A Field Experiment of Spatially-Stable Overviews for Document Navigation

Carl Gutwin[1], Andy Cockburn[2], and Nickolas Gough[1]

[1]Department of Computer Science
University of Saskatchewan
Saskatoon, Canada
gutwin@cs.usask.ca

[2]Computer Science and Software Engineering
University of Canterbury
Christchurch, New Zealand
andy@cosc.canterbury.ac.nz

ABSTRACT

Finding (and re-finding) locations in text documents is a common activity for most computer users – but tools for document navigation are still limited in many ways. Previous research has shown that a spatially-stable overview of the entire document can be substantially faster than any other navigation technique – particularly when revisiting previous locations. However, the overview technique has only been tested in a limited laboratory study, so little is known about whether it works in more realistic contexts. To answer this question, we developed a PDF viewer that incorporates several document-navigation techniques, and carried out two studies. First, we ran a field experiment in which users carried out search tasks using an overview and other techniques – on their own computers in a non-laboratory environment. Second, we ran a smaller field study in which people used our viewer (with choice of navigation techniques) for their own PDF tasks. In the field experiment, the overview was significantly and substantially faster than other techniques, and in the field study, the technique was frequently used for a wide variety of documents. Our work provides confirmation of the value of spatially stable overviews as a basis for document navigation.

Author Keywords

Document navigation; space-filling thumbnails; scrolling.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI).

INTRODUCTION

Finding and re-finding locations in text documents is a common activity for most computer users. There are many ways for people to navigate in documents – such as scroll bars, scroll wheels, touch-based flick scrolling, rate-based drag scrolling, page up/down keys, arrow keys, thumbnail views, outline views, "go to page number" controls, and "Find" dialogs. Despite this range of techniques, the tools provided in most document viewers have two limitations.

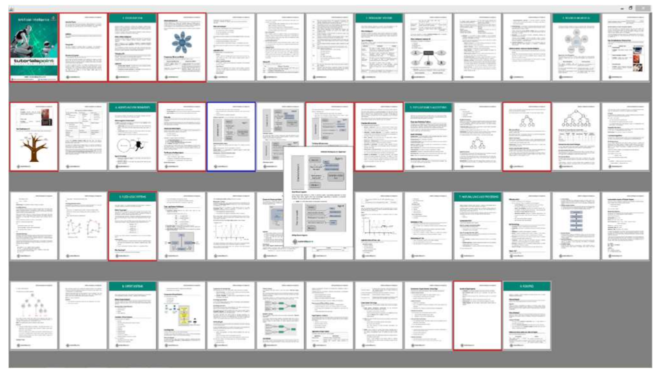

Figure 1. A document overview, showing all document pages at once: the current page is outlined blue, recently visited pages red. Hovering on a thumbnail magnifies the page.

First, it can be difficult to do *visual search for a target page*. In many situations, users remember visual features of a page, rather than its page number or text heading. For example, people may remember a picture on the page (e.g., "the page with the picture of an octopus") or the overall layout (e.g., "the page where the bottom half was empty"). Pages have distinctive "page shapes," and people have considerable perceptual expertise in visual pattern-matching – but current systems do not provide a good way for users to conduct efficient visual search. While some viewers provide a sidebar column of thumbnails that shows 5-10 pages at a time, considerable scrolling is still required.

Second, it can be difficult to *revisit a page*. Going back to previously-visited locations is a common activity (often accounting for a large percentage of page visits [2, 3]). In these situations, users often think of the page in terms of their visitation history (e.g., "the page I was just at"). Current document viewers, however, do not record interaction history, so there are no tools to help users revisit recent (or frequent) pages. Although research on interaction history is now more than 25 years old (e.g., [16]), and although it has been implemented in research prototypes [4], real-world viewers do not make use of this valuable information.

Previous research suggests that thumbnail overview interfaces can address these problems. These interfaces provide two views of a document – an ordinary page view,

CHI 2017, May 06 - 11, 2017, Denver, CO, USA

ACM 978-1-4503-4655-9/17/05…$15.00
DOI: http://dx.doi.org/10.1145/3025453.3025905

and a spatially-stable full-screen overview of the entire document (e.g., Figure 1). The user brings up the overview with a control action (e.g., a right-click); to go back to the page view, the user clicks on any of the thumbnails in the overview (and that is the page shown in the page view).

Overviews have several valuable properties for document navigation. An overview treats a document as a 2D map, whereas scrolling is based on the metaphor of a long 1D piece of paper. The overview therefore provides an *absolute* representation of the document, rather than a relative one. The two key advantages of absolute representations are that they are inherently stable (which facilitates formation of spatial memories), and that they permit direct access to document regions using predictable methods (e.g., clicking on a page thumbnail that lies at a consistent location). Relative document navigation, in contrast, is inherently *un*stable, impeding formation of a spatial document model [22, 23]; in addition, the mechanisms used to acquire specific document regions are unpredictable because they can vary depending on the current document location.

Overviews support visual search because they concurrently show thumbnails of all pages in the document, arranged in a grid. The overview presentation preserves the main visual features of each page (including color, layout, pictures, and whitespace), enabling visual search based on "page shapes." As shown by Figure 1, pages are clearly differentiable even in miniature. Overviews also implicitly support revisitation because of their spatially-stable arrangement – a particular page is always in the same place in the grid, so users can go directly to a page by clicking on its position in the overview.

More than ten years ago, a previous implementation and study of the overview representation – called Space Filling Thumbnails (SFT) – showed the advantages of the overview approach [7]. The overview technique was faster than six other navigation methods, particularly when revisiting pages. Although the previous study provided strong evidence of the value of SFT, the study was limited in four ways: first, it was conducted in a lab setting, which can lead to artificial results; second, it only used mock-ups of the different views, rather than implementing the technique in a real viewer; third, it used tasks that only revisited locations once, whereas more frequent revisitation is common; and fourth, it did not explore the reasons why the overview was faster than the main competitor (a thumbnail sidebar).

To provide a second and more thorough investigation of document overviews, we carried out two studies that examined overviews in more realistic settings. We added several navigation techniques to an open-source PDF viewer, and carried out two studies with 25 people who installed our document viewer on their own computers for one month. The first study was a field experiment [19] which asked people to carry out controlled tasks, but in their everyday computing environment. Participants searched for target pages using one of three interfaces: scrolling (i.e., the scrollbar, scroll wheel, and keyboard navigation), scrolling plus a thumbnail sidebar, and the overview as described above. We used documents of 10, 50, 100, and 200 pages.

The field experiment showed that overviews were significantly faster overall for the three longer documents (50, 100, 200 pages), with a substantial advantage when tasks involved page revisitation. Overall, search times with the overview were 11.3 seconds, compared to 18.7 seconds for scrolling with thumbnails, and 28 seconds for scrolling. In addition to these quantitative results, we also gathered qualitative data about how and why the techniques differed.

