

Popout Prism: Adding Perceptual Principles to Overview+Detail Document Interfaces

Bongwon Suh

Human-Computer Interaction Laboratory,
Computer Science, University of Maryland,
College Park, MD 20742 USA

Allison Woodruff, Ruth Rosenholtz, Alyssa Glass

Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304 USA

ABSTRACT

We present an overview+detail document interface that draws on perceptual principles to help users work with documents. Central to our approach is the use of improved document overviews. Our approach also includes novel highlighting in the full representation of documents, as well as techniques to help users smoothly transition from the overview to the full representation of the document. We present a specific implementation of our design for Web browsing. We also present a qualitative user study that indicates that our perceptual design principles are effective and that users prefer our interface to traditional “find” and highlighting techniques. Our user study additionally reveals interesting tasks and strategies supported in our framework that have implications for overview+detail document interfaces in general.

Keywords

Document interfaces, overview+detail, popout effects, thumbnails

INTRODUCTION

When working with documents, people routinely perform a number of tasks that can be facilitated by an overview+detail interface. Overview+detail interfaces present multiple, coordinated views of data [3]. For documents, the *overview* can be used for orientation and navigation, while the *detail view* shows a portion of the full representation of the document.

As an example of a task that can be improved by an overview+detail document interface, consider the task of finding a keyword in a document. Many desktop applications support a *find* feature that identifies the occurrences of a given word in a document. Conventional “find-in-page” or “find-in-document” interfaces do not include an overview; instead, users enter a keyword and the system automatically scrolls the document to the location of the next instance of that keyword. Because the system jumps to the new location without giving any context about that location, users often become disoriented. By contrast, a document overview+detail interface could allow users to

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 2002, April 20-25, 2002, Minneapolis, Minnesota, USA.
Copyright 2001 ACM 1-58113-453-3/02/0004...\$5.00.

search for words while preserving an understanding of their location within the document.

Despite their potential utility, document overview+detail interfaces are not widespread. For example, no standard commercial Web browser supports document overviews. We believe this lack of success is due to the fact that currently available document overview implementations do not adequately address the requirements of common tasks. For example, *plain thumbnails* (simple scale-reduced images of documents) provide a gestalt of the document, but contain less detail than users would like, e.g., words in plain thumbnails are generally too small to be read.

In this paper, we apply perceptual principles to improve the utility of document overview+detail interfaces. To demonstrate our techniques, we have implemented Popout Prism, an overview+detail system for Web pages. The system includes an overview+detail document interface in which the overview is an *enhanced thumbnail* [17] and the detail view is a portion of the full document. Perceptual principles are applied in both the enhanced thumbnail and the full document views to make elements of interest “pop out” visually. In Popout Prism, these *popout* elements are based on user-selected keywords. In addition to drawing attention to key elements, we employ perceptual principles in other aspects of the design, e.g., to ensure that the transition between the views places a low perceptual and cognitive load on users. We have also taken care in integrating perceptual aspects and user interaction aspects of the design. For example, when the user navigates to new documents, the interface automatically applies the existing keywords so that the popouts appear without requiring further action from the user. The popouts in the detail view are also controlled in a natural way through the use of the overview thumbnail as a scrollbar. The resulting design naturally supports a number of tasks ranging from a find feature that identifies the occurrences of a given word in a document to monitoring applications that require delivering information to the user in the periphery of their attention.

We have conducted a qualitative user study intended to examine the effectiveness of our perceptual principles, to study the strategies users employ with our system, and to explore a range of tasks for which overview+detail document interfaces could be useful. Participants indicated that our popout interface was preferable to plain thumbnails as well as to other alternatives with which they were

familiar, and they identified specific properties that made our interface useful and appealing for a number of tasks.

To summarize, the contributions of this paper include a novel overview+detail interface, perceptual principles for overview+detail interface design, and a qualitative user study of how participants use an overview+detail document interface. As we will see, while the popouts studied here contain keywords, our interface, principles, and study all have applicability beyond simple find interfaces.

In the next section, we describe our interface and implementation. In the following section, we discuss perceptual principles used in our design. We then discuss the method and findings of our user study. Next, we discuss related work, and finally, we conclude.

POPOUT PRISM

In this section, we present Popout Prism's user interface, a sample use scenario, and our implementation.

User Interface

Popout Prism is composed of three panes: an overview on the left hand side that shows enhanced thumbnails of the current Web page, a detail view on the right hand side that shows the current Web page at full scale, and a toolbar along the top that supports standard browsing functions (e.g., back, forward, URL entry) as well as keyword entry for the popouts. In this subsection, we give a summary of user interaction with the system and then further describe the overview and the detail view.

