

Summary Thumbnails: Readable Overviews for Small Screen Web Browsers

Heidi Lam

University of British Columbia
201-2366 Main Mall
Vancouver, BC
hllam@cs.ubc.ca

Patrick Baudisch

Microsoft Research
One Microsoft Way
Redmond, WA 98052, USA
baudisch@microsoft.com

ABSTRACT

In order to display web pages designed for desktop-sized monitors, some small-screen web browsers provide single-column or thumbnail views. Both have limitations. Single-column views affect page layouts and require users to scroll significantly more. Thumbnail views tend to reduce contained text beyond readability, so differentiating visually similar areas requires users to zoom. In this paper, we present *Summary Thumbnails*—thumbnail views enhanced with readable text fragments. Summary Thumbnails help users identify viewed material and distinguish between visually similar areas. In our user study, participants located content in web pages about 41% faster and with 71% lower error rates when using the Summary Thumbnail interface than when using the Single-Column interface, and zoomed 59% less than when using the Thumbnail interface. Nine of the eleven participants preferred Summary Thumbnails over both the Thumbnail and Single-Column interfaces.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General Terms: Human Factors, Design.

Keywords: Web browsing, small screen device, PDA, thumbnail view, overview, semantic zooming.

INTRODUCTION

Web pages are typically designed with the desktop screen in mind, and therefore often use multi-column layouts and preformatted page widths. Such pages can be hard to read on small screens. If rendered as is, the resulting page is typically much larger than the web browser screen and users need to scroll both horizontally and vertically to view it [27].

To avoid the need for horizontal scrolling, the majority of commercially available small-screen web browsers provide a single-column viewing mode that reformats the page by concatenating all its columns, thus displaying it as a single,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2005, April 2–7, 2005, Portland, Oregon, USA.
Copyright 2005 ACM 1-58113-998-5/05/0004...\$5.00.

very long column. Figure 1 shows an example. While this approach tends to work well for helping users read pages, it is of limited use while browsing. Since this approach affects the layout of pages so significantly, users may find it hard to recognize pages familiar from desktop viewing. (See the left side of Figure 3 for a better idea of what this page looked like before the conversion). This display style also significantly increases the required amount of vertical scrolling. As the scrollbar position in Figure 1 indicates, accessing the news story that used to be at a prime location at the top of the page now requires scrolling 8 screens down, past what used to be the menu column of that page.



Figure 1: Viewing the first headline in this single-column browser, requires scrolling 8 screens down.

To reduce the need for horizontal and vertical scrolling and to give users an overview of the page, researchers have proposed displaying web pages as a thumbnail view, i.e., a version of the page that is scaled down to fit the width of the small screen (e.g., [7]). Figure 2a shows an example. In the intended use, users start by viewing a web page in the thumbnail mode, rapidly identify the area of interest, and then zoom into that area for reading. The problem with this approach, however, is that the required size reduction typically renders text on thumbnails unreadable [2], as illus-

trated by the callout in Figure 2a. Due to the lack of readable text, users are often unable to tell similar looking areas apart and thus have to resort to exploring these locations using a tedious zoom-in-and-out or panning strategy.

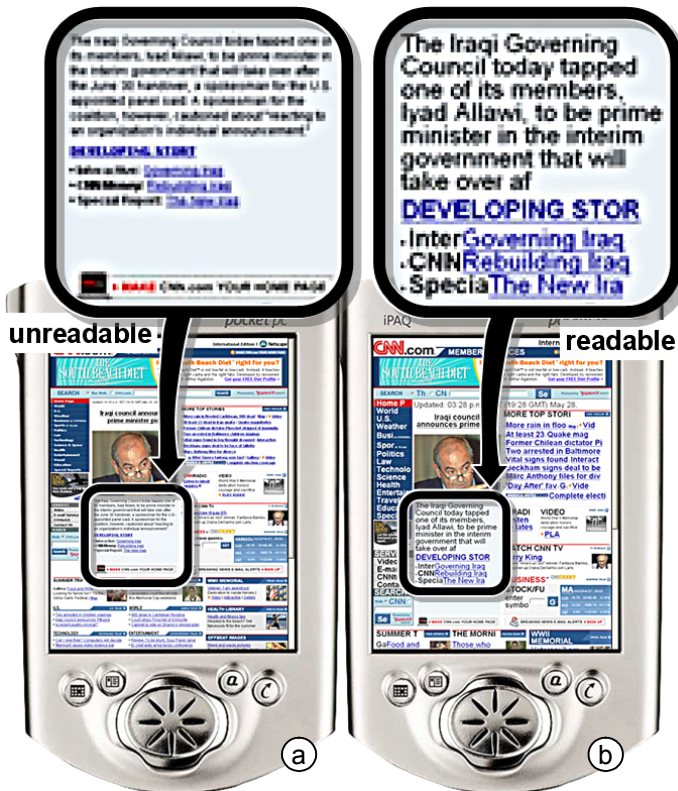


Figure 2: (a) Traditional thumbnail views render text in this news page unreadable. (b) The same page as a Summary Thumbnail contains enough readable text to allow users to identify the area containing the sought content.

