

# Hear It Is: Enhancing Rapid Document Browsing with Sound Cues

Parisa Eslambochilar, George Buchanan, and Fernando Loizides

Future Interaction Laboratory, Swansea University, UK  
Centre for HCI Design, City University, UK  
{p.eslambochilar,f.loizides}@swansea.ac.uk,  
george.buchanan.1@city.ac.uk

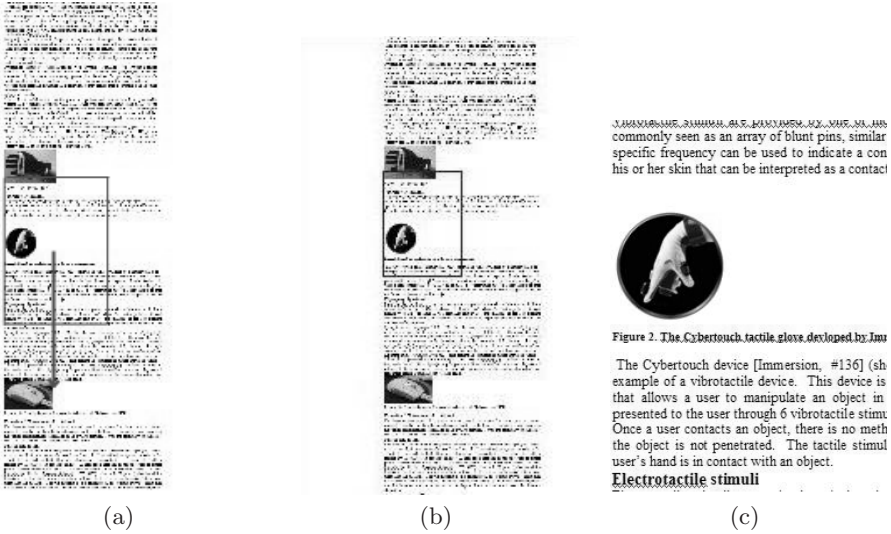
**Abstract.** Document navigation has become increasingly commonplace as the use of electronic documents has grown. Speed-Dependent Automatic Zooming (SDAZ) is one popular method for providing rapid movement within a digital text. However, there is evidence that details of the document are overlooked as the pace of navigation rises. We produced a document reader software where sound is used to complement the visual cues that a user searches for visually. This software was then evaluated in a user study that provides strong supportive evidence that non-visual cues can improve user performance in visual seeking tasks.

## 1 Introduction

Document navigation has received increasing coverage from researchers in recent years. This paper reports a project that develops the Speed-Dependent Automatic Zooming method of navigation [14]. The SDAZ method applies to all digital navigation (e.g. of maps, images, etc.) but has been used with particular success in the task of navigation with digital texts [1,14].

This paper reports an initial investigation into how users' document navigation could be enhanced through the deployment of simple audio cues. We begin by exploring the navigation of documents where vision is most hindered; namely when navigating on a small screen device. There are many cases where this technique would be inappropriate (e.g. where the sound was made within a library's reading room), but when a user is on the move, or sitting at their own computer wearing headphones, additional support can be given to what is known to be a surprisingly challenging task. Such situations are increasingly commonplace as the use of digital library materials moves from being solely on the desktop.

The results show that the audio cue was sufficient for the users to distinguish where the target is. Results show improved timings for the completion of tasks and increased reliability in the successful location of sections. We also regularly witness a lag in reaction times when a user overshoots the target location. The paper commences with a related works section, followed by a description of our sound-enhanced document reader software. We then report an initial user study, conducted to evaluate the impact of this technique on simple navigational tasks. The paper progresses to a detailed discussion of our findings and work that



**Fig. 1.** A user is scrolling through a document using touch-controlled SDAZ, (a) Scrolling through the text document at maximum speed, (b) the user has released the stylus. Zero scroll speed in document. Vertically falling to 100% magnification, and (c) finished zooming. Document is at 100% magnification.

closely relates to it. We conclude with a review of future work in this area, and the opportunities that exist to make further improvements to user’s interaction with digital texts.