Our second analysis was a small field study that ran concurrently with the field experiment. We asked some of the participants to use our custom PDF viewer for their own documents and their own tasks, and we logged their interaction with these documents. In this study, participants could choose to use any of the navigation mechanisms built into the viewer. The log data provides additional evidence that overviews can be a successful technique for real tasks: people used the overview for 31% of the more than 550 page visits, showing that the technique is feasible for real use.

Overall, our studies provide a useful triangulation of earlier laboratory evaluations. Our results add new evidence about the effectiveness of spatially stable overviews, their ability to work in a broad range of usage environments, the underlying design principles that can explain their performance, and their feasibility for real-world tasks.

RELATED WORK

Two main areas of research have influenced our work on spatially stable document overviews: improved methods for scrolling, and spatially stable interfaces.

Improved methods for scrolling

Scrolling has been a fundamental part of interaction since the first visual editors [30], and a wide range of input and output methods have been developed to improve user performance and experience in moving through documents. The comparative efficiency of alternative scrolling methods will be influenced by the nature of the user's navigational task. It is therefore unlikely that one method will be the best for all scrolling activities. For example, if the user finds a particular page number in an index, the most efficient method for moving to that area in the document might be to type the target page into a "go to" field. Alternatively, if the user needs to center the document in the page for reading comfort, the scroll wheel may provide the most convenient method; and manipulating the scroll-thumb might be a suitable method for moving to the start or end of the document.

Hill and Hollan [16] proposed adding "edit wear and read wear" marks into the otherwise blank scroll-trough, which would allow users to see the document regions that are most frequently used. Alexander et al. [4] implemented and evaluated a related method in their 'footprints' scrollbar, which left marks in the scroll-trough to demark regions most recently visited by the user. Their evaluation showed that when the system correctly inferred the document regions that were of interest to the user, footprint marks significantly reduced scrolling time by providing direct shortcuts to document regions. However, they offered little assistance to users when their target region was not selected for explicit marking. Wightman et al. [31] described a related 'TouchMark' system for easing document navigation on touchscreen devices, including support for revisitation.

Many researchers have examined input methods that allow faster and more accurate control over scroll-based target acquisition. The method used to evaluate these techniques typically involved showing a target destination to the user (via highlighting in the scroll-trough and in the document itself), then asking them to move as quickly as possible to the target line in the document. Hinckley et al. [17] showed that Fitts' law [13] accurately models scroll-based target acquisition of this form, with both position-control input devices (such as a scroll wheel) and with rate-control input (such as an isometric joystick, or middle mouse-button dragging). More recently, researchers have scrutinised the transfer functions that translate user actions with an input device into scrolling outcomes [24]: Cockburn et al. [8] proposed and evaluated a transfer function that adapted the scroll wheel's gain based on document length; Quinn et al. [25] examined touch-scrolling transfer functions; and Aceituno et al. [1] examined edge-based transfer functions in which scroll speed depends on the distance that the user drags beyond the window edge.

Spatial memory

When users interact with physical or electronic documents, they build up a spatial understanding of the location of document features [22, 23]. While people find ways to mark salient document regions when using physical documents (e.g., folding a dogear into a page, or using a finger as a temporary bookmark), analogous methods are relatively poorly supported in electronic documents.

Space-Filling Thumbnails (SFT) was proposed to improve users' ability to form and exploit spatial memory for document regions [7]. The technique allows users to toggle between two full-window views of a document – a view that displays a single document page, and a thumbnail view that shows all pages in the document in spatially stable locations. Clicking on a thumbnail immediately shows that page in the main view. The thumbnail view is accessed by pressing a control key. A lab-study evaluation showed that SFT was faster than a variety of other scrolling methods, particularly when tasks involved revisitation.

Although not directed at document navigation, CommandMaps [26, 27] use a similar approach to provide rapid access to all of the commands in an interface via a spatially stable presentation of all commands at once. Traditional interfaces divide commands between different menus or toolbars, reducing a user's ability to form and exploit spatial memory (due to spatial overloading) and requiring extra mechanical actions to traverse through the command hierarchy. Evaluations demonstrated that users quickly learn the correspondence between spatial locations and commands, enabling CommandMaps to reduce command selection times compared to other methods. The studies of CommandMaps progressed from an initial lab-based study of a prototype system [28] to a field study of an actual system used for realistic work.

DESIGN OF DOCUMENT-NAVIGATION OVERVIEWS

The main design objectives proposed for Space-Filling Thumbnails [7] were spatial stability and the visual display of all pages. We summarize these here and add several other factors including the problem of changing window size, how to best support visual recognition, the differences between overviews and general zoom capability, and the choice of a full-screen overview rather than a multi-window view. These design issues guided our implementation of a fully-functional PDF viewer (described below).

Spatial stability and changing window sizes

Previous research clearly shows the value of stability in the development of location memory. When data organization changes from task to task, the time to find an item increases linearly with the number of candidate items, but when data locations are stable, performance quickly improves to a logarithmic function of the number of items [6]. This change occurs because once items are well known, their locations can be retrieved from memory, which follows the logarithmic Hick-Hyman law of choice response [6, 18, 29].

Maintaining perfect spatial stability, however, is almost impossible in a real-world document viewer – for example, when the user moves the viewer to a different monitor, they translate the entire overview to a new location; when they resize the window, the original overview may no longer fit into the available space.

Previous work on this issue, however, provides us with an appropriate design path. Scarr and colleagues studied the resilience of spatial memory to various transformations (i.e., once a spatial arrangement is learned, what is the effect of translating, scaling, or rotating the arrangement) [28]. They found that only rotation caused any substantial loss of performance, suggesting that moving the viewer application within the user's display environment will not be a major problem. Scarr and colleagues also looked at the issue of changing window aspect ratio – they examined user performance after "reflowing" items into a new arrangement, and compared it with performance after scaling the arrangement to fit the new window (maintaining relative spatial locations). The results of this test clearly showed that scaling was much better because it allowed users to continue to use their spatial memory.

As a result, we maintain the layout of our document overview across window-size changes. The grid is set up based on the user's window size when the document is first opened, and these grid dimensions are kept stable; if necessary, the thumbnails are scaled to be able to present the grid in the available space. Of course, the user can request that the grid be re-created, (e.g., if they are switching permanently from a portrait to a landscape monitor setup).

Finally, the overview presents a rich visual space that provides a wide variety of potential landmarks for the user's memory – the graphical elements in a document overview gives a much more memorable visual field than other representations (such as scrollbars).

Supporting recognition and search with "page shapes"

As described above, a main way that users look through documents is by searching for remembered visual characteristics of a page. The SFT system [7] provided thumbnails that showed the overall layout of the page, but small details were often lost – leading to questions, for example, about whether the overview could be used for visual pattern-matching in documents longer than 200 pages.