SUMMARY THUMBNAILS

To address this issue, we propose *Summary Thumbnails*. Summary Thumbnails are thumbnail views enhanced with fragments of readable text as shown in Figure 2b. Similar to traditional thumbnails, Summary Thumbnails preserve the original page layout which allows users to identify the overall page structure, and can help users recognize previously viewed pages. In Figure 2, for example, users may recognize the typical appearance of the homepage of the CNN news site and may find that the page consists of a thin menu column at the left, a news headline column in the middle, and a column of links and auxiliary materials at the right. Unlike traditional thumbnail views, however, the readable text fragments provided by Summary Thumbnails allow users to disambiguate the desired news story from similar looking areas. For example, the image caption shown in this Summary Thumbnail allows users to verify that this really is the story of interest concerning the Iraqi Governing Council. The readable text thereby eliminates the need for any further zooming or panning activity.

In order to fit the space requirements, mainly to fit the page to the browser width, Summary Thumbnails contain less text than the original web page (we will describe our text

reduction algorithm in detail in the “Implementation” section). When zoomed in, however, Summary Thumbnails show the original, unabbreviated version of the page as shown in Figure 3. Zooming interactions involving a change in representation have also been referred to as *semantic zooming* [3]. Note that despite the change of representation during this zoom operation, the Summary Thumbnail and the detail views look similar enough for users to maintain a sense of which areas in the thumbnail correspond to which areas in the detail view.



Figure 3: The zooming interaction provided by Summary Thumbnails is a semantic zoom: zooming in replaces the abbreviated text with complete text.

Summary Thumbnails can be scaled arbitrarily, allowing them to fit the screen size of any target device. Font size can be adjusted independently, which allows adapting Summary Thumbnails to the requirements of the target scenario. In the case of Personal Digital Assistants or Smartphones with a crisp display, setting the font size to the smallest readable font size (e.g., 7 pixels, as done in Figure 2) maximizes the amount of viewable screen content. In other scenarios, reading text of that size would be hard. For example, the display may use very high resolution resulting in very small font, the display may be blurry (e.g., a TV set), or the user may be vision-impaired. In these cases, setting font size to a higher value can help address the problem. Figure 4 illustrates this at the example of a page being scaled to fit a variety of target screen sizes.

In the remainder of this paper, we give a brief summary of the related work and discuss the algorithm and implementation behind Summary Thumbnails in more detail. We then present two user studies in which Summary Thumbnails outperformed the competing Single-Column interface in terms of task time and accuracy, and the Thumbnail interface in terms of navigation effort. We conclude with a summary of our findings and some ideas for future work.



Figure 4: Summary Thumbnails can be adapted to arbitrary screen and font sizes. Here the same web page is shown on mockups of different devices.

RELATED WORK

Summary Thumbnails are associated with two fields of related work, i.e., small-screen browsing techniques and semantic zooming.

There are four general approaches to displaying web pages on small screen devices: device-specific authoring, multi-device authoring, automatic re-authoring, and client-side navigation [4]. The first two approaches obtain high-quality results by authoring for device specifics (e.g., [13]). This requires the effort and cooperation of the individual page authors, and cannot be applied to existing pages.

Automatic re-authoring and client-side navigation do not require the collaboration of page authors and are therefore more widely applicable. Research prototypes that use automatic re-authoring fall into two main categories: page reformatting and page scaling. An example of page reformatting is the aforementioned single-column views (e.g., *Small-Screen Rendering*TM, opera.com). Other examples of techniques based on page reformatting include the *Power Browser* [6], where images and white space are removed, and the *WEST* browser [5], which uses flip zooming, a visualization technique that breaks pages into screen-sized tiles, and presents them as a stack. Difficulties with recognizing layout and leveraging the desktop browsing experience, as we have described them for single-column browsing, are common to all these approaches, since they all impact the page layout more or less significantly.

To avoid the drawbacks faced by page reformatting, researchers proposed approaches that preserve the appearance of the page by scaling the page, resulting in a web page thumbnail [7]. To help overcome the user effort involved in zooming thumbnails, researchers have proposed

extending thumbnail browsers by adding overview plus detail solutions ([18], also, in the *Thunderhawk* browser, thunderhawk.com), callouts for selected content areas (*WebThumb* [23]), user-selected column zoom (*SmartView* [16]), text summaries [11], rapid serial visual presentation display (*Streaming Thumbnails* [8]), and removal of irrelevant tiles (*collapse-to-zoom* [2]). Fisheye-based solutions such as fishnet [1] were shown to be useful for reducing the length of a web page. For reducing the width of a page, however, fisheye-based approaches can force users to scroll horizontally for each line of read text.

The second field of related work is semantic zooming, or techniques that change the representation of content during scaling or zooming [3, 12]. In the context of web browsing on the desktop, semantic zooming has been used to present query results [24, 9], to support keyword searching inside web pages [21, 1], and to improve accessibility [26]. For smaller screens, Gomes et al. [11] developed a system that displays text on PDAs at different semantic zoom levels: from displaying only the title at the lowest level, to displaying the complete original text. Lee & Grice [15] extracted text from XML-based files and displayed them in a customizable viewing style on PDAs. While these systems allow PDA users to view a larger selection of web pages, neither of them preserves the original layout of the viewed pages.