When the user enters keywords, the system locates instances of these keywords in the document and applies a popout transformation to those instances that are currently visible in either the overview or detail view, enlarging and transparently highlighting them. (See Figure 1.) Keywords remain *active* until the user removes them from the keyword field in the toolbar. If the user navigates to a new page while keywords are active, popouts are automatically created for the new page.

The thumbnail overview can be used to navigate in the detail view. A sliding window, drawn as a dotted rectangle on the thumbnail overview, identifies the currently visible portion of the full Web page. This sliding window is tightly coupled with the full Web page so that the thumbnail overview may be used as a scrollbar. Users can drag the sliding window, moving smoothly through the Web page, or they may double-click anywhere in the thumbnail to automatically scroll the full Web page to a specified location.

The overview contains a thumbnail of the Web page displayed to its right. Thumbnails are one-quarter the width of the Web page and are generated immediately after the Web page has been fully downloaded. If the thumbnail would be taller than the vertical height of the overview pane, we vertically compress it to fit in the pane, so that users always have a complete overview of the document. (By default, this happens when the full Web page is more than four screens long; approximately 11% of the pages in

our user study were in this category.) When users do not specify any keywords, the overview is essentially a plain thumbnail. Popouts in the overview appear when keywords are entered and remain visible as long as keywords are active. The popouts are designed in accordance with previous research on enhanced thumbnail summaries [17]. They are enlarged so their text is readable in the thumbnail and are highlighted in a salient color, using a different color for each keyword or keyword phrase. The resulting popouts are transparently overlaid at the location of the original textual element in the underlying thumbnail. Showing elements in context is a very powerful technique, allowing the user to see the spatial relationship between multiple elements, and also allowing the user to see the position of a given element relative to the structure of the document. The colors in the underlying thumbnail are desaturated so it will be less likely to distract attention from the popouts.

The detail view is a scrollable display of the current Web page. When users do not specify any keywords, the Web page appears as it would normally appear in Microsoft® Internet Explorer. When users do specify keywords, the system adds popouts to draw attention to the textual elements that match the keywords and are currently visible. The popouts in the detail view are designed to be similar to the popouts in the overview in style aspects such as color, transparency, and relative size. The upper left hand corner of the popouts in the detail view is aligned to the upper left hand corner of the original text, so that users can quickly find the original text associated with each popout. Unlike the popouts in the overview, popouts in the detail view are not continuously visible while the keywords are active. As the detail view popouts are predominantly useful for locating keywords and may be distracting during other phases of the user's task (such as reading), they are made available selectively. Popouts are shown for five seconds when keywords are first entered and then disappear; according to our informal observations, five seconds is sufficient time to allow users to recognize the existence of matching keywords. Users can make popouts disappear earlier by clicking anywhere on the thumbnail overview. Users can recall the popouts at any time by pressing the left mouse button on the thumbnail overview; the popouts disappear as soon as the mouse button is released. This button assignment means that popouts are visible in the detail view while the user is navigating with the thumbnail overview (e.g., dragging the sliding window) and disappear as soon as navigation is complete.

Scenario

To more clearly illustrate the usability and effectiveness of Popout Prism, we offer a sample use scenario. Suppose a user is browsing through news sites and decides she is particularly interested in economic news related to Federal Reserve rate cuts and the possibility of a recession. In the keyword section of Popout Prism, the user enters keywords 'rates cut recession.' She then leaves these keywords in effect while browsing a variety of Web sites. As she is