With advances in display resolution, however, it is now possible to provide high-quality thumbnail images that allow much more detailed visual recognition. For example, Figure 2 shows the actual size of a thumbnail from a 100-page document; on a high-DPI laptop screen such as Apple's Retina Display, it is relatively easy to identify the details of the diagram on the page, and to read the main headings. In our experience with overview-based document navigation, these high-resolution thumbnails can be important, because the success of a user's visual search depends on how many other pages have similar visual attributes to the target. When the number of candidate pages is large, high resolution displays can help to by allowing users to use finer-grained detail as additional search criteria.

As with the earlier implementation of SFT, our overview also provides an enlarged pop-up thumbnail, displayed when the user hovers over any page in the overview. The increased size of the pop-up provides an intermediate way for users to inspect a page without committing to the selection – particularly useful when candidate pages have been identified by a more general visual scan. The pop-up size can be configured by the user, but is not dependent on the document length – therefore, precise page details can be seen even with long documents.

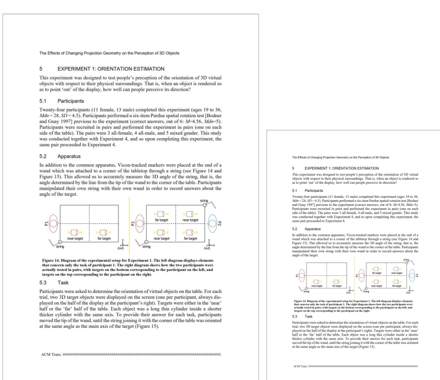
The Effects of Changing Projection Geometry on the Perception of 3D Objects

5 EXPERIMENT 1: ORIENTATION ESTIMATION

This experiment was designed to test people's perception of the orientation of 3D virtual objects with respect to their physical surroundings. That is, when an object is rendered so as to point 'out' of the display, how well can people perceive its direction?

5.1 Participants

Twenty-four participants (11 female, 13 male) completed this experiment (ages 19 to 36, *Mdn* = 28, *SD* = 4.5). Participants performed a six-item Purdue spatial rotation test [Bodner and Guay 1997] previous to the experiment (correct answers, out of 6: *M*=4.56, *Mdn*=5). Participants were recruited in pairs and performed the experiment in pairs (one on each side of the table). The pairs were 3 all-female, 4 all-male, and 5 mixed gender. This study was conducted together with Experiment 4, and so upon completing this experiment, the same pair proceeded to Experiment 4.

5.2 Apparatus

In addition to the common apparatus, Vicon-tracked markers were placed at the end of a wand which was attached to a corner of the tabletop through a string (see Figure 14 and Figure 15). This allowed us to accurately measure the 3D angle of the string, that is, the angle determined by the line from the tip of the wand to the corner of the table. Participants manipulated their own string with their own wand in order to record answers about the angle of the target.

Figure 14. Diagram of the experimental setup for Experiment 1. The left diagram displays elements that concern only the task of participant 1. The right diagram shows how the two participants were actually tested in pairs, with targets on the bottom corresponding to the participant on the left, and targets on the top corresponding to the participant on the right.

5.3 Task

Participants were asked to determine the orientation of virtual objects on the table. For each trial, two 3D target objects were displayed on the screen (one per participant, always displayed on the half of the display at the participant's right). Targets were either in the 'near' half or the 'far' half of the table. Each object was a long thin cylinder inside a shorter thicker cylinder with the same axis. To provide their answer for each task, participants moved the tip of the wand, until the string joining it with the corner of the table was oriented at the same angle as the main axis of the target (Figure 15).

ACM Trans. ########.

Figure 2. Actual-size page thumbnail (left, scaled for 100-page document on 28" monitor; right, for 225-page document). On high-DPI displays, page details and even larger text is visible.

Another new feature in our view is that it marks recently-visited pages with a red border, as shown in Figure 1 – this is a form of "read wear" described earlier [16]. In our implementation, we mark both the pages that have been inspected using the pop-up (in orange) and pages that have been selected (in red). The highlights are cleared once the user leaves the overview for more than one minute (these times can be set by the user).

Fixed overviews vs. arbitrary zoom

Our approach relies upon two fixed levels of zoom – the overview (i.e., zoomed out all the way), and the main page view (i.e., zoomed in all the way). By hard-coding these two levels into the design, we gain important simplicity that allows performance improvements (other controls remain available to configure zoom level of the page display). Variable levels of zoom provide useful utilities for the user, but they address a different purpose to the overview (such as scrutinizing a particular document feature). Providing a shortcut to a fully-zoomed-out view (i.e., the overview) is intended to promote the benefits described earlier – enhanced visual search and improved revistitation through support for spatial memory formation and exploitation.

Full-screen overview vs. multiple views

Some systems provide semantic overviews as secondary views beside or inset into the main page view. This is a useful representation for some tasks – in particular, the task of understanding your current location in relation to the overall structure of the document. However, there are two reasons why a multiple-view solution would be less effective in a document-navigation system.

First, the screen real estate available for document thumbnails is a critical resource that can greatly change the success of visual search based on page shapes. An inset overview would not be able to provide the level of detail needed for visual pattern-matching in longer documents. Second, our observations of document navigation suggest that there are few if any user tasks that require both the detailed view and the overview at the same time. It is more likely that users want to carry out a navigation action (e.g., choose a new location for inspection), and then carry out a task that uses the detail view (e.g., read text on the page). Given the serialized nature of these tasks, it makes sense to multiplex the available window space and give maximal area to the overview.

Overviews as a base representation

One of the main advantages of absolute and stable document representations is that their structure and meaning can be learned and relied upon. As mentioned above, this enables people to use spatial memory as a shortcut to finding previously-visited pages, and provides a landmark-rich substrate on which memory can be anchored. Once the base representation is learned, however, other information can be encoded in the overview. Our visualizations of the user's current page and recent pages are an example, and other information is also possible (e.g., in a collaborative document space, the overview can show others' locations).

Working together with other document-navigation tools

Different document-navigation tools provide solutions to different kinds of user tasks. We are not suggesting that overview-based navigation will solve all search needs – rather, that overviews can support certain common tasks that are currently unsupported. Overviews do not provide support for local and linear navigation (e.g., reading across a page boundary), and we do not suggest that users should go to the overview to move by a single page.

The overviews presented here, however, need not displace other navigational techniques – users are already capable of understanding several different methods for moving around in documents, and it is likely that overviews could be added to current document viewers without disrupting current practices. In addition, other methods for improving existing techniques' support for tasks like revisitation (e.g., marks in the scrollbar to indicate recent visits [4]) can be used in conjunction with an overview approach.

A CUSTOM PDF VIEWER WITH OVERVIEW NAVIGATION

We built a custom viewer based on the open-source IcePDF Java PDF system (www.icesoft.com/icepdf/). The viewer is shown in Figures 1 & 3 (left), and contains most of the main features of commercial PDF viewers (including zooming, rotating, annotating, and printing). The existing viewer also supports several types of document navigation: scrolling (with scroll bar and scroll wheel), keyboard navigation (arrow keys and PageUp/PageDown keys), a Find dialog, and a basic thumbnail sidebar view (Figure 3).