Researchers have found that displaying both the thumbnail and a text summary of a web page better supports page identification among query results [9]. These two elements can be presented separately [9] or integrated, e.g., in enhanced thumbnails where search words “popout” of otherwise unreadable thumbnails [24]. The concept of search term popouts was used to help users find keywords in web pages more efficiently by combining it with an overview-plus-detail approach (*popout prism* [21]) or a fisheye approach (*fishnet* [1]). While search term highlighting/popouts were proven to be effective [1], their applicability is limited to those cases where users can, and are willing to express their information need in terms of query terms.

Summary Thumbnails combine many benefits of the approaches listed above. As thumbnails, Summary Thumbnails preserve page layouts, and allow users to leverage their prior browsing experience. Readable text fragments allow users to disambiguate page content and to identify relevant areas. Further, since these text fragments are offered across the entire page, it takes less effort to skim the page than when using techniques that require users to explore the page using isolated focus areas [18, 23].

IMPLEMENTATION

A standard way of processing web pages for viewing on small screen devices is through a proxy server that transforms pages on-the-fly (e.g., *Thunderhawk* browser, thunderhawk.com, but also [26]). A proxy server is a program that receives web page requests (here from mobile devices), loads the respective pages, converts them, and serves them to the devices that requested them. Running the proxy on a powerful machine, such as a PC or server,

eliminates the need for processing on computationally weak mobile devices. Also, this approach makes it easier to serve different platforms, such as the ones mentioned above.

Our current implementation of Summary Thumbnails implements such a converter program, a standalone executable that runs on a desktop PC. However, since our immediate focus was to enable the user study, our converter still lacks the communication capabilities of a proxy and requires users to load and save pages manually.

Our converter program supports two output formats. First, it can output a converted HTML page. This page has the structure of a Summary Thumbnail, i.e., it contains abbreviated but enlarged text, but it's still as big as the original page. The final web page size reduction is performed by displaying the page on a web browser with scaling capabilities. We used the CSS zoom feature supported by MS Internet Explorer 5.5 and later [2] (msdn.microsoft.com).

Second, to implement our study conditions, we configured our converter to automatically render pages using a Microsoft WebBrowser control (<http://msdn.microsoft.com>), capture its output, and scale down the captured bitmap images to the desired output size using bi-cubic filtering. We created a simple viewer program for viewing and scrolling the resulting bitmaps. We used this viewer to administer all interface conditions in our user studies. Also all screenshots shown in this paper are screenshots of this viewer.

Our converter prototype is based on the MSHTML library, which prevents it from handling pages that contain frames. Besides that, conversion so far has worked for all web pages we sampled. Note however that since our converter program modifies only text, text encoded as bitmaps remains unchanged.

Text reduction

Removing common words: Words that otherwise occur less frequently tend to carry more of the meaning contained in a segment of text [19]. When reducing text, our converter program therefore removes common words first, as defined in a standard word frequency list [25]. The preservation of rare words also helps preserve keywords that a user might be visually scanning for. Alternatively, our converter program can be configured to crop paragraphs, which can be preferable for cases where text is already highly summarized, e.g., news headline. The Summary Thumbnail shown in Figure 2 was generated using the second method.

Preservation of line count: Our initial strategy was to make Summary Thumbnails preserve the overall length of the page. Since larger font in Summary Thumbnails is also taller, however, length preservation would have required removing entire lines, and this resulted in pages that appeared to be incomplete. Our current prototype therefore preserves the total number of lines instead of the absolute length of the page. The resulting Summary Thumbnails are typically longer than the corresponding thumbnails. The actual length increase depends on the amount of text in the

original page. For the 44 pages used in our user studies, the increase in length ranged from 0 to 83% (median 33%). In cases where neighboring columns contain different amounts of text, text growth can affect the vertical alignment at the bottoms of the columns. For web pages we looked at so far, the misalignment seemed acceptable.

Omission of ellipses: It is common practice to replace removed content with placeholders [2] or ellipses to indicating the omission, yet we decided against that option. Due to the high degree of text reduction in Summary Thumbnails this would have add visual noise, be spatially expensive, and would render ellipses mostly meaningless given their ubiquity.

Detailed description of our algorithm

Here is a schematic overview of our algorithm. First, the page is loaded. Then all text that is smaller than a user-defined threshold is enlarged. The result is a page that is still as wide as the original page, but now contains large text, e.g., about 19 pt for a page to be displayed on a PDA. Enlarged text typically occupies more lines than at its original font size. To preserve the line count, our program removes words until the total numbers of lines in the paragraphs are preserved. The resulting page is then saved in HTML format for devices with scaling capabilities, or rendered, scaled, and saved as a bitmap for all other devices. The following two paragraphs describe this process in additional technical detail.

The page is loaded and partitioned into *elements* (`IHTMLElements`, as defined in the Microsoft MSHTML library, msdn.microsoft.com) by recursively traversing the page's Document Object Model (w3.org). Elements can be paragraphs of text, input boxes, option boxes, or other elements specified in a style sheet.

To reduce the text, our prototype iterates through all elements of the page. For each element it performs the following steps: (1) Width, height, and font attributes are extracted. The number of lines of text is not directly available, so it is estimated as cell height/font height. (2) The text is extracted from the element and stored in a string. (3) The font of the string is enlarged. (4) *First reduction:* The width of the string is compared with a first estimate of the available space, i.e., width of the element multiplied by the number of text lines in the element. Lowest ranked words are removed until this space requirement is met. (5) *Second reduction:* If the element contains multiple lines of text, the string is broken to the width of the element. More words are removed until meeting these more stringent space requirements. (6) If the element has not enough room to even accommodate a single word, one word is selected and cropped to fit the space constraints. (7) The `innerText` property of the element is updated. For text that is separated by tags, the original element's `innerHTML` property is traversed. Rather than replacing this text as a whole, all those words that had been removed from the string earlier are now removed from `innerHTML`, to preserve the original

formatting. (8) Appropriate tags are inserted to increase the font size of the text in the element. (9) The resulting page is exported as described above.