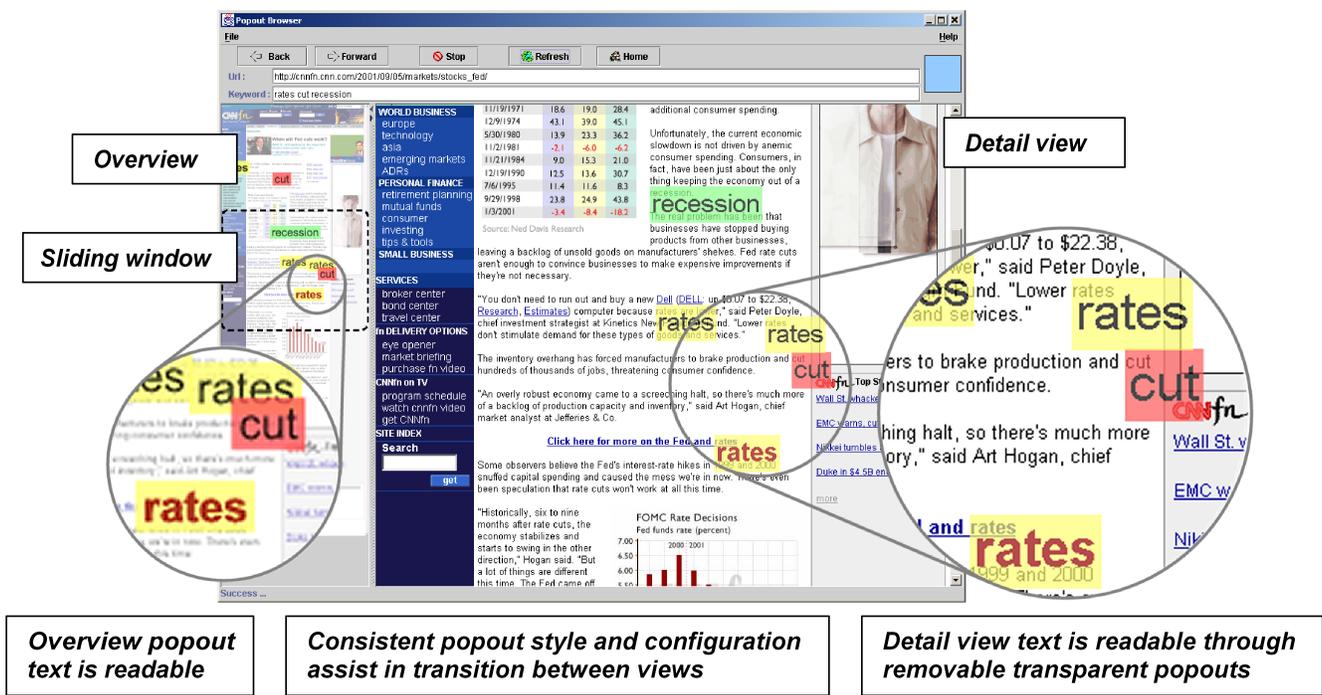


Figure 1. News Article Scenario.

Note that the zoomed portions are added for presentation purposes and are not present in the actual interface. (This figure is reproduced in color on page 000.)

browsing the CNN™ Financial Network Web site, <http://www.cnnfn.com/>, a cluster of popouts in the overview alerts her that some of her keywords appear in a part of the document that is not currently visible in the detail view. She drags the sliding window in the overview to scroll the document to the place the popouts appear. She then looks over at the popouts in the detail view and sees that they are associated with a short summary of an article. Based on the configuration and content of the popouts, and the document text she can read through the popouts, she decides to click through to the article.

When the article has loaded, new popouts based on the same keywords direct her attention to sections of the article that may be of particular interest to her, as shown in Figure 1. She drags the sliding window to navigate directly to the portion of the document that contains the keywords ‘rates’ and ‘cut’ in close proximity (Figure 1, left). She then uses the popouts in the detail view to locate the corresponding words in the document (Figure 1, right) and begins reading.

Implementation

The system is currently written in Java with a C++ JNI interface. To generate thumbnails efficiently, we include Microsoft® Internet Explorer 5.5 as an ActiveX control inside Popout Prism. This scheme enables us to display Web pages rendered as in a standard browser.

Popout Prism is able to generate thumbnails of Web pages shorter than 1000 pixels in height in less than two seconds. For example, it takes approximately two seconds for the system to generate a thumbnail containing several popouts from <http://www.yahoo.com/> on a Pentium III 800 MHz PC

with 256 MB of main memory. Web pages longer than ten screens (>8000 pixels) may take more than 5 seconds. The display of the popouts in the detail view is virtually instantaneous. As with any Web browser, network delays may also impact Popout Prism’s speed. We believe that performance can be improved in future versions of Popout Prism by using threads efficiently, or by doing incremental rendering, so that the thumbnail generation does not disturb page navigation.

In the long term we envision embedding Popout Prism in existing browsers. Therefore, we did not focus effort on core Web browser functionality, such as histories and bookmarks. Finally, the current implementation of Popout Prism does not correctly handle some Web pages that use frames. We hope to address this issue in future versions.

PERCEPTUAL PRINCIPLES

While designing our system, we considered a number of principles from the study of human perception. In this section, we present some of the principles we used to draw attention to popouts, preserve the usability of the detail view, and ensure that users could easily map between views. Each principle is followed by a discussion with references to the perceptual literature and examples of use from our system.

