

Looking Ahead: A Comparison of Page Preview Techniques for Goal-Directed Web Navigation

Aaron Genest¹, Carl Gutwin¹, Adrian Reetz¹, Regan Mandryk¹, David Pinelle² and Andre Doucette¹

¹Department of Computer Science, University of Saskatchewan, 110 Science Place, Saskatoon, SK, Canada, S7N 5C9

{aaron.genest, regan.mandryk, adrian.reetz, carl.gutwin, andre.doucette}@usask.ca

²National Research Council, 46 Dineen Dr., Fredericton, NB, Canada, E3B 9W4
david.pinelle@nrc-cnrc.gc.ca

Abstract. On the World Wide Web, page previews augment hyperlinks to provide extra information about each link's destination. These previews can reduce navigation time and errors in goal-directed navigation tasks when the information provided by the text and context of links is inadequate. A number of different types of page previews have been proposed, and some are already in use; however, little is known about which preview types will consistently help users make good navigation decisions. Our study compares six preview techniques (title, URL, subject category, page genre, genre symbol, and thumbnail), two delivery mechanisms (inline and popup), and two page load times (fast and slow). We found that previews showing the genre of the page (e.g., whether the page is an information page or a search page) yielded significantly faster performance than other preview techniques, and participants also preferred the genre-based previews. Our study is the first to compare the performance of a wide range of page previews in a naturalistic, non-search environment, and provides empirical data that can improve support for goal-directed navigation.

Keywords: Goal-directed browsing, information scent, page previews.


1 Introduction

Goal-directed browsing – in which users look for information that satisfies a specific information need – is a common activity on the World Wide Web (WWW). In these tasks, people try to move towards their goal by navigating through successive Web pages, making a decision at each page about where to go next. If the user has experience with a page or destination, these decisions could be based on prior knowledge, but in the majority of cases, a page will be unfamiliar to the user, so they will have to choose an outgoing link based on information presented on the current page.

This local information about where links will take the user is called *information scent* – that is, proximal information that gives cues about distal information sources (i.e., other Web pages) [4]. The primary means for providing information scent in a Web page is through the link text itself – and in many cases, link text does an ade-

quate job of indicating where the link goes: e.g., “For more information, visit the [SuperCorp website](#)”. However, many other pages have links with poor information scent, and provide few clues about where links will take the user. Many blogs, news sites, and review sites suffer from this problem of hyperlinks with low scent.

When links have low information scent, it becomes difficult for users to make good navigation decisions; as a result, people need more time to decide which link to follow, and may make more navigation errors. In high-cost environments (such as a slow network link, or an expensive roaming data plan), extra time and more errors can quickly lead to frustration and added expense.

To deal with this problem, Web designers can incorporate better information scent into the links on their pages (e.g., [18, 25]); however, automatic provision of additional scent could be an important alternative for existing pages, or for pages that do not follow ideal design principles. Insertion of additional scent is a realistic possibility – because browsers can retrieve information about the destination page, it is possible to add a *destination preview* as part of the information scent for each link on the page. Most browsers already provide a limited form of this kind of preview, by showing the link URL in the window’s status bar. There are also several other types of previews that are possible. For example, some pages augment links with an icon to indicate that the link will take the user offsite (e.g., “[example link](#) ”), and page scripts can allow the display of a thumbnail image of the actual destination page on mouseover (as in Figure 1B). Other possibilities include showing the page title, a text summary of the page, or a word or symbol indicating the page type or subject.

Although destination previews are already possible, little is known about what preview type will work best – that is, which kind of preview will most consistently help users make good decisions about where to go. Prior work has tested previews for search results pages, but no research has looked at preview techniques in naturalistic Web pages such as blogs or reviews. To address this limitation, we carried out an empirical study that compared the performance of six different preview types in realistic goal-directed navigation tasks, and tested them with fast and slow page load times and in both pop-up and inline formats.

Three of our destination previews can be derived from the destination page itself: URL, page title, and thumbnail. Three others involve human or automatic classification of pages into categories – in our study, either the subject category of the page or the page’s genre (e.g., information, shopping, or search).