QUALITATIVE USER STUDY

Our first user study was a qualitative study. Its purpose was to solicit initial user responses to our design and obtain a sense of appropriate tasks for use in the following quantitative study.

Study Design

Participants: We recruited 9 participants internally (7 male). All were mobile device users.

Interfaces: The participants used three different interfaces: a Thumbnail interface (Figure 2a), a Single-Column interface (Figure 1), and a Summary Thumbnail interface (Figure 3). All interfaces were displayed on a laptop computer using a viewer program that offered a net display area of 240x320 pixels for actual page content. All three interfaces allowed participants to scroll vertically through the web pages using the scrollbar. All interfaces fitted web pages to the screen width to remove the need for horizontal scrolling. Thumbnail and Summary Thumbnail interfaces were early design prototypes; while we explained the zooming capabilities to the participants, these prototypes did not yet support zooming; zooming support was not added until the quantitative study.

Procedure: Participants were presented with the same news page (<http://news.bbc.co.uk>) displayed on the interfaces in random order. They were told to scroll through the page and pick a news story they deemed interesting and elaborate on how it was represented on the respective interface. We encouraged participants to “think-aloud” during this process. For each interface, we conducted a brief interview and asked the participants to list pros and cons of the individual interfaces, and what page types they expected each interface to be most and least suitable for. We asked additional questions about application scenarios using a questionnaire. Overall, the study took 45 minutes per participant. Participants received a minor gratuity for their participation.

Results

Thumbnail interface: According to participants, the major advantages of the Thumbnail interface were that it preserved page layout (4 reports), that it provided an overview (8 reports), and that it provided a sense of branding (2 reports). However, six participants said the text—and for some pages images as well—could end up being too small to be useful. Two participants expressed concerns about the need for zooming before the text information on the page could be read.

Six participants judged the Thumbnail interface as useful in situations where the layout of the pages aided navigation, e.g., in maps and previously visited sites, or where the images conveyed enough information for decision-making, e.g., shopping sites that made heavy use of images. Seven participants judged the interface as inappropriate for pages

that required users to actually read the page, as opposed to just scanning it (e.g., portal pages, lists of search results).

Single-Column interface: Six participants rated the legible text offered by the Single-Column interface favorably, and one said he felt more confident before clicking a hyperlink using this interface than when using any of the other two study interfaces. Four participants liked that this interface avoided the need for horizontal scrolling. Eight complained about the altered layout and three participants described the transformed web page as being “ugly”, or “unrecognizable”. Four participants expressed dislike for the large amount of vertical scrolling this interface required; one participant liked the fact that all information could be viewed by vertical scrolling only.

Seven participants judged this interface as useful for linear reading (e.g., news, search results), and six said they found it inappropriate for tasks that relied heavily on the original layout of the page (e.g., map, bus schedules, and carefully designed pages). One participant expressed concerns that the Single-Column interface would be inappropriate for tasks that required users to compare different parts within the page (e.g., performing price comparisons while making a purchase decision).

Summary Thumbnail interface: Eight participants found and liked that the Summary Thumbnail interface preserved page layout better than the other two interfaces. Seven mentioned the text in the Summary Thumbnail interface to be legible, and judged this superior to the thumbnail interface. However, one participant was concerned that this interface would show too little text information for him to be able to select areas for further investigation and was concerned that the abbreviated text content offered by this interface might be misleading.

All nine participants judged the interface as suitable for tasks that relied on the layout of the page (e.g., maps, tables, and professionally designed pages), but not for linear reading (e.g., reading of news articles or search lists).

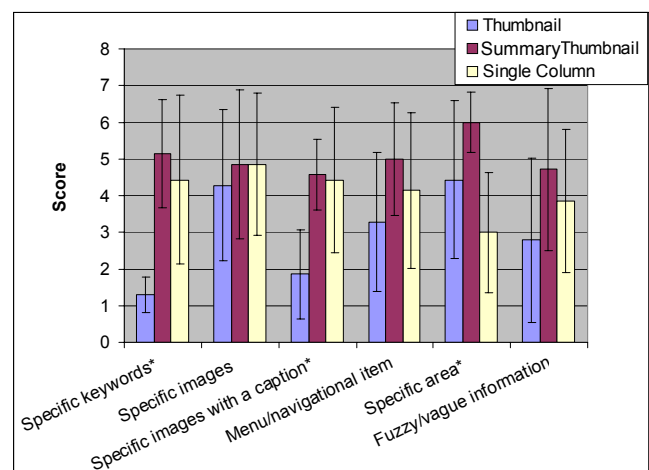


Figure 5: Questionnaire results. Scores are averages of 9 results (1=useless, 7=extremely useful, * one or more differences significant).