Drawing Attention to Popouts

We want the popouts to draw the user’s attention so that the user will be alerted to the presence of relevant popouts and also so that the user will be easily able to locate popouts in both views. We use color and size modifications to draw the user’s attention to the popouts.

Principle 1: Colors that are unusual compared to other colors in the display draw more attention; conversely colors that are similar to others in the display draw less attention

A quantitative model of visual search says that colors are more likely to draw attention if they are outliers to the distribution of surrounding colors in the display [15]. This suggests that popouts will stand out more when they are more distant in color space from surrounding colors, and further, it suggests that reducing the variability of surrounding colors in the display will make popouts draw more attention.

Accordingly, we use the above referenced model of visual search to choose popout colors that are distant from surrounding colors in the display. Additionally, we desaturate the underlying thumbnail. This reduces its contrast, thus reducing the variability of its colors. In a sense, we are reducing the attentional demands of the underlying thumbnail, so as to better allow the popouts to draw attention. Desaturating the background also allows us to use less saturated colors for the popouts while still drawing attention. Less saturated colors for the popouts have the advantages that they facilitate reading and are more soothing. Using a quantitative model of visual search to choose the colors of the popouts allows us to balance the competing goals of having the popouts draw attention and ensuring their readability.

Principle 2: Enlarging elements draws attention to them

Recall that in the thumbnail overview, we enlarge the popout text to make it more readable. Providing readable text places less cognitive load on the user than alternative designs that rely solely on color-coding to convey semantic information, since such designs require the user to mentally map colors to keywords. The enlargement of the popouts for readability has the added benefit of drawing attention to them, as it tends to make them larger than other objects – particularly text – in the thumbnail (see, e.g., [16] for a discussion of the impact of size on attention).

Because the popouts are readable and salient, the overview in our system contains sufficient detail to help the user decide where they want to go – a key attribute for document navigation.

Preserving the Usability of the Detail View

Using popouts in the detail view is more complex than using them in the overview because we need to ensure that the popouts do not compromise the overall readability of the document. We recognize that persistent modifications to the page will reduce the usability of the document and then discuss how to strategically add emphasis.

Principle 3: Persistent emphasis in the page distracts readers

Humans seem to have an “attentional control setting” – essentially, a setting for what they are (and are not) looking for while performing a given task. They have difficulty switching their attentional control setting instantaneously,

even in order to ignore an irrelevant and distracting cue [7]. Highlighting or other emphasis of keywords is likely to be selected by the user’s attentional control setting and therefore distract the user from activities such as reading or skimming. This is particularly true since the emphasis is often relevant to the user’s task. Further, forms of emphasis in the text are known to slow a reading task, at least for older users [5].

Therefore, although our transparent popouts generally allow reading of the underlying content of the page, we remove the emphasis – i.e., popouts – from the detail view when the user is not using them to navigate. Also observe that Principle 3 indicates that removing emphasis from the detail view makes sense even for modifications that are less obtrusive than popouts, such as simple highlighting [9,10].

Principle 4: Making objects appear draws attention

In our system, in addition to using color and size to draw attention in the detail view, we make use of the fact that an “onset,” or appearance of an object in the display, draws attention [18]. This is somewhat of an oversimplification, as it depends upon the features of items relevant to the task [7] and the level of focus of the user on the task [19], but for many situations this is a good rule of thumb. Thus the popouts particularly draw attention when they first appear, as well as when the user makes them reappear by clicking the left mouse button.

Making it Easy for Users to Map between Views

We apply perceptual principles to make the transition from the overview to the detail as smooth as possible. Recall that while the user navigates, the popouts are visible in both the overview and the detail view. Here we discuss principles for helping users map from popouts in one view to popouts in the other view. Note that these principles are generally relevant to brushing techniques [2].

Principle 5: People are good at recognizing the same pattern shown at two different scales

Because the overview and the detail views are at different scales, one might be concerned that it would be difficult for users to map from a group of objects in one view to the same group of objects in another view. However, research suggests that perceptual learning for object recognition is size-invariant [8]. Therefore, popout configurations are similar in the overview and the detail view to make it easy for the user to map between them.

Principle 6: Using the same emphasis style in two views makes it easier to map between the views

We use the same styles (e.g., colors) for popouts in both the overview and the detail view, thereby reducing the cognitive load required to switch views and allowing the user to maintain the same attentional control setting for both views.