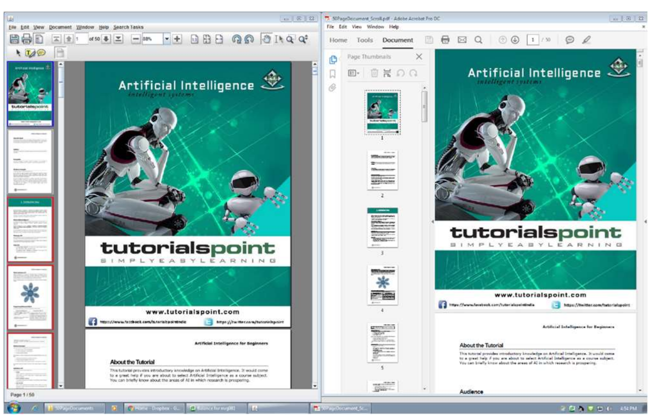

Figure 3. Thumbnail sidebar view in our viewer (left) and the Adobe Acrobat viewer with default thumbnail size (right).

We added an overview navigation system to the basic viewer, and we also enhanced the thumbnail sidebar view. The overview provides a full-screen grid of thumbnails that shows the entire document, as discussed above. Users show the overview by pressing the right mouse button; once the overview is visible, users can go back to the page view either by left-clicking on a page (which opens that page in the main view) or right clicking (which returns them to the page they were on when they opened the overview). In the overview (Figure 1), the current page is highlighted with a blue border, and recently-visited pages are highlighted with a red border. Pages that have been recently inspected (with the enlarged popup shown on hover) have an orange border.

The thumbnail sidebar view works consistently with commercial viewers – the thumbnails can be scrolled independently of the main view, and clicking on a thumbnail takes the main view to that page. We augmented the existing thumbnail sidebar with the "current page" and "recently visited" border highlights. We also added code to the custom viewer to control the presentation of experimental tasks (described below) and to log all navigation actions (both during experimental tasks and during everyday use).

STUDY METHODS

We carried out two experiments with the custom PDF viewer – a field experiment that involved controlled manipulation of conditions in a field setting, and a related field trial in which the PDF viewer was used for the participants' own work, substantially without experimental controls. Here we describe the methods for the field experiment (the setting and apparatus for the field trial was the same). Our main goal in the field experiment was to examine the performance of overview-based document navigation in comparison to regular scrolling and thumbnail-enhanced scrolling – when installed on a wide variety of computers and monitors in people's ordinary usage environments. A second goal was to find out more about the reasons for the performance characteristics of different navigation methods.

Participants, Procedure, and Apparatus

24 participants were recruited by e-mail from two university communities (9 women, 15 men, mean age 28.1 years). All participants were highly experienced with mouse-and-windows computers (more than two hours use per day) and with document viewers (more than 30 minutes use per day).

Participants joined the study by visiting a web page with instructions and links to an informed consent form and a demographics questionnaire. Participants downloaded and installed our custom PDF viewer on their own computer, and sent their e-mail address to us for communication during the study. During the following month, we sent messages to the participants asking them to carry out tasks with the viewer.

When asked to carry out an experimental task, participants selected that task from a "Search Tasks" menu in the viewer. Each task loaded a specific document with a specific length (included with the viewer) and asked the participant to find a series of target pages using one of the three study interfaces (Scrolling, Scrolling+Thumbnails, or Overview). Target pages were shown at the top right of the screen (at 50% size). The experiment tasks are shown for the three interface conditions in Figures 4-6. The test documents were edited to remove page numbers, and the "Find" functionality was turned off for the search tasks. In addition, the recent-visits highlights were turned off for the Overview and Scroll+

Thumbnails conditions, to prevent these visualizations from providing an easy way to find the targets.

Once the participant found the page, they clicked on it in the main page view, and the system went on to the next search task for that interface. There were four target pages, which were revisited different numbers of times: two pages were revisited four times, and two pages were revisited twice (for a total of 12 search trials per task). Once the task was complete, the participant was directed to visit a web page with a post-task effort survey (based on the NASA-TLX).

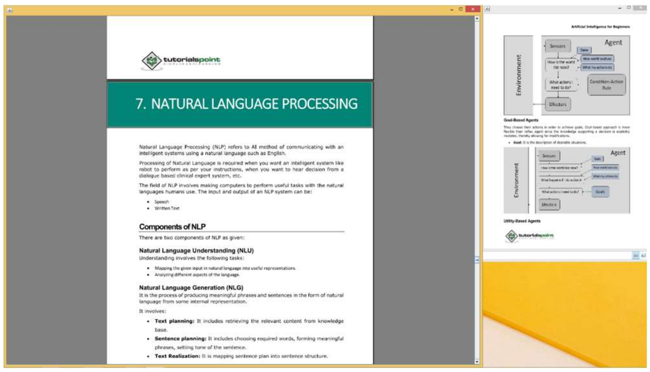

Figure 4. Search task in the Scrolling interface.

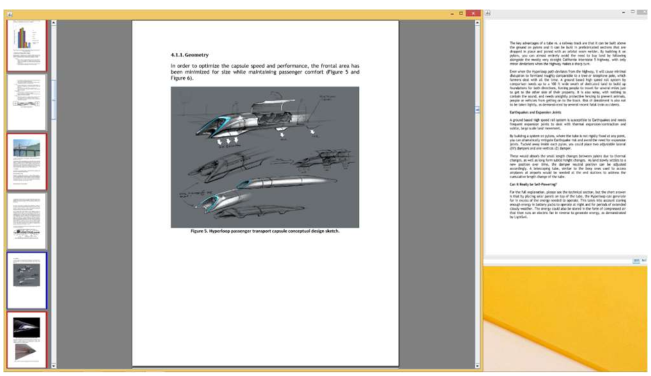

Figure 5. Search task in the Scroll+Thumbs interface.

Figure 6. Search task in the Overview interface; participants could move back and forth from overview to page view.

Participants carried out tasks with each interface and with four different document lengths (10, 50, 100, and 200 pages). The e-mail messages asked participants to carry out the tasks for a particular length, for all interfaces. Participants were allowed to choose the order of the interfaces, and the time of day when they completed the tasks. To ensure that each specific document was used equally with all interfaces, the assignment of documents to interfaces was manipulated remotely during the study. Performance data was recorded by the viewer and sent back to a secure web server for collection and analysis. Participants were asked to do the tasks at approximately one-week intervals; therefore, the four tasks were completed over about one month.