Our study showed that there were significant differences between the preview types: performance with genre-based previews was significantly faster than with the other preview types, and participants preferred genre-based previews. In addition, thumbnails had the slowest task times of all the methods, and were the least-liked technique. In this paper, we provide details of these findings and present guidelines for the design of Web pages and page preview techniques.

2 Related Work

Information foraging theory considers how people navigate as they seek information in spaces like the WWW [4]. An important part of information foraging is the idea of

information scent – ‘proximal clues to distal content’ that help the user understand the likely outcome of each navigation action [4, 20]. When information scent is poor, users exhibit indecision and hesitation when deciding between hyperlinks. When they do not have the information that they need to make good navigation decisions, users often randomly click on links, and must backtrack to correct their mistakes [25].

Poor information scent is common on the Web, and evidence for the errors that this can cause can be seen in recent studies of backtracking behaviour [2, 14, 23]. For example, Obendorf et al. found that backtracking accounted for 31% of page revisits, and they also suggest that people actively spend time trying to avoid backtracking [19]. In recent years, developers have released software aimed at reducing backtracking and related navigation errors. For example, Berkun [1] stated that navigation errors helped to motivate the development of the Explorer bar, which helps manage favorites, browsing history, and search results.

Navigation errors can be frustrating and potentially expensive since each page view can cost time or money. For example, navigation errors are particularly time consuming for visually impaired users: it can take up to 120 seconds to interpret a new page [10]. Similarly, users on pay-per-bit networks or networks with bandwidth constraints pay substantial penalties when they make hyperlink selection errors.

Providing users with information about the destination of a hyperlink can improve navigation [9, 17, 26]. Adding this preview information enhances information scent so users can make more accurate decisions about target pages [25]. Nielsen [18] suggests that Web designers should strive to improve information scent by using links and category text that accurately characterize the destination page, but others have taken this idea further, considering several different presentation techniques, including thumbnails, symbols, and genre information [16, 24, 26].

Weinreich and Lamersdorf [24] list several types of information that can be used to provide previews of destination pages, including basic information about the page (such as the title, author, language, or last visit), the direction of navigation (forward or backward), the expected access time, the target file type, the behavior of the link (will the link open a new window, replace the current one, or change a frame), the location of the page (remote or local), and the status of the link. Drori [7] found that a combined preview of the title, relevant text from the target page, common words, and keywords in context was effective at improving completion times during search tasks.

Another preview approach is to use high-level semantic categories to describe target pages. Matsuda and Fukushima [16] developed a system that classifies pages based on keywords and the page’s internal structure. They classify pages according to common page genres, such as ‘product catalog,’ ‘investigation report,’ and ‘advertisement.’ Other research developed similar genre classifications, with different criteria, including fine-grained information about the style and layout of page content [15], and information about the placement of links on pages and their URLs [22]. Clustering search results based on subject categories can improve speed and accuracy for search tasks [3]. Generating subject or genre classifications accurately and efficiently is still an active research area [27], but explorations into syntactic classification [16] and semantic hierarchical classifications [8] have shown promise.

A third common preview technique is to provide thumbnails that show a miniature representation of the destination page. Some researchers have found that thumbnails improve recognition during bookmarking tasks [5, 13]. More recently, Woodruff et al.

[26] evaluated an enhanced thumbnail technique that highlights page text, and found that people were able to complete certain search tasks more quickly with the enhanced thumbnail than with either text alone or standard thumbnails.

Researchers have also developed preview techniques to dynamically modify the current page to provide information about destination pages. Faaborg et al. found that determining the user's goals and adaptively altering or adding to the content of the current Web page can reduce completion times in goal-directed navigation [9]. Miura et al. [17] expand potential page targets within the current browser window to create documents within documents. However, this technique can create large unmanageable spaces that force users to scroll extensively within a page.

There are several examples of page previews already used on the Web: most browsers show URLs in the status bar when the user hovers over a link; plug-ins and page scripts are available for adding thumbnails; and Web sites such as Wikipedia add a symbol to indicate that a link will take the user to an external page.