## 2 Related Work

SDAZ is a simple technique that automatically varies the current level of zooming used in a view with the speed of the user’s movement. Rapid movement results in the document view zooming outwards to view a larger area in less detail, while when the user’s motion slows, the view zooms in to see a small area in fine relief (see Fig. 1). This intuitive model for a navigator first was proposed and implemented by Igarashi and Hinckley [14] for browsing large documents and maps on Desktop PCs. Cockburn et al. [1] suggested and successfully evaluated the refinements of this approach for navigating digital documents such as PDF files on Desktop PCs. Eslambolchilar and Murray-Smith [6,7] suggested a control-centric approach to control both zooming and panning by tilting alone or stylus alone on mobile platforms.

Even within the mobile domain, there is currently minimal research on how to use the very same cues for other navigation tasks. [9] presents the use of multimodal feedback in a document browsing task based on tilt-controlled SDAZ on a mobile platform. They assigned audio and haptic (tactile) feedback to important structures in the document such as headers, figures, tables etc.. They

showed that audio or tactile feedback can support intermittent interaction, i.e. allowing movement-based interaction techniques to continue while the user is simultaneously involved with real life tasks.

Researchers in mobile devices have been exploring many other methods of enhancing navigation in general. One popular area of research is exploring what utility non-visual cues provide for navigation in the physical world. One example is Warren et al's use of music to indicate a user's position compared to their chosen destination [17]. In that case, if the target is to the user's left, the music is panned in that direction; if the user turns to face in its direction, the music gradually moves to the user's centre of attention.

Surprisingly, the progress made in using non-visual cues for navigation in the physical world has made little impact on document navigation. We therefore explore the utilisation of sonification to aid the document navigation process. Sonification is defined as the use of non-speech audio to convey information [10]. More specifically, sonification is the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation [13].

It is known that during visual search of documents, users often initially overshoot a target, and subsequently need to back-track to return to their chosen location. This knowledge has even been used to successfully backtrack automatically [16]. Whereas optical search in a dense visual target is problematic, we anticipated that an audio cue could assist the user adjust their navigation strategy in the presence of a potential target: e.g. slowing down their scrolling rate to identify the headings on display.

### 3 Sound-Enhanced Document Navigation

Rather than embark on providing a large number of different techniques that would be difficult to evaluate systematically at the same time, we focused on one particular element of user navigation that is relatively well documented. Researchers have indicated [2,5,15], that users utilise document headings for multiple tasks during digital navigation. Headings provide a cue for locating other material (e.g. from previous reading, a user may know that a particular detail is found in Section 4.3), and are also used to support skimming a document (determining if a section is likely to contain text of interest). Recent research has demonstrated that enhancing the visibility of headings makes a number of improvements to document navigation [3].

We followed this well-established direction of research by providing simple cues for the presence of headings. As an intuitive model of the sonification process, we can imagine the text on the screen to be embossed on the surface. This embossed type excites some object (an elastic band or sharp striking of speed ramps, for example) as it is dragged over the text. This physically motivated model is similar in nature to the model-based sonifications described by Hermann and Ritter [12]. As each single heading passes into view, a simple "click" sound is heard by the user. If a user navigates rapidly up-and-down a document, they will experience a sequence of clicks as different headings become visible.



**Fig. 2.** One of participants who plays with the sonic-enhanced document navigator on a PDA

For the implementation of our sound-enhanced document navigator we pursued the refinements of SDAZ approach suggested by [6,7] and Cockburn et al. [1]. We developed a stylus-controlled SDAZ-based document reader on a Personal Digital Assistant (PDA) device. This document reader utilises sonic cues (Sonification, see Sec. 2) to assist the reader’s navigation. The current reader is implemented in C++, and uses the fmod API<sup>1</sup> and GAPI Draw library<sup>2</sup> to load and display each document. A document is then parsed to identify the location of each heading. These locations are interpreted into scroll positions, and as the user scrolls up and down the document, a click is produced when a heading passes into view.

## 4 User Study

To evaluate the impact of our sonic cues on user navigation we undertook a laboratory-based study. Our participants consisted of 24 university undergraduate and postgraduate students in a computer science department. We carried out a within-subjects design, allowing for 12 participants in each group to use one of two available variations of the SDAZ software on a PDA device with a touch screen facility, using a stylus to navigate. The PDA used throughout was a Dell AX70 (Fig. 2).

The first group used SDAZ without sonic cues whereas the second utilised the available sonic cue additions to the SDAZ software. Users were given a pre-study questionnaire to establish familiarity with small screen devices and to list the difficulties experienced in using these devices for reading, navigating and searching within documents on a small screen device.