Experimental Conditions

The three interfaces used to carry out the tasks were:

- *Scroll* (Figure 4), using the normal page view and any of the built-in scrolling mechanisms (scrollbar, scroll wheel, PageUp/ PageDown keys).
- *Scroll+Thumbnails* (Figure 5), using both the scrolling mechanisms above plus the thumbnail sidebar. These thumbnails had their own scrollbar.
- *Overview* (Figure 6), using the grid view described above. Scrollbars were removed from the page view, but the scroll wheel and keyboard navigation were available.

Easy vs. Hard Page Shapes

Our experience with document overviews suggested that some pages would be more difficult to find than others when using visual pattern-matching. We therefore chose two types of targets for each search task: “easy” targets had obvious visual characteristics that were not frequently repeated among the pages of the document; “hard” targets had visual characteristics that were common to many pages. Figure 7 shows examples of easy (top row) and hard (bottom row) targets: the targets are the leftmost image, and representative other pages from the document are shown to the right.

Figure 7. Above: “easy” target (leftmost) with representative other pages from the document to its right. Below: “hard” target (leftmost) and other pages to its right.

Study Design and Data Analysis

The in-the-wild nature of the field experiment meant that some participants did not complete all of the tasks. The incomplete mapping between participants and conditions prevents the creation of a standard linear model for ANOVA. In order to make use of all of the participants’ data, we used a Hierarchical Mixed Model (HMM), which allows a similar analysis but accommodates missing data from some participants [10]. The analysis was carried in R using the *nlme* library, using participant as a random factor.

The study therefore used a mixed between/within repeated-measures design, with three factors:

- *Interface*: Scrolling, Scroll+Thumbnails, Overview
- *Document Length*: 10, 50, 100, or 200 pages
- *Visit*: first, second, third, or fourth
- *Target Difficulty*: "easy" or "hard" as described above

The dependent measure was search completion time; for subjective responses, the measures were responses to the NASA-TLX.

Participants' Computing Environments

Participants were asked to provide details of their computing environments. 16 stated that they completed tasks using a desktop computer+monitor arrangement, 5 reported using a laptop, 2 a laptop+monitor, and 1 gave no answer. Three types of operating system were reported, with 9 using some version of Microsoft Windows, 4 using Mac OS X, and 8 using a version of Linux; 3 gave no answer. Display sizes ranged from 16 to 28 inches. Several display resolutions were also reported: 1440×1080 (2), 1680×1050 (7), 1080p (9), 2800×1800 (2) and 3200×2000 (1); 3 gave no answer.

Variations in hardware setup, and other demographic variables such as age, sex, and computing experience, had no discernable effect on the results reported below.

RESULTS (FIELD EXPERIMENT)

Task Performance

Overall, there were substantial differences in search time for different interfaces, different document lengths, and different visits to a target. Analysis using the Hierarchical Mixed Model (HMM, see Data Analysis above) showed strong main effects of *Interface* ($F_{2,2303}=15.62$, $p<.0001$), *Length* ($F_{3,2303}=36.20$, $p<.0001$), *Visit* ($F_{3,2303}=13.36$, $p<.0001$), and *Difficulty* ($F_{1,2303}=7.13$, $p=.0076$) on search time.

These data are shown in Figures 8-10. Overall, the Overview was fastest (11.28s per search, s.d. 19.2), followed by Scroll+Thumbnails (18.73s, s.d. 38.4) and Scrolling (28.07s, s.d. 51.2). However, these differences must be interpreted in light of the significant interactions described below.

The HMM showed significant interactions between *Interface* and *Length* ($F_{2,2303}=6.44$, $p=.0016$), between *Interface* and *Visit* ($F_{6,2303}=2.21$, $p=.040$), and between *Length* and *Visit* ($F_{3,2303}=6.54$, $p=.0002$). These interactions are illustrated in Figure 10. For the 10-page document, there is little difference between the interfaces, but the difference increases with document length. In addition, the figure shows that revisitation saves proportionally more time as documents grow longer.

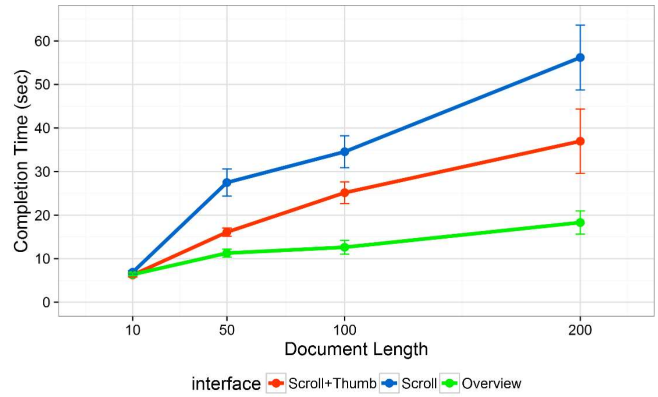

Figure 8. Completion time (±s.e.) by interface and length.

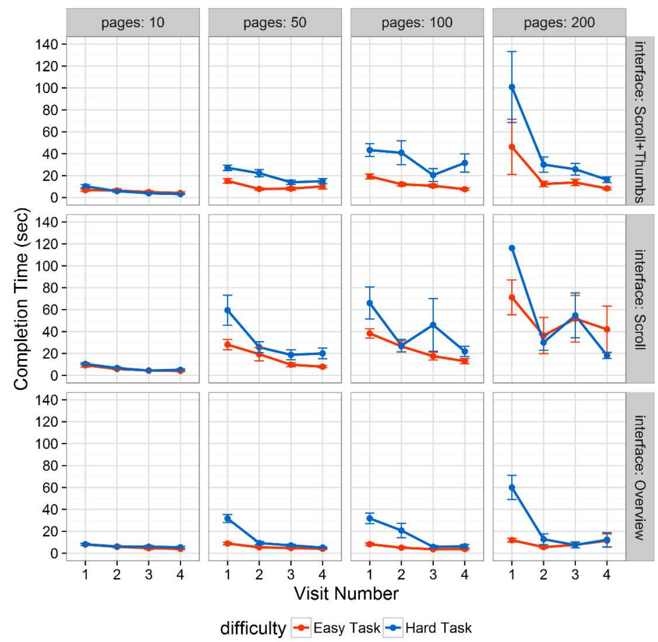

Figure 9. Completion time (±s.e.) by interface (rows), length (columns), visit number (x-axis), and difficulty (color).

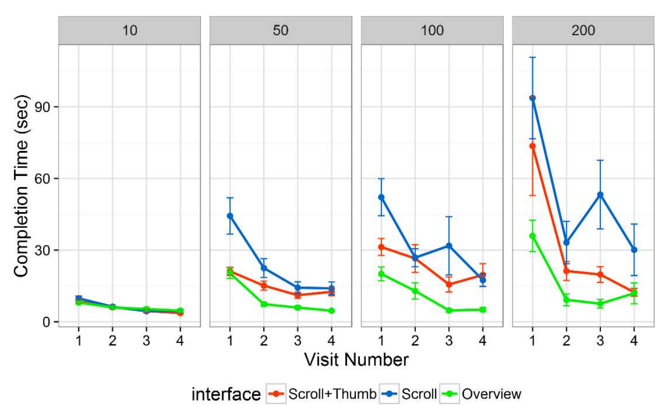

Figure 10. Completion time (±s.e.) by doc length (columns), interface (color), and visit number (x-axis).