Previous studies have compared performance differences between page preview techniques, but these evaluations have focused on Web search scenarios [8, 16, 26]. For example, Woodruff et al. [26] compared the performance and subjective differences when people use thumbnails, enhanced thumbnails, and text summaries in Web searches. Dumais et al. [8] compared several different search interfaces, which combined different preview techniques, including category names, text summaries, and page titles, using both inline and hover presentations.

Unlike past studies, we are interested in evaluating preview techniques when they are used with realistic Web pages. Our goal is to extend the understanding of preview performance to tasks and situations that more closely reflect general goal-directed navigation on the WWW.

3 A Comparison Study of Page Preview Techniques

To determine performance and preference differences between different types of page previews, we carried out a study that asked participants to make navigation decisions in realistic Web pages with six different preview techniques. In particular, we wanted to answer the following questions:

1. Does the preview type affect the time needed to navigate to the correct page?
2. Does the preview type affect the number of navigation errors?
3. Which preview types do users prefer?
4. Do people use the previews or the page content when making navigation decisions in low-scent situations?
5. Does changing the presentation method (inline or popup) change the performance of any preview type?
6. Does changing the page load time change the way people use previews?

To answer these questions, we manipulated three independent variables in our study: preview content, presentation method, and page load time.

3.1 Independent Variables

We tested six different kinds of preview content: destination page title, URL, thumbnail image, subject category, genre category, and genre symbol.

Title. Title previews showed the full title of the destination page, as retrieved from the page's HTML header.

URL. The full URL of the destination page.

Thumbnail image. Thumbnail previews showed 128x128-pixel images of the destination page (similar to that shown in Figure 1B).

Subject category. Subject previews are short tags (one to three words) describing the semantic category of the destination page (e.g., 'hockey' or 'politics'). We selected subject categories from Google Directories (www.google.com/dirhp); for example, www.golfknickers.com has a directory path of Shopping > Sports > Golf > Apparel, so we used the tag 'Golf apparel shopping.'

Text genre category. Text genre categories are a single word that describes the genre of the destination page; they classify the page type, but not the page subject. Eight genre categories – developed through a series of pilot studies – were used in the study: login, news, shopping, search results, blog, forum, corporate, and information.



Fig. 1. A) Sample of page with inline genre text links; B) Popup thumbnail preview



Fig. 2. Genre symbols and names

Genre symbols. Genre symbols use the same categories as genre text, but use a symbolic representation (see Figure 2).

Two presentation mechanisms, inline and popup, were used in the study. Inline presentation showed the preview content at all times, located to the right of target links (Figure 1A). Inline URLs and titles were truncated to 30 characters. Popup presentation showed previews when the mouse pointer was over the link (Figure 1B).

Two page load times were simulated to determine whether delay influenced the performance or preference of the preview types. In the slow page-load condition, there was a 10-second delay between clicking on a link and the delivery of the destination page. The fast-load condition had no delay.

3.2 Task

The previews were deployed in a series of custom-built Web pages that resembled simple blog or news sites. In all pages, we manipulated the content so that the links on the page had poor information scent. We equalized the amount of text shown in links, and the amount of information provided by the surrounding text and the link text. Each hyperlink was reduced to one word. These manipulations were determined through pilot studies where users judged the link scent on each page.

Source pages provided the same set of sidebar and topbar links (these links appeared live, but returned the participant to the source page). There were no animations or advanced visual objects (e.g., Flash) on the source pages.

Participants were asked to carry out short navigation tasks in which they had to decide which link on a source page would move them towards a given information goal. Participants carried out 30 of these tasks in different experimental conditions, but in all cases the source page had links with poor information scent. An example source page is shown in Figure 1A. The task for this page was to select the link that would best allow them to “find out the average impact resistance of a rubber sidewalk slab.” Participants were told to decide as quickly as possible while minimizing errors. The information goal was always available in a separate browser window.

Articles were actual posts found in news and blog sites. The text was approximately the same length for each article, and all pages could be read without scrolling the browser window. Links were added or removed to achieve a uniform four links per article, and only one link was the correct choice for each task.