Following the pre-study questionnaire all participants were given the same 50 page document containing plain text and 5 headings only. The document dimensions were 20000 height by 250 width (measured in pixels). The first heading contained the word “Bob”, the second the word “Sam”, third “Tom”, fourth “Tim” and the final heading contained the word “Kim”. These headings, which

<sup>1</sup> Music and sound effect system [www.fmod.org](http://www.fmod.org)

<sup>2</sup> Graphic api draw [www.gapi.org](http://www.gapi.org)



**Fig. 3.** Heading locations respective to document

were located in different pages of the document (see Fig. 3), were distinguishable clearly from a close zoom but semi-obscured when scrolling fast from a distant zoom. At maximum zoom, the title was almost, yet not impossible, to distinguish from the text. Participants were introduced to the system and allowed time to familiarise themselves with its navigation capabilities.

The participants were then given a set of 8 tasks to undertake, one at a time, using the software assigned to them. They were encouraged to complete each task in as little time as possible. The study was aimed at exploring behaviour for realistic within-document information seeking tasks. The first task was identifying the number of headings in the document. This simulated an information seeker trying to get an overview of document content during document triage before more in depth search and reading. The subsequent four tasks were to find specific headings, just as seekers would try to locate relevant sections they had previously found. The final three tasks were to find the first, second and final headings. These tasks emulated an information seeker locating the introduction, the subsequent section and conclusions of document. It should be noted here that after completing the first task participants were also asked if they could recall the heading titles and order. It was found that 22 participants could remember the first and final titles but were uncertain about the titles and order of the remaining three. Participants' navigational patterns were recorded in real time by the software. This data included position within the document, zoom values and speed of scrolling. Notes were taken during the tasks and participants were asked to comment on their behaviour after every task was completed. After all tasks were completed, we conducted a semi-structured interview asking participants to express the positive aspects they found, as well as any difficulties that hindered their tasks from being achieved efficiently. Participants from both groups were then presented with the alternate software and asked to choose between the two in terms of usefulness.

## 5 Results

We evaluated the navigation results for three following sets of tasks: finding number of headings in the document, locating a heading by content), and locating the heading by order (e.g. second, last). Before comparing the figures we take a look at a common behaviour seen in the navigational behaviour of the participants using the audio cues.

### 5.1 Overshooting and Backtracking

A common occurrence between the participants with audio assistance is that of overshooting titles while scrolling and needing to backtrack to land on the




Fig. 4. Overshooting and backtracking in sonic-enhanced document navigator

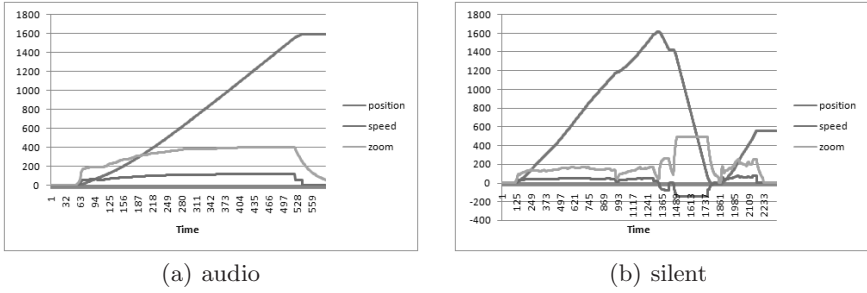
title. This is due to the reaction time between the participant hearing the audio cue and the time it took to release the stylus from the screen. Participants when backtracking were also moving at a decreased speed so as not to miss the title again and consequently at a much lower zoom percentage. We notice this behaviour clearly from all the participants using the audio cues and all have commented on it as a disadvantage. One participant noted that it “would be nice to have an indication of a heading coming up rather than going too far”. It should also be noted that this also affects the timings for audio cues in a negative way, since users, if provided with a predictive model to land successfully on the titles are likely to improve performance. The common overshooting and backtracking behaviour can be seen in Fig. 4.

## 5.2 Finding Number of Headings

The first task that users were required to undertake was to identify the total number of headings in the example document. This was an overview-style task that also familiarised the participant with the document.