Perception of Effort

After each task (interface and document length), participants filled out an effort questionnaire based on the NASA-TLX. Mean response scores are shown in Figure 11. Effort scores approximately follow the performance results above: there are few differences between interfaces for 10-page documents, but larger differences for longer documents. In particular, the Overview is consistently scored as requiring lower effort and as causing less frustration than Scrolling or Scroll+Thumbnails.

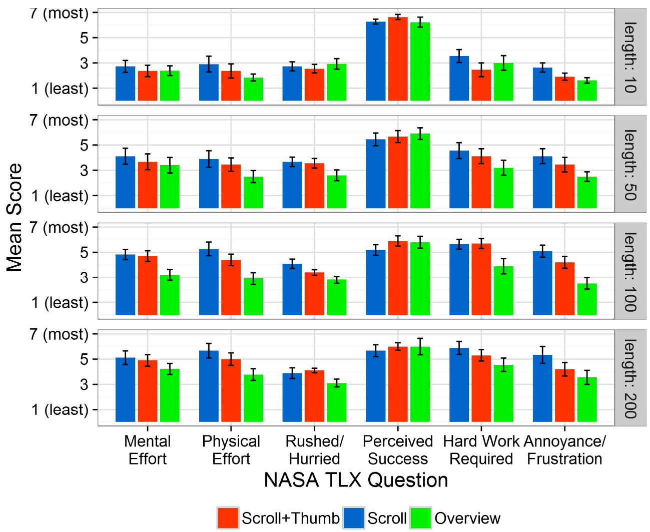

Figure 11. Mean TLX Scores (±s.e.) by interface and length.

Search Strategies and Participant Comments

Participant comments and discussions over the course of the study revealed strategies used with the different interfaces, and difficulties in these conditions.

Scroll+Thumbnails supports page matching, but helps less with revisitation

Participants had mixed comments about the Scroll+ Thumbnails interface. First, it was clear that the thumbnails helped users carry out visual matching – for example, one person said: "I used the scrolling thumbnails bar to scroll through pages. I remembered the specific features of a target page, and looked for that feature while scrolling through the interface with mouse." The limited number of thumbnails, however, was seen as a problem: as one person stated, "For the large documents, it was annoying to use the thumbnails because you could see so few of them at one time. It still felt like scrolling to me."

Interestingly, participants used similar language to describe their solution strategies for revisitation searches with Scroll+ Thumbnails, and none of them mentioned using a different strategy when finding pages for the third or fourth time. For example, one participant said as an overall comment about this interface "When there were lots of pages, it took a while to find the desired page;" another stated "I found it difficult with large files and specially, in target pages where there were limited features to focus on."

Scrolling poorly supports page matching and revisitation

Participants were consistently negative about their experiences with the regular Scrolling interface. There were several expressions of frustration with this method for the longer documents – for example, "When it came to the 100 and 200 pages experiment, it was hard to finish since I had to scroll most of the pages in that experiment to find the target page;" "I found it most difficult in large documents when my current location was too far from the target page. It was so frustrating to scroll through so many pages".

Participants mentioned several strategies that they adopted to try and find or re-find pages with the Scrolling interface. Although people tried to use visual pattern-matching (e.g., "I tried to see what could be easier to spot in that page, for instance, picture etc."), this became difficult with longer documents (e.g., "When there were a lot of pages, you sometimes passed the correct page because you were scrolling fast in an attempt to get through it quicker"). Several participants also stated that they attempted to remember the location of the scroll thumb for a particular page (e.g., "I also tried to keep track of a page's location in the scroll bar. It took me close to target page then I used scrolling"). Other participants mentioned using the background image on their desktop as a marker to help them remember these locations – as described below, maintaining this memory of scrollbar locations was sometimes difficult.

In some cases, participants were forced to use the content of the document in order to help them find a target. For example, one participant said that "Most documents had chapters or sections, so I tried to remember which section each page was in." This strategy was not mentioned for either of the other techniques.

It is difficult to remember page locations in linear views

As mentioned above, several people stated that they tried to remember the location of the scroll thumb when they found a page for the first time. This was not easy to do, since there are few landmarks or visual textures in the scrollbar to help form an accurate memory (as the participant above noted, this strategy was only useful as an approximation). In addition, several participants related instances where they simply forgot to fix the location in their memory before going on to the next task – and once the scroll thumb had moved, the opportunity to remember the previous location was lost. This problem occurred both in the main page view and in the thumbnail sidebar.

Overview supported both pattern-matching and revisitation

Comments were generally very positive for the Overview representation. Participants clearly indicated that they were able to carry out both visual pattern-matching and spatial-memory-based revisitation. Commenting on their strategy for finding a page for the first time, one person stated that the Overview "was super easy and comfortable. I did not have to scroll through lots of pages just to find one page. When in this interface, all the pages were shown at a time, which actually saved me a lot of browsing time." Another person said "you can see all the pages in one view" as the reason why this interface was effective for finding targets.

Participants also recognized that they were able to remember page locations for revisitation, instead of having to carry out another visual search. For example, one person remarked on the value of the overview's richer visual representation: "I remembered the relative position of the target page with reference to other pages, and notable features like colors, image, bold texts etc. to quickly revisit that page." Another participant said that they remembered the pages based on their "physical position in a two dimensional plane."

"Hard" targets are difficult, and require different strategy

Targets where the page shape of the target was not unique in the document were seen as more difficult for all three interfaces – for example, participants disliked targets where there were "limited features to focus on," such as targets that were "just a page full of text."

When faced with these targets, participants stated that they had to adapt their strategy for visual page search. Whereas with an "easy" target people could look for the most obvious visual features of the page, "hard" targets required people to consider a second level of visual detail. For example, if a target page looked superficially similar to many pages, participants had to start looking for other features – such as where bold words appeared in the paragraph, the length of a paragraph, or the shape of the right edge of the text block.

RESULTS FROM FIELD TRIAL OF EVERYDAY USE

We asked participants from the field experiment to also try using the custom PDF viewer for their everyday document viewing in addition to the experimental tasks. There was no restriction or requirement on how they should use the tool, and all interfaces were available at all times. Fourteen of the study's participants used the system for at least one of their own tasks over the one-month period. A summary of this use is shown in Table 1.