Questionnaire results and overall preference: Figure 5 shows the results of the questionnaire in which participants rated how appropriate they judged the individual browsers for six given types of web browsing tasks. We performed a Kruskal-Wallis test on the satisfaction ratings for the three browsers across the six task types, and found three significant differences. The majority of participants judged the Summary Thumbnail interface as more useful than Thumbnail interface for keyword search ($\chi^2=11.4, p=0.003$) and captioned image searches ($\chi^2=9.0, p=0.01$), and more useful than the Single-Column interface when looking for a specific area on a previously visited page ($\chi^2=7.8, p=0.02$).

In the final ranking, 8 of 9 participants ranked the Summary Thumbnail interface first; one preferred the Single-Column interface.

QUANTITATIVE USER STUDY

In this study, we evaluated the performance of the same three small screen interfaces (Summary Thumbnail, Thumbnail, and the Single-Column interfaces). In addition, we also included a web browser using a desktop-sized display window (“Desktop interface”) as a common reference. The participants’ task was to locate information in web pages originally designed for the desktop. Our main hypotheses were that the Summary Thumbnail interface would require less scrolling than the Single-Column interface and less zooming than the Thumbnail interface and that this would lead to measurable task time and accuracy differences. The experiment was a within-subjects design with a single manipulation (Interface). The dependent variables were task time and error rate. In addition, we recorded zooming and scrolling events.

Participants

11 participants (7 males) from the local community were recruited for the study. Ages ranged from 23 to 52 years, median 42 years. All had prior experience with desktop web browsers, and 4 out of 11 of the participants had previously used a mobile device for web browsing. Two of them indicated that they preferred using a 2D spatial view with horizontal and vertical scrolling for web browsing on their PDAs while the other two preferred the single-column mode on their devices.

Interfaces

In the experiment, participants viewed web pages using the following four interfaces: a Thumbnail interface, a Summary Thumbnail interface, a Single-Column interface, and a Desktop interface. The first three were similar to the prototypes used in the qualitative study reported earlier, but with additional interaction functionality. In addition to using the scrollbar, participants could now vertically scroll using keyboard or mouse wheel. The Thumbnail and Summary Thumbnail interfaces also allowed participants to zoom into a 1:1 scaled view of the original page by clicking the respective area in the thumbnail view as shown in Figure 3. While zoomed in, participants could scroll the page in all four directions by pressing the arrow keys on their keyboard or by using the scrollbars. The Desktop in-

terface displayed the original web pages without reformatting or scaling, and only vertical scrolling was provided via scrollbar and keyboard. Single-column views were generated using an Opera browser supporting Small-Screen Rendering™ (opera.com). All four interfaces were able to display all pages used in the study without requiring horizontal scrolling—horizontal scrolling occurred only in the zoomed-in views of the two thumbnail-based interfaces.

All interfaces were presented on an 18” LCD screen running at 1280x1024 pixel resolution. As shown in Figure 6, the top area of the 818x827 pixel study environment contained a textual description of the task, while the remaining display area was used to display the respective interface. For the thumbnail-based interfaces and the Single-Column interface, the window contained a simulated PocketPC device with a display area of 240x320 pixels. Since we included the Desktop condition as an approximate upper bound on performance on these tasks, we did not intentionally limit the browsing space used by these interfaces. In fact, our Desktop offered a viewing space of 800x600 pixels (Figure 6b).



Figure 6: Program used to administer the quantitative experiment running the (a) Summary Thumbnail, (b) Desktop, and (c) Single-Column interface.

Task

Each trial started with a written task description displayed at the top of the study environment (Figure 6). When ready, the participant clicked a button to reveal the web page described and to start the timer. The participant's task was to locate the information described in the textual description and place a marker on that location. Markers were placed by clicking into that area or by dragging an existing marker with the mouse. Size and shape of the area considered a correct answer varied across pages, but all were large enough to allow easy placement of the marker. When satisfied, the participant pressed a button to complete the trial.

Example: The page in Figure 6 shows an imdb.com review for the movie Shrek 2. The task description at the top of the window reads "You are thinking about renting a movie for the night, and remember your friend once recommended 'Shrek 2'. You want to see what the rating is before heading to the video store. You have used this movie review/database site in the past, so you went to the home page, searched for 'Shrek 2' by name, and navigated to the following page. You expect to see the rating information somewhere in the main section of the page, and you know they will be next to those bright yellow stars. <Click 'Ready To Start' to look for the rating of the movie 'Shrek 2'>."

How these tasks were created

To obtain a balanced set of web pages and a description of an actual information need we went through the following three-step procedure instead:

First, we collected web pages and task descriptions by interviewing 12 university student volunteers. These volunteers did not participate in the actual study. During the interviews, the volunteers randomly selected three to five web pages from their recent browser history. For each of these pages, they told us why they visited the page, how they got to the page, and where they expected target information to appear on the page before it was displayed. We gathered a total of 45 pages. Figure 7 shows Summary Thumbnails of some of these pages.

Next, we manually aggregated the gathered information into task descriptions, each consisting of a web page and a brief task and background description. Figure 6 shows one example.

Finally, we had two internally recruited pilot participants perform all 45 tasks using the desktop interface. We removed a page because both pilot participants found its description ambiguous. Based on the pilot participants' task times, we divided the remaining 44 pages into the four sets of 2 (training) + 9 (timed) trials, such that the sets balanced both average task time and page types (e.g., image, text box, main section of the page etc.). During the study participants performed all four trial sets—each one on a different interface. Presentation order and the assignment of trial sets to interfaces were counterbalanced.