The 12 participants using the audio cue all completed the task correctly. Their navigational behaviour was linear, progressing through the document in an orderly manner (see Fig. 5(a)). There were no instances of backtracking: all motion was downward through out the document. The mean zoom level was 307% across this task for all participants ( $sd=48.03\%$ ). The participants assisted by the audio cues immediately commented saying that the audio cues are “very handy for skim reading titles”. They were sure that they did “not miss titles and therefore useful sections on the way”.

In contrast, the 12 participants using the silent interface all demonstrated some backtracking behaviour: six showing localised backtracking of less than 100 pixels, the others showing large-scale movement (see Fig. 5(b)). For example, silent participant 4 moved forward to the end of the document, reversed the



**Fig. 5.** Between group navigation comparison for the first task

entire length of the document, and then again moved forward to the end. The mean zoom level was a mere 129% ( $sd=37.09\%$ ).

We analysed the differences for this first task by applying Student’s  $t$ -test for equal variance with degrees of freedom ( $df$ ) set at 22. For the average zoom level per participant, this yielded  $p=0.0079$  ( $t=2.9242$ ). The same test for the maximum zoom level gave  $p=0.58$  ( $t=0.5614$ ), indicating that there was no significant difference in the maximum zoom level. As the zoom level is related to the current speed of movement, similar results could be expected for the comparison of speed. Ignoring the direction of movement, the average movement rate for the audio group was 85.48 pixels/second ( $sd=15.51$ ), whilst the average for the silent group was 36.57 ( $sd=10.29$ );  $p=0.0154$  ( $t=2.2677$ ).

We also analysed the time taken to find all headings in both scenarios. Data show that participants in the audio group ( $sd=4.67$ , average=19.2 seconds) performed significantly faster than participants in the silent group ( $sd=47$ , average=86 seconds);  $p=0.009$  ( $t=-4.91$ ). Similarly, the total distance traveled by the audio group ( $sd=841.4$ , average=47800 pixels) was shorter than the silent group ( $sd=85291$ , average=101000);  $p=0.0408$  ( $t=2.1728$ ).

### 5.3 Locating by Heading Content

There were four tasks that required the participants to find a given heading, based on the content. The data is presented in Table 1. Data in columns “Audio” and “Silent” present the average and standard deviation. Data in column “Significant” present the  $p$ -value and  $t$ -value using Student’s  $t$ -test ( $df=22$ ). Zoom levels are in percentages and the speed given in pixels per second.

The initial goal of the four was to find the heading that contains “Bob”. There was no significant difference in the traveled distance between the audio and silent group.

The second heading navigation task was to find the heading “Sam”. There was no significant difference in the traveled distance between the audio and silent group. Backtracking behaviour was more common, and observed in 70% of participants in the audio group. However, 40% of participant in the silent group also showed this behaviour.

**Table 1.** Data analysis for locating headings by content. “Audio” column: average (sdev), “Silent” column: average (sdev), and “Significant” column: Student’s t-test; p-value (t-value).

Bob	Audio	Silent	Significant
Time (secs)	4.85 (0.86)	13.7 (24.34)	0.17 (t=1.4)
Zoom (Average)	86.62 (44.57)	57.41 (58.15)	0.091 (t=1.38)
Zoom (Max)	229.04 (94.33)	189.63 (139.89)	0.21 (t=0.81)
Speed (Average)	7.38 (1.75)	6.11 (3.17)	0.12 (t=1.21)
Speed (Max)	66.19 (21.64)	55.73 (31.32)	0.17 (t=0.95)
Sam	Audio	Silent	Significant
Time (secs)	18.4 (18.41)	11.6 (4.21)	0.23 (t=1.24)
Zoom (Average)	166.23 (60.30)	98.42 (28.85)	0.0009 (t=3.51)
Zoom (Max)	341.65 (107.30)	235.05 (52.39)	0.002 (t=3.09)
Speed (Average)	24.13 (10.68)	22.82(5.75)	0.35 (t=0.37)
Speed (Max)	94.75 (24.38)	69.79 (15.74)	0.003 (t=2.98)
Tom	Audio	Silent	Significant
Time (secs)	16.7 (8.85)	16.4 (5.51)	0.93 (t=0.084)
Zoom (Average)	199.67 (35.21)	147.98 (35.01)	0.00078 (t=3.61)
Zoom (Max)	359.29 (90.76)	250.04 (75.04)	0.002 (t=3.21)
Speed (Average)	43.94 (12.52)	39.68(11.00)	0.19 (t=0.88)
Speed (Max)	106.21 (25.52)	75.71 (22.04)	0.0024 (t=3.13)
Kim	Audio	Silent	Significant
Time (secs)	25.71 (3.7)	31.2 (12.7)	0.158 (t=1.46)
Zoom (Average)	247.82 (43.00)	178.46 (67.18)	0.0032 (t=3.01)
Zoom (Max)	413.19 (77.80)	286.54 (113.35)	0.0021 (t=3.19)
Speed (Average)	73.26 (62.13)	48.30(15.11)	0.095 (t=1.35)
Speed (Max)	139.95 (76.65)	85.68 (34.97)	0.018 (t=2.23)