Not all systems use the same emphasis style in both views (e.g., [10]), and in fact, using the same style is not necessarily an obvious design choice. For example, when

views are at different scales, different colors may seem appropriate for emphasis in each view. Yellow is an appropriate highlight color for dark text and therefore might be desirable in a detail view. However, yellow is not salient at smaller scales, due to low sensitivity of the visual system to yellow at high spatial frequencies [12]; therefore, a color like saturated red (which would be highly undesirable to highlight readable text) might be chosen in an overview. We specifically choose colors that balance readability and saturation, and our popouts are enlarged so they are not at as small a scale as the highlights in systems such as Reader's Helper [10]. This allows us to use the same colors in both overview and detail.

METHOD

We conducted a qualitative user study to examine whether our visual design was effective, to learn more about the tasks for which Popout Prism's functionality would be useful, and to observe the strategies users employed.

Procedure

Participants used Popout Prism for 3-5 days during the month of August, 2001. Because Popout Prism did not provide all of the functionality of a standard Web browser, such as bookmarks, participants were permitted to use other Web browsers as well during the study. A few participants used Popout Prism as their primary browser while others used it in more specific situations. Outside the interviews, each participant visited over 120 pages in Popout Prism and performed nearly 20 new keyword searches on average.

We interviewed each participant twice, once very early in their use of the system and once later in their use of the system, so that we could see which strategies participants developed immediately and which evolved over time. Each interview was videotaped and consisted of direct observation of participants using Popout Prism to perform their normal Web tasks, as well as a semi-structured interview.

Using the videotapes, we made notes on the interviews and clustered the results by theme. Note that some topics were raised by only a few participants. Therefore, the fact that a subset of participants made an observation does not imply that other participants did or did not agree.

Participants

Participants were twelve members of the Xerox PARC community, six men and six women. Approximately half the participants had a technical background and approximately half did not (e.g., administrative assistants, business development staff, a public relations manager). Participants spanned a wide range of ages and varied widely in terms of the average amount of time they spend using the Web in one week. Participants' typical Web tasks include reading news, reading documentation, visiting portals, and conducting literature searches.

FINDINGS

In this section, we present results from our user study. We have selected seven topic clusters, which have in turn been divided into three groups for expository reasons. These

groups are participant responses to the interface, participant strategies for using the interface, and participant comparisons with other systems.

Participant Response

Participants articulated their impressions and preferences at a variety of levels. These included general attitudes, opinions on usability, and assessments of the most useful aspects of the system.

High-level Impressions. Response was generally positive. Participants described Popout Prism using words such as "fun" and "cool." They said it was useful, "handy," and "a definite time-saver." A few participants did observe that while they liked it and found it useful, it was a nice-to-have tool rather than a must-have tool. The main complaints were about the lack of standard browser functions, the memory requirements of the tool, and the perceived slow speed of the tool. Only a few participants commented on Popout Prism's screen real estate requirements, and these participants indicated this was only a minor issue for them.

Overall, ten of the twelve participants said they would like to use Popout Prism again if it were a standard tool. A few participants were interested in continuing to use it in its current research-prototype state, commenting, "I'm hoping you don't take it away after a week" and "I'm not giving it back."

Usability. Participants generally found the tool easy to use. Training was minimal (usually either a short demo or a brief review of a short user manual). Participants immediately grasped the main concepts such as popouts and navigating with the thumbnail overview and quickly developed strategies for using the features. Some participants had minor confusion about mouse-click events, e.g., what made popouts appear and disappear.

Participants found popouts easy to locate in both the thumbnail overview and in the detail view: "I know exactly where I'm going... I know exactly here [points to thumbnail] and I know exactly here [points to full Web page]. There's no doubt, absolutely." Participants seemed to immediately understand the correspondence between popouts in the thumbnail overview and popouts in the detail view.

Participants did not generally feel that the thumbnail overview interfered with their tasks. The majority of participants did not feel that the popouts in the detail view interfered with their tasks, although a few participants did say that it was difficult to read in the detail view when large numbers of popouts were present.

Utility. We asked participants to compare the utility of thumbnail overviews with popouts to that of thumbnail overviews without popouts, and we asked participants to assess the utility of popouts in the detail view as well as the utility of the navigation mechanism. We also asked participants about the types of pages and tasks for which they found Popout Prism useful.

Most participants indicated that the thumbnail overviews with popouts were the most useful feature of the system. One participant said she found this part of the tool so useful that she physically moved her chair to orient to it. By contrast, plain thumbnails (thumbnail overviews without popouts) were judged to be significantly less useful than either the overview or the detail view with popouts, several participants saying plain thumbnails were not useful at all.

Most participants said that the popouts in the detail view were useful, primarily when used in conjunction with the thumbnail popouts to navigate to a popout or popouts. Correspondingly, the ability to navigate with the overview was also judged highly useful when popouts were present, although many participants did enjoy using the sliding window to navigate even when no popouts were present.