Figure 7: Summary Thumbnails of seven of the 44 pages used in the quantitative user study.

Procedure

At the beginning of the study, participants filled in a questionnaire with demographic information. Participants then completed 2 training and 9 timed trials using the first interface, after which they filled out a short questionnaire about that interface. Then they repeated the procedure with a different sets of pages on the remaining three interfaces. The presentation order of pages and their assignment to interfaces was counterbalanced. Finally, participants were asked about their overall preference.

Hypotheses

Our hypotheses were:

1. *Displaying web pages in a way that preserves their original layout allows users to locate information faster and more accurately.* For tasks where participants were able to find desired information based on page layout, we expected the desktop interface to perform best, followed by the Summary Thumbnail interface, and the Thumbnail interface. The Single-Column

interface should affect page layout more than the other interfaces and should therefore perform worse on such pages.

2. *The presence of readable text reduces the need for zooming navigation.* We expected the Summary Thumbnail interface to require less zooming than the Thumbnail interface. The Desktop and the Single-Column interfaces obviously required no zooming.
3. *Shorter pages require less scrolling.* We consequently expected the Desktop interface to require the least amount of scrolling. Since Summary Thumbnails were slightly longer than the corresponding thumbnails, the Thumbnail interface should require less scrolling than the Summary Thumbnail interface. We expected the Single-Column interface to require by far the highest amount of scrolling.
4. *Effect of incomplete text fragments are tolerable.* While incomplete text fragments on the Summary Thumbnails could potentially lead to misinterpretation and a higher error rate in those trials, we expected these effects to be minor.

Results

Unless otherwise stated, we used single-factor, Analysis of Variance to analyze the data. We performed post-hoc analyses with Bonferroni corrections for multiple tests to further explore significant main effects. We used an alpha value of $p=.05$ across all these tests.

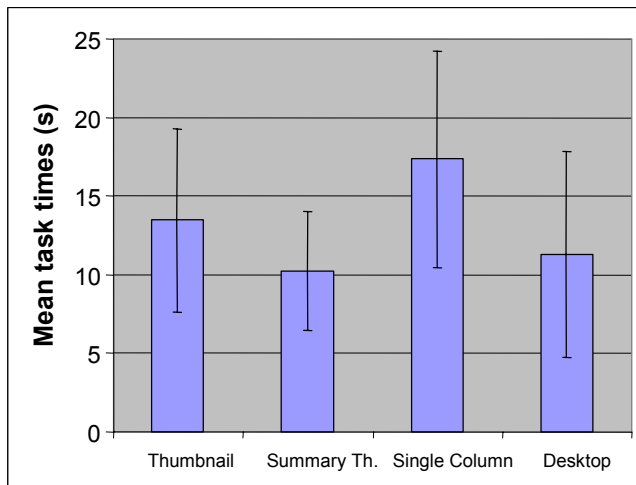


Figure 8: Mean task times for each interface (+/- standard error of the mean).

Task completion time

Task time results are shown in Figure 8. ANOVA indicated a significant main effect of interface ($F(3,40)=3.12, p=0.04$). Post-hoc analyses showed that Summary Thumbnail trials were significantly faster than Single-column trials. The difference was 41%.

Task Accuracy

Figure 9 shows the error rates, i.e., cases where participants had incorrectly placed the marker, thus failed to locate the intended target.




Figure 9: Error rates: percentage of erroneous trials across all 11 participants.

Zooming

Since there were only two interfaces that allowed zooming, an unpaired, two-tailed t-test was used to analyze the results. Participants zoomed 59% less often when using the Summary Thumbnail interface than when using the Thumbnail interface ($t(20)=2.1, p<0.001$, Figure 10). Three participants did not zoom in at all in the Summary Thumbnail trials.




Figure 10: Zooming event counts (+/- standard error of the mean).

Zooming accuracy: Next, we looked at what percentage of participants' zoom-in interactions were successful, i.e., the resulting detail view did contain the target area. The difference (62% with the Thumbnail and 51% with the Summary Thumbnail interface) was not statistically significant.

Scrolling

The vertical scrolling results are shown in Figure 11a. A single-factor ANOVA revealed a main effect of Interface ($F(3,40)=12.5, p<0.001$). Post-hoc analyses showed that the Single-Column interface required significantly more vertical scrolling than any of the other three interfaces (e.g., 6.9 times more than the Summary Thumbnails). When zoomed in, horizontal scrolling results are shown in Figure 11b. Participants scrolled 88% less horizontally when using the Summary Thumbnail interface than when using the Thumbnail interface ($F(1,20)=15.3, p<0.001$).

We observed that participants scrolled back and forth when they had trouble orienting themselves. Scroll direction change results may therefore give a sense of participants' level of confusion while searching for the targets (Figure 11c and d). A single factor ANOVA of the scrolling data for the four interfaces revealed a main effect ($F(3,40)=4.3$, $p<0.01$) for the vertical direction change. Post-hoc analyses showed that this parameter is 4 times lower in the Summary Thumbnail than the Single-column trials. The Summary Thumbnail trials contained horizontal scroll direction change 82% less than the Thumbnail trials ($t(20)=3.9$, $p<0.001$).