The third heading to find was “Tom”. There was no significant difference in the traveled distance between the audio and silent group. Similar backtracking behaviour to “Sam” heading was observed, but even stronger in this case: 85% of participants in the audio case showed this behaviour.

In the last heading navigation task, participants were asked to locate a heading that contained “Kim”. Due to backtracking in the audio mode there was a difference in the traveled distance: the average traveled distance in the audio group was 49100 pixels (sd=3506) vs 41600 pixels (sd=4315) in the silent group. The t-test result was p=0.042 (t=2.157).

#### 5.4 Locating by Heading Order

There were three tasks that required the participant to find a given heading, based on the order of its appearance. Basic results are seen in Table 2. The data layout in this table is similar to Table 1. Zoom levels are measured in percentage and the speed is measured in pixels per second.

The initial goal of the three was to find the first heading in the document. The average traveled distance in the audio group was 2080 pixels (sd=693) vs 3140



**Table 2.** Data analysis for locating headings by content. “Audio” column: average (sdev), “Silent” column: average (sdev), and “Significant” column: Student’s t-test; p-value (t-value).

1st Heading	Audio	Silent	Significant
Time	4.33 (0.745)	4.95 (5.089)	0.6769 (t=0.422)
Zoom (Average)	109.61 (14.09)	79.29 (32.11)	0.0052 (t=3.08)
Zoom (Max)	210.73 (50.25)	140 (64.72)	0.0061 (t=2.99)
Speed (Average)	20.75 (3.68)	18.52 (8.07)	0.197 (t=0.897)
Speed (Max)	67.62 (16.32)	42.09 (15.08)	0.0006 (t=4.24)
2nd Heading	Audio	Silent	Significant
Time	2.44 (0.374)	2.73 (0.666)	0.20 (t=1.32)
Zoom (Average)	178.18(29.55)	123(34.23)	0.000017(t=4.21)
Zoom (Max)	30.46 (62.25)	203.73 (61.69)	0.0011 (t=3.99)
Speed (Average)	43.92(5.99)	33.66(8.83)	0.0015(t=3.33)
Speed (Max)	93.46 (19.02)	63.03 (19.40)	0.0017 (t=3.70)
Last Heading	Audio	Silent	Significant
Time	5.71 (1.22)	7.04 (2.11)	0.037 (t=1.87)
Zoom (Average)	331.37 (46.44)	261.16 (51.83)	0.001 (t=3.50)
Zoom (Max)	469.39 (49.21)	437.06 (48.86)	0.045 (t=1.85)
Speed (Average)	90.11(11.78)	73.78(15.95)	0.0047(t=2.843)
Speed (Max)	140.33 (14.09)	128.55 (13.28)	0.031 (t=2.07)

pixels in the silent group (sd=642). Student’s t-test analysis produced  $p=0.58$  ( $t=0.57$ ) meaning no significant difference was observed in the average traveled distances between two groups. No significant difference was observed in our other log data.

The next heading navigation task was to find the second heading in the document. While the timing difference is not conclusive, the navigational patterns do suggest a difference in behaviour. There was no significant difference in the traveled distance between the audio and silent group; average traveled distance in the modes was: audio 10600 pixels (sd=1004); silent 9900 pixels (sd=191). Student’s t-test produced  $p=0.43$  ( $t=1.746$ ).

In the last-heading navigation task, there was no significant difference in the traveled distance between the audio and silent group; average traveled distance in the modes was: audio 50300 pixels (sd=4301); silent 49200 pixels (sd=5869). Student’s t-test produced  $p=0.6$  ( $t=2.07$ ). As with our overview task, the overall performance figures gave the advantage to the audio mode. However, in this task, backtracking was performed by all audio mode participants, whereas eight silent mode participants demonstrated some backtracking. The distribution giving two values of less than five, applying the Chi-squared test is not appropriate.

