the stream of information, in the hopes of determining what effect her behavior had on the system's behavior.

A relatively high frequency was also associated with transitioning from Perceiving the System's State to Executing an Action (244). Thus, Mary was frequently unable to determine what the system was doing, and instead of waiting for additional information, she chose to execute an action that would take her elsewhere.

These last two findings suggest that designers need to focus on providing more information about what the system is doing. User testing should be conducted to determine the critical pieces of information about a system's behavior that are missing from current auditory interfaces. In addition, the findings suggest that the existing information about what the system is doing may need to be made more salient. Using different voices for different kinds of information (for example, system status versus content), as suggested by James (1998), may help alleviate this problem.

Furthermore, the largest frequency in the matrix was associated with Interpreting the System's State, followed by continuing to Interpret the System's State (1,599). An inspection of the transcripts indicated that Mary often had difficulty determining whether the information that she desired was, in fact, on the page to which she was currently listening. Accordingly, she frequently made numerous successive verbalizations about trying to determine whether the desired information was on the current page. As Barnicle (2000) noted, this situation is exacerbated by the fact that auditory interfaces present information in an inherently serial manner, which tends to force the user to listen to a great deal of irrelevant information. This concern is consistent with the findings of the earlier analyses and the general recommendations pertaining to serial presentations, such as using overviews or annotation strategies, that apply here as well.

To underscore the prevalence of these types of transitions, it is important to note that the last three types of transitions accounted for roughly 38% of the total number of transitions. Accordingly, improvements that reduce the number of these unexpected transitions could significantly improve the usability of auditory computer interfaces.

Conclusion

This study examined the behavior of a visually impaired individual as she found information on the Internet through use of a self-voicing application. Care should be taken not to generalize the results because only a single user working with a single type of self-voicing application was studied. Pending future confirmation through other experimental means, however, it appears that such auditory computer interfaces could be improved by (1) redesigning the feedback mechanisms that inform users of the system's status and (2) modifying the interface to help users identify goal-relevant information without having to move serially through a document.

References

Alty, J. L., & Rigas, D. I. (1998). Communicating graphical information to blind users using music: The role of context. *Proceedings of ACM CHI 98 Conference on Human Factors in Computing Systems* (pp. 574–581). New York: ACM Press.

Asakawa, C., & Takagi, H. (2000). Annotation-based transcoding for nonvisual web access. *Fourth Annual ACM Conference on Assistive Technologies* (pp. 172–179). New York: ACM Press.

Barnicle, K. (2000). Usability testing with screen reading technology in a Windows environment. *Proceedings of the 2000 International Conference on Intelligent User Interfaces* (pp. 102–109). New York: ACM Press.

Blattner, M. M., Sumikawa, D. A., & Greenburg, R. M. (1989). *Human-Computer Interaction*, *4*, 11–44.

Blenkhorn, P., & Evans, D. G. (2002, March). *Augmenting the user interface of standard applications for improved screen reading for blind people*. Paper presented at the 17th Annual International Conference on Technology and Persons with Disabilities [Online]. Available: http://www.csun.edu/cod/conf/2002/proceedings/55.htm

Bobby [computer software]. (2002). Lexington, MA: Watchfire.

Coyne, K. P., & Nielsen, J. (2001a). *Beyond ALT text: Making the web easy to use for users with disabilities* [Online]. Available: http://www.nngroup.com/reports/accessibility/testing

Coyne, K. P., & Nielsen, J. (2001b). *How to conduct usability evaluations for accessibility* [Online]. Available: http://www.nngroup.com/reports/accessibility/testing

Earl, C. L., & Leventhal, J. D. (1999). A survey of Windows screen reader users: Recent improvements in accessibility. *Journal of Visual Impairment & Blindness, 93*, 174–177.

Fleiss, J. L. (1981). *Statistical methods for rates and proportions* (2nd ed.). New York: John Wiley & Sons.

Gaver, W. W. (1989). The SonicFinder: An interface that uses auditory icons. *Human-Computer Interaction, 4*, 67–94.

Gaver, W. W. (1993). Synthesizing auditory icons. *Proceedings of ACM INTERCHI'93 Conference on Human Factors in Computing Systems* (pp. 228–235). New York: ACM Press.

Gerber, E. (2002). Surfing by ear: Usability concerns of computer users who are blind or visually impaired. *AccessWorld*, *3*(1), 38–43. Available: http://www.afb.org/aw/AW030107.asp

Gerber, E., & Kirchner, C. (2001). Who's surfing? Internet access and computer use by visually impaired youths and adults. *Journal of Visual Impairment & Blindness, 95,* 176–181.

Goose, S., & Möller, C. (1999). A 3D audio-only interactive web browser: Using spatialization to convey hypermedia document structure. *Proceedings of the Seventh ACM International Conference on Multimedia* (pp. 363–371). New York: ACM Press.

James, F. (1998). Lessons from developing audio HTML interfaces. *Third Annual ACM Conference on Assistive Technologies*, 27–34.

Kurniawan, S. H., Sutcliffe, A. G., & Blenkhorn, P. L. (2003). How blind users' mental models affect their perceived usability of an unfamiliar screen reader. In M. Rauterberg, M. Menozzi, J. Wesson (Eds.), *INTERACT 03* (pp. 631–638). Amsterdam: IOS Press.