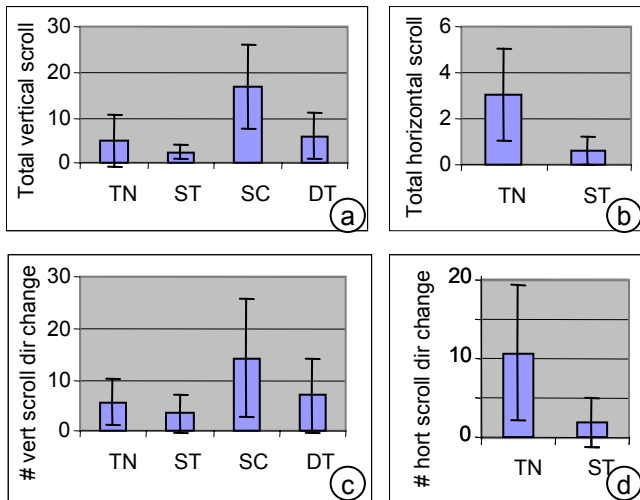


Figure 11: Scrolling results. (a) total amount of vertical scrolling for the interfaces (unit=1000 pixels); (b) total amount of horizontal scrolling while zoomed in (unit=1000 pixels); (c) total number of vertical scroll direction change events; (d) total number of horizontal scroll direction change events while zoomed in (TN=Thumbnail; ST=Summary Thumbnail; SC=Single-column; DT=Desktop).




Figure 12: Participant's preferences for use on their personal browsing device

Overall preference

At the end of the study, we asked participants “Which of the three web browser: Thumbnail, Summary Thumbnail, and Single-column would you install on your mobile device?” As shown in Figure 12, nine out of 11 participants preferred the Summary Thumbnail interface ($\chi^2(2)=12.2$, $p=0.002$), while the remaining two selected the Single-Column interface.

DISCUSSION

The results of the study indicated a strong participant preference for the Summary Thumbnail interface over the Single-column and the Thumbnail interfaces on small screens when browsing web pages originally designed for the desktop.

Our results suggest that better preservation of page layout can help participants browse. Trials using the Summary Thumbnail interface were 41% faster than those using the Single-Column interface where layout was considerably altered. Participants also made more mistakes when using the single-column interface than when using any of the other interfaces. Our observations during the study indicate that the layout conversion of the Single-Column interface may have affected participants' ability to recognize some page elements, such as horizontal tabs or horizontal menus, as well as overall page structures. In fact, the Single-Column trials contained more back-and-forth scrolling than all other trials, indicating that participants had difficulties in orientation. This confirmed our first hypothesis, wherein we postulated that layout information would help visual search.

Our study results also supported our second hypothesis: layout information in itself is not always sufficient for locating content—readable text is required as well. By offering fragments of readable text, the Summary Thumbnail interface was able to reduce the amount of zooming by 59% when compared to the Thumbnail interface. Three participants even located all trial targets without ever zooming into the detail view. Another indicator of direct access to target information was scrolling. Summary Thumbnail trials contained less horizontal scrolling and scrolling direction changed less often than in the Thumbnail trials. These results may indicate that participants could disambiguate the page content better using the Summary Thumbnails than with “plain” thumbnails.

Interestingly, participants scrolled 51% less when using the Summary Thumbnail interface than when using the Desktop interface. While this result may seem surprising at first, it is easily explained: Since the Summary Thumbnail interface (as well as the Thumbnail interface) was running in the “portrait” aspect ratio typical of *handheld devices*, it allowed participants to see more of the length of the page than the Desktop interface, which used a landscape aspect ratio. The differences in scrolling did not lead to any significant effects in task time or error rate. Still, we are surprised to see that the Desktop interface did not clearly outperform the Summary Thumbnail interface. One possible interpretation of this finding is that the reduced amount of scrolling and the reduced amount of text participants were confronted with on the Summary Thumbnail compensated for the obvious benefits of the Desktop interface.

Our concern that text cropping in the Summary Thumbnails interface would lead to a higher error rate was not confirmed. Instead, participants made fewer errors with the Summary Thumbnail interface compared to the Single-Column interface, where all of the original text was avail-

able. Also, participants seemed to find it easier to orient themselves with the Summary Thumbnail interface than the Single-Column interface, as indicated by the scroll direction change results.

CONCLUSIONS AND FUTURE WORK

In this paper, we presented Summary Thumbnails, a technique for displaying web pages on small screen devices that combines the benefits of thumbnail-based web browsers with the benefits of readable text. Our user study results indicate that Summary Thumbnails were found to be more effective in supporting web browsing than the single-column browsing technique that currently dominates the commercial market of small screen devices. Also, the vast majority of participants preferred Summary Thumbnails over traditional thumbnail views as well as single-column browsing.

As future work, we plan to combine Summary Thumbnails with a more powerful zooming interaction (e.g., *collapse-to-zoom* [2]) and the auto-cropping of photos [22]. In addition, we plan to explore the applicability of Summary Thumbnails for illustrating search results and as an alternative to the Thumbnail view in the file system viewer.

Acknowledgments

We would like to thank Mary Czerwinski, Ed Cutrell, Sue Dumais, Maneesh Agarwala, Dane Howard, Dan Vogel, Xing Xie, and Chong Wang for their support.