After participants selected a link from the source page, they were shown a form asking what influenced their decision (the article’s content, the link preview information, both, or neither). Participants were then told whether they had chosen the correct page; if they were incorrect, they were reminded of the information goal and directed back to the source page. Participants continued until they chose the correct link.

To simulate expert performance with the symbols in the ‘genre symbols’ condition, participants could see a reminder window with the symbols and their names.

3.3 Study Design and Procedures

The study used a mixed factorial design with both between- and within-participants factors to reduce the length of the study to under an hour per participant. The within-participants factor was preview type (page title, URL, thumbnail image, subject cate-

gory, genre text, genre symbol). The two between-participants factors were presentation method (inline or popup) and page load time (fast or slow), since pilots suggested that both factors would have easily measurable differences between conditions.

Each participant used each technique on five Web pages (for a total of 30 pages). Each technique was shown as a block: a participant would see five pages in a row for each preview type before moving on to a different type. The order of the six preview techniques was counterbalanced using a Latin square, but tasks and pages were presented in the same order for all participants, balancing the six preview techniques across the 30 task pages.

Presentation method and page load time were between-participants factors. Half the participants saw preview techniques inline, and the other half saw them as popups. In addition, half the participants were assigned to the slow-loading condition, and half to the fast-loading condition. In the slow-load condition, participants were redirected to a page that indicated that the content was loading during the 10-second delay; the loading page was not used in the fast-loading condition.

Participants began by completing a demographic survey and an informed consent form. Participants read a short description of the study procedure and viewed examples of each of the preview techniques using the appropriate presentation method. Following the orientation, participants completed the set of 30 navigation tasks using the Firefox Web browser. At the end of the session, participants completed a post-experiment questionnaire that asked them to rate and rank the techniques according to preference and usefulness. The questionnaire also asked whether participants used the text of the page, the previews, or both in their navigation decisions. Participants were also permitted open-ended responses.

Twenty-four participants, 14 female and 10 male, were recruited from a local university. Participants ranged in age from 18 to 28 years (mean of 22). All participants reported either using a computer 'often' (1) or 'every day' (23), and reported using a Web browser 'occasionally' (1), 'often' (1), or 'every day' (22). As there were two between-participants factors, there were four groups of users (6 users per group): popup-delay, popup-no-delay, inline-delay, and inline-no-delay. Kruskal-Wallis tests on 4-independent samples showed no differences between the groups in terms of their expertise with computers ($\chi^2=3.0$, $p=.392$), video games ($\chi^2=2.6$, $p=.451$), Web browsers ($\chi^2=2.1$, $p=.553$) or e-mail ($\chi^2=2.1$, $p=.554$), or with their choice of primary Web browser ($\chi^2=6.2$, $p=.101$) or primary input device ($\chi^2=1.8$, $p=.614$).

Computer logs were used to track completion time, error rates, and participants' mouse movement over target links. Completion time was measured from the moment the task page loaded to the moment the participant selected a target link. In the delay condition, the 10-second delay was not included in the completion time measure. Pages were served from a machine connected on a local area network, so additional latency was negligible. Incorrect links selected during a task were logged as errors.

We performed a Repeated Measures Multivariate Analysis of Variance (MANOVA) on the completion time and error data for the six preview types, with presentation method and page load time included as between-subjects factors. All main effects and interactions were tested at $\alpha=.05$ and Least Significant Differences were used for all post-hoc pairwise comparisons. The sphericity assumption was not violated, thus the degrees of freedom were not adjusted. Questionnaire data were analyzed using non-parametric statistical techniques appropriate for rankings data.

4 Results

4.1 Empirical Results

Does preview technique affect task completion time? Preview technique had a significant effect on the task completion time ($F_{5,100}=3.50$, $p=.006$, $\eta^2=.15$; see Figure 3A). Pairwise comparisons showed genre text yielded significantly faster performance than all other preview types (all $p<.025$), while there were no differences between the other preview types. Table 1 shows the means and standard errors for each preview type.