Table 1. Summary of non-experiment use of the viewer

Document Length	# Opened	# People	Total Page Visits	Scrollbar, Wheel, Keys	Overview	Thumbs
1..9	38	11	202	134	68	0
10..19	14	5	74	59	15	0
20..49	7	4	45	40	5	0
50..99	8	4	117	72	36	9
100..199	6	4	90	57	33	0
200..250	1	1	28	13	15	0
	74	**14**	**556**	**375**	**172**	**9**

Although this is only a small amount of data, some initial patterns of use emerged that could be followed up in a larger field trial. First, people did use the overview interface – the fraction of page visits accounted for by the overview interface was higher than expected (31% of the 556 page views). This total may be artificially high due to participant interest in a novel interface – but given the low rates of adoption that are often seen in studies where users are already familiar with an existing interface, we see this willingness to use the overview as a positive sign. In addition, people used the overview in all document lengths – although we were surprised that there was not a greater proportion of overview use in long documents. Third, there was very little use of the thumbnail sidebar (the nine page visits were all from one user).

The real-world use of the viewer, however, showed that there are a wide variety of document tasks that people carry out, and that these affect the interface features that people use. For example, several participants mentioned that they often opened documents just to print a particular page, or to cut and paste a quotation – and these tasks did not involve the kinds of search and revisitation that we included in the earlier study tasks, and were usually easy to accomplish using only the standard scrollbar and scrollwheel. In addition, one participant also noted that when he opened a document to read it, the linear document model worked well, and he did not need to navigate in a random-access fashion. When participants did use the overview, however, they saw it as valuable – one participant stated that he used the overview almost exclusively for his own work, but only after realizing its benefit from the study tasks. Finally, three participants mentioned their use of the read-wear marks (and one commented that these would have been very helpful in the study tasks where revisitation was more frequent).

The main practical conclusion that we draw from the everyday use of the system is that the overview interface worked well enough for people to be able to use it for a large number of page visits. This shows that the approach is feasible and usable in real-world use over a wide variety of implemented systems and real document lengths.

DISCUSSION

Results of the field experiment and field study support the findings of the SFT lab study conducted a decade ago [7]. The overview allowed users to complete document navigation tasks much more quickly than other methods, particularly when revisiting document regions.

Importance of replication and experimental triangulation

Research in most scientific fields is incremental. Scientists carefully replicate their colleagues' studies, seeking to understand their generality and boundaries of application. Computer Science is a somewhat anomalous field in that the word '*incremental*' is often used as a criticism – the implication is that instead, we should be seeking paradigm shifts in our research; an objective that Meyer [20] characterized as a 'mania'.

Our results agree with those of the earlier lab study, and they did so even though participants were using their own machines in their own homes (or offices). Prior to conducting the field study, the agreement between studies was far from certain. For example, Fitchett et al. [11] conducted a field evaluation (like ours) seeking to replicate findings of an earlier lab study [12]. Yet their findings differed substantially – an interface feature that had large performance advantages and strongly positive subjective assessments in the lab was more-or-less ignored when evaluated in the field. Many other researchers have noted concerns regarding discrepancies between experimental outcomes when investigating the same interfaces using different methodologies [5, 9, 21], leading to recommendations for methodological triangulation [14, 15].

Future empirical work might seek further triangulated evidence of SFT's performance when different experimental methods are used to cue the users' page retrieval. Our field study, like previous lab studies of SFT, used images of the target page to the retrieval, and these images were continually available to the user during the task. While we believe this task cuing method is appropriate (particularly for page revisitation, where users have a memory for the visual shape of page content), method triangulation could provide further evidence of the techniques' relative strengths.

Why does SFT work?

Traditional scrolling mechanisms (with the exception of 'goto page') are all *relative* means for attaining a target region in a document. Therefore, the user's input to acquire any particular document region differs dependent on where they initiate their action. The scroll thumb does provide an absolute *depiction* of the location of the current view within the document (e.g., when the thumb is in the middle of the scroll trough, the view is at the middle of the document). However, it is difficult to use this depiction for absolute location selection – the behavioral outcome of clicking on any particular location in the scroll trough depends on the relative location of the thumb (the document is scrolled towards its start if the thumb is below the click; and towards its end if the click is above). Even zoom controls are predominantly relative (e.g., control-scroll wheel, or zoom-in/out keys and buttons).

We believe that the dominance of relative methods for document navigation causes the problems that users have in forming an understanding of the spatial location of document features, as observed in prior studies [22, 23].

Overviews, in contrast, provide a consistent and *absolute* method for viewing and acquiring document features. A user who wishes to navigate to a particular page can display the overview, find the target page, point to it and click on it at a certain x, y coordinate in the view. For any subsequent navigation to that page, clicking in the same x, y location will retrieve the same page. This facilitates rapid spatial decisions, and predictable anticipatory actions by the user.

Deployment issues

Although the results for the overview are promising, there are some important issues to understand and address prior to any practical deployment. First, like the previous lab studies of SFT, our studies only examined static documents in which content does not change. Real documents, in contrast, can change in length as new pages are added or deleted. Further studies are needed to examine how performance with overviews degrades when editing changes cause a lack of stability in thumbnail locations. We suspect that the combination of consistent spatial vicinity plus visual cues from the thumbnails' appearance will allow overviews to outperform other scrolling methods in the presence of editing changes, but this requires validation.

Second, there are related issues to address in adapting overview presentations to substantial changes in aspect ratios when windows are resized or mobile devices are reoriented (between landscape and portrait display modes). The homescreen icon layout on mobile devices encounters a similar design problem – if the spatial layout is consistent across orientations, users gain from spatial predictability, but the aesthetics may be compromised by blank spaces and malformed icons; conversely, if the spatial layout is inconsistent, icon locations may be unpredictable (slowing performance), but the aesthetics may be improved. As mentioned earlier, Scarr [28] examined related issues, noting strong performance advantages for spatial stability.

Third, a critical deployment issue concerns consistency with existing scrolling behaviors. Obviously, users can only benefit from overviews if they use them, and the benefits increase with level of use because spatial memories develop with experience. However, users are extremely familiar with scroll wheels, scrollbars, and touch-scrolling gestures, and the more users rely on these well-practiced and habitual utilities, the less they will experience the potential benefits of overviews. Overcoming stasis in interaction paradigms is highly challenging, but it can and does occur – for example, the now-ubiquitous scroll wheel was first popularized by the Microsoft IntelliMouse in 1996, more than a decade after the widespread introduction of the mouse; and touch-scrolling (with disappearing scrollbars) was popularized by the iPhone in 2007, more than a decade later. The frequent use of the overview in our small field study, however, can be interpreted as providing an initial indication that the benefits of the overview may help it overcome the "interface inertia" typically seen in real-world deployment.