By the first interview, most participants had specific ideas about situations in which Popout Prism would be useful. For example, in terms of page characteristics, participants mentioned that Popout Prism was useful on longer pages (approximately 69% of the pages in our user study were more than one screen long), pages with which they were not familiar, and pages with poor structure. Participants mentioned that it was less useful (or not useful at all) on short pages, pages they knew well, pages they intended to read in entirety, and pages on which the word they were searching for appeared very frequently.

Participants suggested Popout Prism would be useful for a number of tasks such as reading news, reviews, documentation, or mailing lists, or for doing searches. A few participants suggested that popouts would be useful for providing information about pages other than the current page, e.g., a site map could be composed of thumbnails with popouts. Nearly half the participants volunteered that the popout functionality would be useful to them in other applications such as mail tools, word processors, spreadsheets, development environments, and PDF viewers.

Participant Strategies

The intuitive interface design enabled users to adapt the system to their daily tasks with little effort. The interviews revealed that users very quickly developed strategies for using visual properties of popouts, navigating, and choosing keywords.

Visual Search Strategies. Participants expressed pleasure that they could see an overview of the entire page at once, including all popouts, and they took advantage of several visual properties of the popouts to determine relevance. Here, we discuss three key visual properties: density, position in the document, and spatial configuration. Participants used the density of popouts to assess the relevance of a page. For example, when participants clicked through on search engine results to visit pages where they were hoping to find a particular piece of information, they said popouts helped them decide whether or not to stay at a page; if there were only one popout, they might conclude the topic was not covered in depth and try

another site, whereas if there were several popouts, they would be more likely to stay and look. For example, one participant said if no popouts were present, “I can just move on. I don’t have to even scroll down.” Density was also a useful cue about regions within documents; participants said they would save time by ignoring parts of documents that did not have popouts. They also said they could ignore popouts based on their position in the document, e.g., if they could tell from the thumbnail that the popout appeared in a navigation bar, or if the popout were at the end of a document, they might ignore it. The spatial configuration of popouts was also useful, e.g., participants said they would go to a “gathering” of popouts, or would look for places where multiple keywords co-occurred.

Overall, participants felt the visual properties of popouts saved time and effort. For example, one participant said that when he visits pages recommended by Google™, he normally reads the titles of pages; with Popout Prism, rather than reading the titles, he would instead glance at the grouping of popouts on the thumbnails to decide if the page were relevant to him.

Navigation Strategies. Since the thumbnail overview with popouts helped users assess what parts of a page to visit, participants found it very helpful for navigation. A general approach was for participants to identify a region of interest using popouts, and then navigate directly to that region using the sliding window. As one participant said, “... the thumbnail is really handy for saying, well, that’s where I go.” Participants said that they stayed oriented while using the sliding window, in comparison to using traditional find interfaces (discussed further below) or to scrolling.

Keyword Strategies. Participants used a variety of strategies to choose and enter keywords for popouts. Most participants left keywords active at least some of the time as they navigated across sites, sometimes because they wanted to search for the same keywords at multiple sites and sometimes because they did not want to bother changing keywords until they were ready to enter new ones. A couple of participants observed that having to enter keywords was tiring. Some participants suggested smart keywords, e.g., keywords that were personalized, or site-specific keywords that would automatically come up whenever they browsed to a given site.

A couple of participants talked about specifying different types of keywords simultaneously, some for the current search, and others for background monitoring. For example, one participant entered his name as a keyword so that anytime he came across a page with his name on it, he would be notified. As one participant said, “[Y]ou’re viewing a site and you weren’t really viewing it for this purpose but [Popout Prism] happens to tell you that oh, today’s slashdot has an article on MEMS and like, oh, cool, I wasn’t looking, I wasn’t planning on reading about MEMS, but you know, I’ll read about MEMS today.” Since popouts appear automatically on navigation to a site

and because they can be ignored, they are well-suited to such peripheral monitoring tasks.

The interaction of Popout Prism with search engines is particularly interesting. Some users entered the same keywords in both tools. A couple of users articulated a multi-level strategy, entering broad search terms in the search engine and more specific terms in Popout Prism; the search engine finds documents with certain terms that must be present, and within these results, Popout Prism highlights additional terms of interest. To illustrate this usage, one user entered “VB mysterious crashes version DLL” into the search engine and then “MDAC” for the popouts, because he had a suspicion that the source of the mysterious crashes was MDAC, and he wanted to see if MDAC appeared on the search results page.