REFERENCES

1. Baudisch, P., Lee, B., and Hanna, L. Fishnet, a fisheye web browser with search term popouts: a comparative evaluation with overview and linear view. In *Proc. AVI 2004*, pp 133-140.
2. Baudisch, P., Xie, X., Wang, C., and Ma, W-Y. Collapse-to-Zoom: Viewing Web Pages on Small Screen Devices by Interactively Removing Irrelevant Content. To appear in *Proc. UIST 2004*.
3. Bederson, B.B., Hollan, J.D., Perlin, K., Meyer, J., Bacon, D., and Furnas, G., Pad++: a zoomable graphical sketchpad for exploring alternate interface physics, *Journal of Visual Languages and Computation*, 7, 1996, p3-31
4. Bickmore, T., and Schilit, B, Digestor: Device-Independent Access to the World Wide Web. In *Proc. Seventh Intl. WWW Conference 1997*, pp. 655-663.
5. Bjork, S., Bretan, I., Danielsson, R., and Karlgren, J. WEST: A Web Browser for Small Terminals. In *Proc. UIST'99*, pp. 187-196.
6. Buyukkokten, O., Gracia-Molina, H., Paepcke, and Winograd, T. Power Browser: Efficient Web Browsing for PDAs. In *Proc. CHI 2000*, pp. 430-437.
7. Chen, L., Xie, X., Ma, Y. W., Zhang, H-J., Zhou, H. Feng, H. DRESS: A Slicing Tree Based Web Representation for Various Display Sizes. Technical Report MSR-TR-2002-126, December 2002.
8. Czerwinski, M. P., van Dantzich, M., Robertson, G., and Hoffman, H. The Contribution of Thumbnail Image, Mouse-over Text and Spatial Location Memory to Web Page Retrieval in 3D. In *Proc INTERACT'99*, pp. 163-170.
9. Dziadosz, S., and Chandrasekar, R. Do Thumbnail Previews Help Users Make Better Relevance Decisions about Web Search Results? *Proc SIGIR'02*, pp. 365-366.
10. Gajos, K. and Weld, D. SUPPLE: Automatically Generating User Interfaces. In *Proc. IUI'04*, pp. 93-101.
11. Gomes, P., Tostao, S., Goncalves, D., and Jorge, J. Web Clipping: Compression Heuristics for Displaying Text on a PDA. *Mobile HCI '01* (poster paper), 2001.
12. Good, L., Bederson, B., Stefik, M., and Baudisch, P. Automatic Text Reduction For Changing Size Constraints. In *CHI'02 Extended Abstracts*, pp. 798-799.
13. Jacobs, C., Li, W., Schrier, E., Barger, D., and Sale-sin, D. Adaptive Grid-Based Document Layout. In *Proc. SIGGRAPH'03*, pp. 838-847
14. Lawton, D. T. and Feigin, E. J. Streaming Thumbnails: Combining Low Resolution Navigation and RSVP Displays. In *Proc. CHI'00*, pp. 159-160.
15. Lee, K. B., Grice, R. A. Getting and Giving Information: An Adaptive Viewing Application for the Web on Personal Digital Assistants. In *Proc. 21st annual international conference on Documentation*, pp. 125-132.
16. Milic-Frayling, N. and Sommerer, R. Smartview: Enhanced document viewer for mobile devices. *Technical Report MSR-TR-2002-114*, Microsoft Research, Cambridge, UK, November 2002.
17. Nygren, E., Lind, M., Johnson M., and Sandblad, B. The Art of the Obvious. In *Proc. CHI '92*, pp. 235-239.
18. O'Hara, K., Sellen, A. and Bentley, R. Supporting Memory for Spatial Location while Reading from Small Displays. In *Proc CHI'99*, pp. 220-221.
19. Salton, G. and McGill, M.J. *Introduction to Modern Information Retrieval*. New York: McGraw-Hill, 1983.
20. Stone, M. C., Fishkin, K., and Bier, E. A. The Movable Filter as a User Interface Tool. In *Proc. CHI'94*, pp. 306-312.
21. Suh, B., and Woodruff, A. Popout Prism: Adding Perceptual Principles to Overview + Detail Document Interfaces. In *Proc. CHI'02*, pp. 251-258.
22. Suh, B., Ling, H., Bederson, B. B., and Jacobs, D. W. Automatic Thumbnail Cropping and Its Effectiveness, In *Proc. UIST'03*, pp. 95-104.
23. Wobbrock, J., Forlizzi, J., Hudson, S., Myers, B. Web-Thumb: Interaction Techniques for Small-Screen Browsers. In *Proc. UIST'02*, pp. 205-208.
24. Woodruff, A., Faulring, A., Rosenholtz, R., Morrison, J., and Pirolli, P. Using Thumbnails to Search the Web. In *Proc. CHI'01*, pp. 198-205.
25. Word frequency list from http://www.comp.lancs.ac.uk/ucrel/bncfreq/lists/6_1_1_all_tag.txt
26. Hanson, V. L. and Richards, J. T. A Web Accessibility Service: Update and Findings. In *Proc. SIGACCESS '04*, pp.169-176.
27. Jones, M. Marsden, G., Mohd-Nasir, N., Boone, K. Buchanon, G. Improving Web Interaction on Small Displays. In *Proc. WWW8 conference*, pp. 1129-1137.