Does preview technique affect error rates during tasks? Although error rates for the six preview types showed similar trends as completion time, there were no statistically significant differences between the preview techniques with respect to their error rates ($F_{5,100}=1.16$, $p=.337$, $\eta^2=.06$), see Figure 3B.

Does page load time affect task completion time, error rate, or mouse movements? There were no main effects on completion time ($F_{1,20}=.315$, $p=.581$, $\eta^2=.02$) or error rate ($F_{1,20}=.749$, $p=.397$, $\eta^2=.04$) by varying the simulated network delay. The 10-second delay between clicking the link and discovering whether or not the answer was correct appeared to have no effect on user behavior (see Table 1).

For participants who saw the preview techniques as popups, we investigated whether page load time impacted the number of times participants hovered over the links before making a decision. An RM-ANOVA on the preview types with delay as a between-subjects factor showed no significant difference ($F_{1,10}=0.67$, $p=.432$, $\eta^2=.06$).

Does the presentation method affect completion time? There were no effects of the presentation method on completion time ($F_{1,20}=2.15$, $p=.158$, $\eta^2=.10$) (see Table 1).

Does the presentation method affect the error rate? Error rate did vary significantly with respect to presentation method ($F_{1,20}=4.62$, $p=.044$, $\eta^2=.19$). Participants who were presented with the preview techniques inline had a higher error rate than those who saw popups (means and standard errors are presented in Table 1).

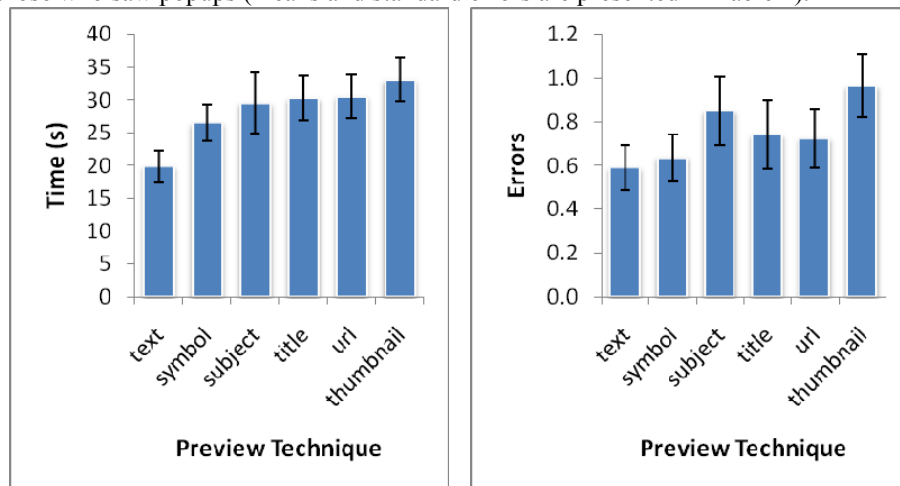


Fig. 3. A) Mean time (\pm SE) per task by preview technique; B) Mean errors (\pm SE) per task by preview technique

Table 1. Means \pm SE for time (s) and errors by page load time and presentation method

	Presentation Method		Page Load Time	
	Time	popup	24.6 (36.1)	delay
	inline	32.1 (36.1)	no-delay	26.9 (36.1)
Errors	popup	0.6 (0.1)	delay	0.8 (0.1)
	inline	0.9 (0.1)	no-delay	0.7 (0.1)

Does preview type interact with the between-subject factors (page load time, presentation method)? There were no significant interactions between preview type and page load time (Time: $F_{1,20}=1.0$, $p=.425$, $\eta^2=.05$; Error: $F_{1,20}=1.11$, $p=.358$, $\eta^2=.06$), preview type and presentation method (Time: $F_{1,20}=.69$, $p=.635$, $\eta^2=.03$; Error: $F_{1,20}=1.24$, $p=.295$, $\eta^2=.06$), or presentation method and page load time (Time: $F_{1,20}=1.37$, $p=.256$, $\eta^2=.06$; Error: $F_{1,20}=.18$, $p=.675$, $\eta^2=.01$).