A similar multi-level search was suggested by another participant. He frequently visits his personalized stock portfolio, a page with information about user-specified stocks containing a matrix with pricing information, etc., and pointers to relevant news articles. He has a number of stocks that he cares about only peripherally (e.g., the stock of a friend’s startup, or companies he uses to monitor industry trends) and a smaller number of stocks that he owns and cares about a great deal. He wants to see the entire stock portfolio, but can use popouts to highlight the stocks that he owns. In this way, he can focus on the stocks of highest relevance to him while viewing them in the context of the other stocks on the page.

Comparison with Alternative Keyword Search Interfaces

Participants compared Popout Prism to other tools with which they were familiar. Roughly half the participants said they commonly use the traditional find interface to locate terms in Web pages, and all were familiar with find either from Web browsers or other applications. All participants preferred popouts to find, citing three main reasons. First, most participants cited the advantages associated with the popouts in the thumbnail overview: a sense of context, density, and the ability to see all occurrences of a keyword at once. One participant said that with find, he needs to make two passes through the document, one to get an overview, and the second to look at individual instances; having the overview changed his search process to a single pass. Second, about half the participants complained about getting disoriented when using find, and said they disliked repetitively jumping from one place to the next. As one participant said, “It’s very disorienting to say find and then... jump to an entirely different part of the document... Every time I search, I jump to the next spot, I have to sort of spend a little bit of time, you know, figuring out where I am.” Because of these effects, participants said popouts were a good replacement for find: “Anywhere I would normally use find, I would be very happy to have the thumbnail show me where it is in the document all at once.” Additionally, participants said that they would use popouts in instances in which they would not use find: “I found myself using this a

lot more than I would find... I think that means I like this better.” Third, participants found popouts in the Web page easier to spot than the highlighting done by find.

One participant was a frequent user of Google™ Toolbar™ highlighting [9], and three participants were familiar with the feature but used it infrequently. These participants compared Popout Prism favorably to Google Toolbar highlighting, generally citing the same types of advantages that it had over find. The frequent user of Google Toolbar highlighting said, “Especially like when I showed you the mailing list... that was kind of like a godsend... I didn’t have to rely on Google’s highlighting and ... go uh-uh-uh there’s a highlight there [mimics scrolling through page searching for highlights in actual Web page as he scrolls], uh-uh-uh there’s a highlight there. I didn’t have to do that. That was really nice.”

RELATED WORK

A number of previous systems use reduced-scale document representations such as plain thumbnails, slightly modified thumbnails, and abstract representations (e.g., [14], [4], and [11], respectively). Previous work suggests that enhanced thumbnails showing a portion of a Web page have a number of advantages over these other representations as document summaries (communicating gist, genre, etc.) for Web search tasks [17]; in this paper we have built on this work to illustrate the benefits of using enhanced thumbnails of full Web pages as document overviews (providing context for the entire document).

Researchers have also enhanced full size documents. The Google Toolbar [9], for instance, highlights user-specified keywords within pages as the user visits them. Thus, the Google Toolbar enables “persistent search” as does Popout Prism. However, as the Google Toolbar is not an overview+detail system, it does not have the benefit of an overview to help the user easily navigate to important parts of the document, and our user study indicates that users can become disoriented as they search the document for keywords. Several overview+detail document viewers highlight text in the detail view of the document [6,10,11,13]. However, in both the Google Toolbar and these overview+detail systems, text is highlighted in place and/or made bold. A study reported in [13] showed a lack of utility for this sort of enhancement in their overview+detail system. Additionally, as we have discussed, human perception research suggests that the enlarged, transparently highlighted text of Popout Prism is better at drawing the user’s attention, and that the automatic fading of the popouts interferes less with reading in the detail view than persistent enhancement.

There are a number of overview+detail document interfaces, and several of these systems use a thumbnail overview to navigate through the detail view, as in our system [1,10]. Of these, the Adobe® Acrobat® Reader® does not enhance either the overview or the detail view. The Reader’s Helper [10] has a “Thumbar” that provides users with an overview of a Web page, with interesting

topics (automatically selected using personalization techniques) highlighted in both views. Though the Reader's Helper has some similarities to Popout Prism, we feel there are some key perceptual and usage differences. These result from its optimization for a very specialized task and greatly reduce the effectiveness of its design choices for keyword tasks. For example, no text is readable in the Thumbar overview (text is greeked and topic locations are indicated by a colored line), the personalized topics are application settings as opposed to keywords selected on-the-fly, and, at least in the examples in [10], all topics are indicated by the same color. The keywords in Popout Prism are readable in both views, keywords are user-selected at run-time, and different colors are used for different keywords, making it easy for users to identify and navigate to particular words of interest. Furthermore, we have argued that our popouts are better at drawing attention than the Reader's Helper's highlights. Finally, in the examples in [10], the thumbnail highlights are in red, while the detail highlights are in yellow. Though there are valid reasons to choose different colors for views at different scales, perceptual research suggests that using identical emphasis styles in both views makes it easier for the user to transition between the views.