## 6 Discussion

These studies show that the audio cue improved participants’ performance in locating headings, i.e. participants in the audio group performed significantly

faster and the distance they traveled was shorter than participants in the silent group. Also, two participants in the silent group reported 4 titles and two reported 6 titles. All users in the audio group reported correct number of headings. Participants liked the selected audio cue for SDAZ and found it intuitive.

One participant stated that “I like that you don’t have to read something to know that it’s there.” and another participant commented “I need to pay more attention if I use it without sound; I have to go back and verify if they are correct.”

In locating by heading content there was no significant difference in the traveled distance between the audio and silent group. Backtracking was a common navigational behaviour where participants were asked to search for headings by contents or order. The further down the heading was in the document (e.g. “Tom”, “Kim”), participants in the audio group took more advantage of the zoom feature. The average zoom level was significantly different when the headings were in further distance from the beginning of the document.

The issue of within-document navigation has previously been investigated within digital library research. One such example is the enhanced scrollbar technique suggested by Byrd [4]. This used a visual cue – coloured indicators on the scrollbar – to indicate the presence of words that matched a previous query. This technique could be combined with our method, or our audio cues simply used in place of the visual tiles used by Byrd.

David Harper [11] also uses visual cues to suggest the presence of words matching a query, but Harper’s technique is to provide a visual document profile above the document display that complements the main document view and is separate from the usual scrollbar navigation.

Under evaluation, Byrd’s technique has not yielded conclusive results, and it may not be an effective solution. Harper, on the other hand, did demonstrate some effect. Clearly secondary visual cues can have a positive effect on user navigation, so this supports our hypothesis that supplementary information can give users tacit task support. However, there is insufficient data at present to compare our non-visual approach to the techniques of Byrd or Harper.

The DL community’s research on within-document navigation is, at present, young, and we are not aware of any major work in this domain other than Byrd’s and Harper’s. In general, DL-centred HCI work has focussed on primary task performance, rather than the effect of secondary cues, and thus there is also a lack of more general work to contrast Harper, Byrd’s and our own work against. The techniques used for within-document navigation could well be used to assist users in other tasks. One possible example would be when to support then user when they are triaging document result sets to identify the texts they will actually view in detail.

## 7 Future Work

Our tasks encouraged participants to do skim reading, paying little attention to the content of the document. Although participants in the audio group reported that with audio they cannot miss targets, they did not like overshooting targets

and backtracking. One possible solution is having an indication of a heading coming up or predicating what is coming next. For this we can adopt the predictive model described in [8]. The initial study in that paper suggests that a predictive model could reduce users' overshoots when landing on a target: a common problem in audio/haptic interaction.

Another area for future work is to compare the functionality of different cues. In the previous section, we noted the potential of haptic cues. One more simple comparison is to look at the comparative performance of secondary visual cues. For example, a sign may flash on the screen when a heading comes into view. The likely problem of that technique would be that users' visual attention would have to constantly readapt and refocus between different points – most likely reproducing the existing problems of visual search. However, a systematic understanding of the impact of using secondary visual cues versus the non-visual techniques we have discussed in this paper would be of immense value both within document navigation specifically, and for navigation research generally.

In our laboratory, we have briefly investigated other means of providing cues for navigation. One alternative method that seems particularly promising is haptic, or tactile, feedback. This has the advantage that in a public space, sensory cues can be given that have minimal impact on those sharing the same space. One simple example we developed is used where a user is navigating with a “Shake” input device [18]. A “Shake” device is matchbox-sized, and can control navigation using tilt sensors (e.g. tilting forward may move down a document, and tilting forward move up the text). Shakes also include a vibration device that can be triggered by software. Hence, as the human reader uses the shake to navigate, a vibration is given when a heading moves into view, as an alternative to the sonic cue we used here. In principle, the same method could be used with an adapted mouse, for example.

## 8 Conclusion

Our straightforward addition of sonic cues to document navigation yielded clear benefits for our users. When sound was used, user behaviour measurably changed, and navigation times fell. Users were more able to use the zooming features of SDAZ, and moved rapidly at a high zoom level between points where the audio cue marked the presence of a heading. When the “click” was heard that a heading was in view, they rapidly closed in to view the document in detail. In contrast, without this indicator, movement strategies were more controlled and conservative. Users maintained a low level of zoom, and often resorted to backtracking to ensure that they had not missed anything.