CONCLUSIONS

Efficient methods of document navigation are essential for effective office work. Current document applications support a broad range of utilities for scrolling and traversing through documents, but the behavior of nearly all of them is relative to the user's current document location. As a result, current scrolling interfaces provide little support for two main ways that users navigate – visual page searching, and revisitation. Prior research and lab-study findings have suggested that document overviews may improve the efficiency of navigation, particularly when revisiting document regions. The research presented in this paper deepens the analysis and understanding of the main design principles underlying document overviews; it also introduces new overview features, and through two field evaluations, it extends and supports prior findings on the efficiency of overviews. Our work contributes to the now-strong case for including overview displays in commercial document applications.

REFERENCES

1. Aceituno, J., Malacria, S., Quinn, P., Roussel, N., Cockburn, A. and Casiez, G. The design, use, and performance of edge-scrolling techniques. *International Journal of Human-Computer Studies 97* (2017), 58--76.
2. Alexander, J. Understanding and Improving Navigation within Electronic Documents *PhD Thesis. Computer Science and Software Engineering*, University of Canterbury, Christchurch, 2009.
3. Alexander, J. and Cockburn, A. An empirical characterisation of electronic document navigation. in *Proceedings of Graphics Interface (GI 2008)*, Canadian Information Processing Society, (2008), 123-130.
4. Alexander, J., Cockburn, A., Fitchett, S., Gutwin, C. and Greenberg, S. Revisiting Read Wear: Analysis, Design, and Evaluation of a Footprints Scrollbar. in *Proceedings of CHI'09: ACM Conference on Human Factors in Computing Systems*, ACM Press, (2009), 1665-1674.
5. Chiasson, S., Biddle, R. and Oorschot, P.C.v. A second look at the usability of click-based graphical passwords *Proceedings of the 3rd symposium on Usable privacy and security*, ACM, Pittsburgh, Pennsylvania, USA, 2007, 1-12.
6. Cockburn, A. and Gutwin, C. A Predictive Model of Human Performance with Scrolling and Hierarchical Lists. *Human-Computer Interaction 24*, 3 (2009), 273-314.
7. Cockburn, A., Gutwin, C. and Alexander, J. Faster Document Navigation with Space-Filling Thumbnails. in *Proceedings of CHI'06 ACM Conference on Human Factors in Computing Systems*, ACM Press, (2006), 1-10.
8. Cockburn, A., Quinn, P., Gutwin, C. and Fitchett, S. Improving Scrolling Devices with Document-Length-Dependent Gain. in *Proceedings of CHI'12: ACM Conference on Human Factors in Computing Systems*, ACM, (2012), 267-276.
9. Duh, H.B.-L., Tan, G.C.B. and Chen, V.H.-h. Usability evaluation for mobile device: a comparison of laboratory and field tests. in *Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*, ACM, (2006), 181-186.
10. Field, A., Miles, J. and Field, Z. *Discovering Statistics with R*. Sage, 2012.
11. Fitchett, S., Cockburn, A. and Gutwin, C. Finder Highlights: Field Evaluation and Design of an Augmented File Browser. in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, (2014), 3685-3694.
12. Fitchett, S., Cockburn, A. and Gutwin, C. Improving Navigation-Based File Retrieval. in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, (2013), 2329-2338.
13. Fitts, P.M. The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. *Journal of Experimental Psychology 47* (1954), 381-391.
14. Greenberg, S. and Buxton, B. Usability Evaluation Considered Harmful (Some of the Time). in *Proceedings of CHI'08: ACM Conference on Human Factors in Computing Systems*, ACM Press, (2008).
15. Hertzum, M. User Testing in Industry: A Case Study of Laboratory, Workshop, and Field Tests. Kobsa, A. and Stephanidis, C. eds. *Proceedings of the 5th ERCIM workshop*, 1999, 59-72.
16. Hill, W., Hollan, J., Wroblewski, D. and McCandless, T. Edit Wear and Read Wear. in *Proceedings of CHI'92 ACM Conference on Human Factors in Computing Systems*, (1992), 3-9.
17. Hinckley, K., Cutrell, E., Bathiche, S. and Muss, T. Quantitative Analysis of Scrolling Techniques. in *Proceedings of CHI'02 ACM Conference on Human Factors in Computing Systems*, (2002), 65-72.
18. Hyman, R. Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology 45* (1953), 188-196.
19. McGrath, J. Methodology Matters: Doing Research in the Behavioral and Social Sciences. in Baecker, R.a.G., J and Buxton, W and Greenberg, S ed. *Readings in Human-Computer Interaction: Towards the Year 2000*, 1995.
20. Meyer, B. Incremental research vs. paradigm-shift mania. *Commun. ACM 55*, 9 (2012), 8-9.
21. Nielsen, C.M., Overgaard, M., Pedersen, M.B., Stage, J. and Stenild, S. It's worth the hassle!: the added value of evaluating the usability of mobile systems in the field *Proc. Nordic conference on HCI*, ACM, 2006, 272-280.
22. O'Hara, K. and Sellen, A. A comparison of reading paper and on-line documents. in *Proceedings of CHI'97 ACM Conference on Human factors in Computing Systems*, (1997), 335-342.
23. O'Hara, K., Sellen, A. and Bentley, R. Supporting Memory for Spatial Location While Reading from Small Displays. in *Extended Abstracts of CHI'99 Conference on Human Factors in Computing Systems*, (1999), 220-221.
24. Quinn, P., Cockburn, A., Casiez, G., Roussel, N. and Gutwin, C. Exposing and understanding scrolling transfer functions. in *Proceedings of the 25th annual ACM symposium on User interface software and technology*, ACM, (2012), 341-350.
25. Quinn, P., Malacria, S. and Cockburn, A. Touch scrolling transfer functions. in *Proceedings of the 26th annual ACM symposium on User interface software and technology*, ACM, (2013), 61-70.
26. Scarr, J., Cockburn, A., Gutwin, C. and Bunt, A. Improving command selection with CommandMaps. in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, (2012), 257-266.
27. Scarr, J., Cockburn, A., Gutwin, C., Bunt, A. and Cechanowicz, J. The Usability of CommandMaps in Realistic Tasks. in *Proceedings of ACM CHI 2014 Conference on Human Factors in Computing Systems*, ACM, (2014).

28. Scarr, J., Cockburn, A., Gutwin, C. and Malacria, S. Testing the robustness and performance of spatially consistent interfaces. in *Proceedings of the 2013 SIGCHI Conference on Human Factors in Computing Systems*, ACM, (2013), 3139-3148.
29. Seow, S.C. Information Theoretic Models of HCI: A Comparison of the Hick-Hyman Law and Fitts' Law. *Human-Computer Interaction 20*, 3 (2005), 315-352.
30. Sutherland, I. Sketchpad: A Man-Machine Graphical Communication System. in *AFIPS Spring Joint Computer Conference*, (1963), 329-346.
31. Wightman, D., Ginn, T. and Vertegaal, R. TouchMark: flexible document navigation and bookmarking techniques for e-book readers *Proceedings of Graphics Interface 2010*, Canadian Information Processing Society, Ottawa, Ontario, Canada, 2010, 241-244.