CONCLUSION

We have presented an overview+detail document interface that draws on perceptual principles to help users effectively work with documents. Specifically, our interface offers improvements in the areas of document overviews, highlighting in the full representation of the document, and visual coupling between the two views. We have presented a specific implementation of our design for Web pages.

We have also presented a user study that confirms our expectations about the usability of our interface. Our user study additionally reveals interesting tasks and strategies that are supported in our framework. Finally, our user study indicates that our interface has advantages over traditional find and highlighting techniques, as well as over plain thumbnails.

We intend to pursue several directions in the future. In addition to refining the implementation and adding new features, we also intend to develop further design principles and conduct quantitative studies.

ACKNOWLEDGMENTS

We are grateful to the participants in our study for using experimental software and sharing their insights. We thank Paul Aoki, Ben Bederson, and Lance Good for helpful discussion and comments. We are indebted to Maarten van Dantzich of Microsoft Research for advice on interfacing with Internet Explorer. This work was performed while Bongwon Suh was an intern in the User Interface Research group at Xerox PARC.

REFERENCES

1. Adobe® Acrobat® Reader®. <http://www.adobe.com/>
2. Becker, R.A., and Cleveland, W. S. Brushing scatterplots. *Technometrics*, 29(2):127-42, 1987.
3. Card, S.K, Mackinlay J.D., and Shneiderman, B. *Information Visualization: Using Vision to Think*. Morgan-Kaufmann, San Francisco, 1998.
4. Cockburn, A., and Greenberg, S. Issues of page representation and organisation in Web browser's revisitation tools. *Austr. J. Info. Sys.* 7(2):120-7, 2000.
5. Connelly, S.L., Hasher, L., and Zacks, R.T. Age and reading: the impact of distraction. *Psych. & Aging*, 6:533-41, 1991.
6. Eick, S. G., Steffen, J. L., and Sumner, E. E. SeeSoft – a tool for visualizing line oriented software statistics. *IEEE Trans. Software Eng.*, 18(11):957-68, 1992.
7. Folk, C.L., Remington, R.W., and Johnston, J.C. Involuntary covert orienting is contingent on attentional control settings. *J. Exp. Psych: HP&P*, 18:1030-44, 1992.
8. Furmanski, C.S. and Engel, S.A. Perceptual learning in object recognition: object specificity and size invariance. *Vision Research*, 40:473-84, 2000.
9. Google™ Toolbar™. <http://toolbar.google.com/>
10. Graham, J. The Reader's Helper: a personalized document reading environment. *Proc. SIGCHI '99*, 481-8, 1999.
11. Hearst, M. TileBars: visualization of term distribution information in full text information access. *Proc. SIGCHI '95*, 59-66, 1995.
12. Mullen, K. T. The contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings. *J. Physiology*, 359:381-400, 1985.
13. Ogden, W.C., Davis, M. W., and Rice, S. Document thumbnail visualization for rapid relevance judgments: when do they pay off? *TREC 1998*, 528-34, 1998.
14. Robertson, G., Czerwinski, M., Larson, K., Robbins, D. C., Thiel, D., and van Dantzich, M. Data Mountain: using spatial memory for document management. *Proc. UIST '98*, 153-62, 1998.
15. Rosenholtz, R. A simple saliency model predicts a number of motion popout phenomena. *Vision Research*, 39:3157-63, 1999.
16. Treisman, A., and Gelade, G. A feature-integration theory of attention. *Cog. Psych.*, 12:97-136, 1980.
17. Woodruff, A., Faulring, A., Rosenholtz, R., Morrison, J., and Pirolli, P. Using thumbnails to search the Web. *Proc. SIGCHI 2001*, 198-205, 2001.
18. Yantis, S., and Jonides, J. Abrupt visual onsets and selective attention: evidence from visual search. *J. Exp. Psych: HP&P*, 10:601-21, 1984.
19. Yantis, S., and Jonides, J. Abrupt visual onsets and selective attention: voluntary versus automatic allocation. *J. Exp. Psych: HP&P*, 16:121-34, 1990.