4.2 Subjective Results

Which preview techniques were most liked and disliked? Participants rated their agreement with the statement “I liked this preview technique” on a 5-point scale (see Figure 4). The ratings were significantly different for the six preview techniques ($\chi^2(5)=15.6$, $p=.008$), and pairwise comparisons showed that participants liked thumbnails less than all other techniques (all $p<0.05$). There were no significant differences in the ratings between users in the slow page load versus the fast page load condition; however, users who saw the preview techniques inline liked page titles less than those who saw the previews in the popup condition ($p=.027$).

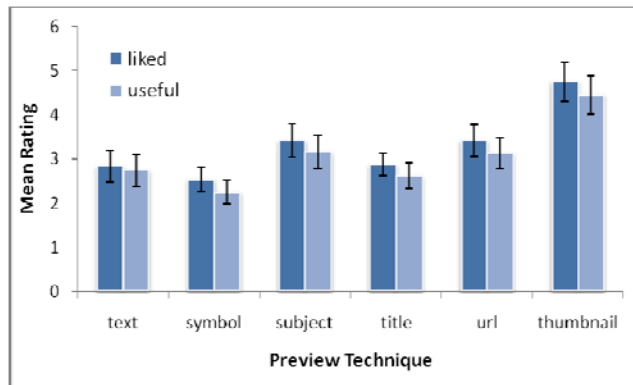


Fig. 4. Mean ratings (\pm SE) for how liked and how useful the preview techniques were on a 5-pt scale (lower rating = more liked and more useful)

In addition, participants chose the preview techniques that they liked the most and the least. The most liked techniques were genre symbol (9 people) and genre text (7). In followup comments, participants suggested that they found the symbolic representation faster than reading text: one participant said “once you learn the symbols it’s the

fastest and easiest approach” and another stated “recognizing symbols when searching was a lot simpler than reading words.”

The least liked technique was thumbnails (9 people), followed by URLs (8). Comments suggested that the fidelity of the thumbnails made it difficult to gain useful information: one person said “I could usually not tell what the page was, or even the type of page.” and another said “The thumbnails were small and I could only [get] a general impression of what the website was.” Reasons for disliking URLs included difficulty in reading and interpreting the URLs, and the lack of useful information: “It does not really tell me whether it is related to what I am searching for.”

Did participants find the preview techniques to be useful? Participants rated their agreement with the statement “This preview technique was useful” on a 5-point scale (see Figure 4). The ratings were significantly different for the six preview techniques ($\chi^2(5)=12.4$, $p=.030$), and pairwise comparisons showed that participants thought thumbnails were less useful than all other techniques (all $p<0.05$).

Did participants use the preview techniques or the page content to make their navigation decisions? After each link selection, we asked participants to report what information they used to make their navigation decision. Participants most often stated that they used previews (51%), followed by page content (22%), then both (16%), and neither (11%).

4.3 Summary of Results for Preview Types and Presentation Method

Genre Information. Participants were significantly faster when using the genre techniques, and fastest with text-only genre labels. There were few individual pages where this preview type performed badly. We believe that this indicates an important aspect of goal-directed navigation: users want to know what they can *accomplish* at a destination page, and that information is not currently conveyed by other page preview techniques. When describing the genre techniques, people used words like “simple”, “easy”, and “quick”, and many participants identified the relationship between the task and the genre text as being particularly helpful. One participant wrote, “if I was to look for something to buy, it would obviously be in the link that had the genre text ‘shopping’ rather than ‘forum’.”

Subject Categories. Subject categories appeared to work well with some tasks, and poorly with others, and overall this technique was slower than genre information. Inspection of individual pages suggested that although the category accurately represented the topic of the destination pages, it was not often a good fit for the task.