## References

1. Andy Cockburn, J.S., Wallace, A.: Tuning and testing scrolling interfaces that automatically zoom. In: CHI 2005: Proceedings of the SIGCHI conference on Human factors in computing systems, pp. 71–80. ACM Press, New York (2005)

2. Buchanan, G., Loizides, F.: Investigating document triage on paper and electronic media. In: Kovács, L., Fuhr, N., Meghini, C. (eds.) *ECDL 2007*. LNCS, vol. 4675, pp. 416–427. Springer, Heidelberg (2007)
3. Buchanan, G., Owen, T.: Improving skim reading for document triage. In: *Proceedings of the Symposium on Information Interaction in Context (IiX)*. British Computer Society (2008)
4. Byrd, D.: A scrollbar-based visualization for document navigation. In: *DL 1999: Proceedings of the fourth ACM conference on Digital libraries*, pp. 122–129. ACM Press, New York (1999)
5. Cool, C., Belkin, N.J., Kantor, P.B.: Characteristics of texts affecting relevance judgments. In: *14th National Online Meeting*, pp. 77–84 (1993)
6. Eslambolchilar, P., Murray-Smith, R.: Tilt-based Automatic Zooming and Scaling in mobile devices—a state-space implementation. In: Brewster, S., Dunlop, M.D. (eds.) *Mobile HCI 2004*. LNCS, vol. 3160, pp. 120–131. Springer, Heidelberg (2004)
7. Eslambolchilar, P., Murray-Smith, R.: Control centric approach in designing scrolling and zooming user interfaces. In: Brewster, S., Oulasvirta, A. (eds.) *International Journal of Human-Computer Studies (IJHCS)*, Special issue on Mobility. Elsevier, Amsterdam (2008)
8. Eslambolchilar, P., Murray-Smith, R.: Model-based Target Sonification in Small Screen Devices: Perception and Action. In: *Handbook of Research on User Interface Design and Evaluation for Mobile Technology*, February 2008, pp. 478–506. Idea Group Reference (2008)
9. Eslambolchilar, P., Murray-Smith, R.: Interact, excite, and feel. In: Schmidt, A., Gellersen, H., van den Hoven, E., Mazalek, A., Holleis, P., Villar, N. (eds.) *Second International Conference on Tangible and Embedded Interaction, TEI 2008*, Bonn, Germany, February 2008, pp. 131–138. ACM, New York (2008)
10. Gaver, W.W.: *Auditory Interfaces*. In: *Handbook of Human-Computer Interaction*, 2nd edn. (1997)
11. Harper, D.J., Koychev, I., Sun, Y., Pirie, I.: Within-document retrieval: A user-centred evaluation of relevance profiling. *Information Retrieval* 7(3-4), 265–290 (2004)
12. Hermann, T., Hansen, M., Ritter, H.: Principal curve sonification. In: *Proceedings of International Conference on Auditory Displays, ICAD 2000*, USA, April 2000, pp. 81–86 (2000)
13. Hermann, T., Hunt, A.: The discipline of interactive sonification. In: Hermann, T., Hunt, A. (eds.) *Proceedings of the International workshop on interactive sonification*, Bielefeld, Germany (January 2004)
14. Igarashi, T., Hinckley, K.: Speed-dependent automatic zooming for browsing large documents. In: *UIST 2000: Proceedings of the 13th annual ACM symposium on User interface software and technology*, pp. 139–148. ACM Press, New York (2000)
15. Liu, Z.: Reading behavior in the digital environment. *Journal of Documentation* 61(6), 700–712 (2005)
16. Sun, L., Guimbretière, F.: Flipper: a new method of digital document navigation. In: *CHI 2005: extended abstracts on Human factors in computing systems*, pp. 2001–2004. ACM Press, New York (2005)
17. Warren, N., Jones, M., Jones, S., Bainbridge, D.: Navigation via continuously adapted music. In: *CHI 2005: CHI 2005 extended abstracts on Human factors in computing systems*, pp. 1849–1852. ACM Press, New York (2005)
18. Williamson, J., Murray-Smith, R., Hughes, S.: Shoogle: excitatory multimodal interaction on mobile devices. In: *CHI 2007: Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 121–124. ACM Press, New York (2007)