Kurniawan, S. H., Sutcliffe, A. G., Blenkhorn, P. L., & Shin, J. (2003). Investigating the usability of a screen reader and mental models of blind users in the Windows environment. *International Journal of Rehabilitation Research*, *26,* 145–148.

Landis, J., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics, 33,* 159–174.

Laux, L. (1998). Designing web pages and web applications for people with disabilities. In C. Forsythe, E. Grose, and J. Ratner, (Eds.), *Human factors and web development* (pp. 87–96). Hillsdale, NJ: Lawrence Erlbaum.

Leventhal, J. D., & Earl, C. L. (1997). Accessing Microsoft Windows with synthetic speech: A user survey. *Journal of Visual Impairment & Blindness*, *91*, 12–16.

Mayhew, D. (1999). *The usability engineering lifecycle*. San Francisco: Morgan Kaufman.

Morley, S., Petrie, H., O'Neill, A., & McNally, P. (1999). Auditory navigation in hyperspace: Design and evaluation of a non-visual hypermedia system for blind users. *Behaviour & Information Technology, 18*, 18–26.

Newburger, E. C. (2001). *Home computers and Internet use in the United States: August 2000.* Special Studies. Washington, DC: U.S. Census Bureau.

Nielsen, J. (1998, May). Cost of user testing a website. *Alertbox* [Online]. Available: http://www.useit.com/alertbox/980503.html

Nielsen, J. (2001, August). First rule of usability? Don't listen to users. *Alertbox* [Online]. Available: http://www.useit.com/alertbox/20010805.html

Norman, D. A. (1988). *The design of everyday things.* New York: Basic Books.

Parasuraman, R. (1986). Vigilance, monitoring, and search. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and human performance. Vol. 2: Cognitive processes and performance* (pp. 43.1–43.39). New York: John Wiley & Sons.

pwWebSpeak 32 (Version 3.0, release 35) [computer software]. (1999). Trenton, NJ: Productivity Works.

Raman, T. V. (1997). *Auditory user interfaces: Toward the speaking computer*. Boston: Kluwer Academic.

Ryder, J. W., & Ghose, K. (1999). Multi-sensory browser and editor model. *Proceedings of the ACM Symposium on Applied Computing* (pp. 443–449). New York: ACM Press.

Savidis, A., Stephanidis, C., Korte, A., Crispien, K., & Fellbaum, K. (1996). A generic direct-manipulation 3D-auditory environment for hierarchical navigation in non-visual interaction. *Second Annual ACM Conference on Assistive Technologies,* 117–123.

Slatin, J. M., & Rush, S. (2002). *Maximum accessibility: Making your web site more usable for everyone.* Reading, MA: Addison-Wesley.

Taylor, H. (2000a). How the Internet is improving the lives of Americans with disabilities. *The Harris Poll, #30.* Los Angeles: Creators Syndicate.

Taylor, H. (2000b). Many people with disabilities feel isolated, left out of their communities, and would like to participate more. *The Harris Poll, #34*. Los Angeles: Creators Syndicate.

Vanderheiden, G. C., Boyd, W., Mendenhall, J. H., & Ford, K. (1991) Development of a multisensory nonvisual interface to computers for blind users. *Proceedings of the Human Factors Society 35th Annual Meeting*, 315–318.

Virzi, R. A. (1992). Refining the test phase of usability evaluation: How many subjects is enough? *Human Factors*, *34*, 457–468.

World Wide Web Consortium. (2002). *Web Accessibility Initiative* [Online]. Available: http://www.w3.org/WAI

World Wide Web Consortium. (2004). *Evaluation, repair, and transformation tools for web content accessibility* [Online]. Available: http://www.w3.org/WAI/ER/existingtools.html

Zajicek, M., Powell, C., Reeves, C., & Griffiths, J. (1998). Web browsing for the visually impaired. In A. Edwards, A. Arato, & W. Zagler (Eds.), *Computers and assistive technology, ICCHP ’98* (pp. 161–169). Vienna: Austrian Computer Society.

Keith S. Jones, Ph.D., *assistant professor, Department of Psychology, Texas Tech University, Box 42051, Lubbock, TX 79409–2051; e-mail: <keith.s.jones@ttu.edu>.* ***J. Shawn Farris, Ph.D.,*** *senior human factors associate, Mail stop 073-3313, 285 International Parkway, Lake Mary, FL 32746; e-mail: <shawn.farris@convergys.com>.* ***Peter D. Elgin, M.A.,*** *graduate student, Kansas State University, 492 Bluemont Hall, Manhattan, KS 66506; e-mail: <elgin@ksu.edu>.* ***Brent A. Anders, M.A.,*** *mediated education coordinator, Kansas State University, 58 College Court, Manhattan, KS 66506; e-mail:*

<superman@ksu.edu>. ***Brian R. Johnson, M.A.,*** *graduate student, Department of Psychology, Texas Tech University, Lubbock, TX 79409; e-mail: <b.johnson@ttu.edu>. Address correspondence to Dr. Jones.*

Previous Article | Next Article | Table of Contents

Search JVIB | JVIB Policies | Contact JVIB | Subscriptions | JVIB Home

www.afb.org | Change Colors and Text Size | Contact Us | Site Map |

Site Search [ ] Go

About AFB | Press Room | Bookstore | Donate | Policy Statement