URLs. Participants’ performance with URLs was comparable to performance with subjects, titles, and thumbnails. This technique was most successful when there was a fortuitous match between an aspect of the task and an element of the preview. For example, one task asked people to find out how a person’s name could be easily changed, and the URL of the correct target contained “www.namechange.com;” similarly, a task asked participants to determine a property of a type of concrete block, and the correct destination page’s title was “Building Materials.” These fortuitous correspondences were infrequent; URLs usually provided very little information that was relevant to the task. Although URL preview information is included in most browser implementations, no users identified it as their favorite technique and several ranked it

last. Our results also suggest that ‘re-written’ URLs (such as tinyurl.com or dynamic javascript links) will detract from the browser’s existing default page preview technique (showing link URLs in the status bar).

Page Titles. User performance with page titles was also relatively poor, and this technique was not preferred by any participants. One problem with page titles is that title information is set by Web designers, so some pages can have titles that contain poor information scent—either accidentally or through intentional misdirection. The pages used in this study were carefully chosen to include valid titles, but the titles were not necessarily good matches for users’ tasks.

Thumbnails. Completion times using thumbnails was the slowest of all techniques. People also gave thumbnails the lowest ratings for usefulness and preference – significantly lower than all other preview types. There was, however, a wide variance in performance with this technique. Thumbnails appeared to work well when the destination page contained distinctive visual elements (e.g., a company logo), and when the title of the page was large enough to be readable. When destination pages were less recognizable, however, thumbnails performed very poorly: in several tasks, all of the four destinations had similar visual appearances, making it difficult for participants to make a choice based on thumbnail alone. The value of having readable text in a thumbnail suggests that the technique could be improved by creating better thumbnails – for example, using textually-enhanced thumbnails (e.g., [26]).

Presentation Method. The popup presentation of previews performed better overall than inline previews – there was no time difference, but popups led to fewer errors. The good performance of popups is valuable for designers, since many available preview techniques cannot feasibly be placed inline, since they clutter the display and may change the author’s intended page layout.

5 Discussion

In this section, we identify lessons for designers, and comment on issues that must be considered in a more widespread deployment of page previews.

First, our work provides guidelines that should be considered by designers of Web pages and browser add-ons.

- *Previews are valuable in low-scent environments.* It is clear from the study that all the preview techniques led to decisions that were much better than chance, and all of our participants reported using the previews in their decisions. Therefore, browsers should continue to make standard previews available, and should continue to investigate previews that are more effective than simple URLs.
- *Genre categories are valuable.* The genre of a page provided our participants with useful information for goal-directed browsing, but genre is not used in any current system. There is still considerable work to be done to achieve accurate and automatic genre classification, but this is an area that shows considerable potential for improving navigation support.
- *Use popup presentations.* User-controlled delivery of previews, such as a popup, is better than inline presentations for several reasons: they led to fewer errors, they do not alter the source page, and they do not add scent where none is needed. They

give Web designers more control over the look and feel of their pages because previews are delivered on demand without changing the page layout.

- *Thumbnails alone perform poorly.* Our study reinforces earlier findings [24] suggesting that thumbnail images alone should not be used as primary previews – they performed poorly and were disliked by participants. This is an important result, since several existing tools provide exactly this type of preview. Designers who wish to use thumbnails should include additional information such as title along with the image.

Second, our experiences with page previews suggest several questions that must be considered when deploying page preview techniques more widely.

- *Will previews clutter or detract from some Web pages?* In situations where the page presents adequate information scent, or where the page author has carefully designed the visual appearance, extra information could distract the user or clutter the page. This possibility argues strongly for popup presentations, which avoid showing preview information until explicitly requested, and which do not change layout.
- *Will automatic information scent clash with intrinsic scent?* Designers may have carefully added information scent to links in their page in order to create a particular impression on the user, and may not want additional previews. Although in general it is always legitimate to provide more information to users who want it, this situation raises the larger issue of whether preview information can be trusted.
- *Can preview information be manipulated?* If the authors of destination pages know how previews are being generated, they could manipulate their pages to change their preview. In some cases, this understanding will be beneficial, in that it will lead to better previews (e.g., readable text in a thumbnail, accurate page titles). However, authors could also create a falsely attractive preview to try and get more traffic.
- *Would combination previews work well?* Several types of previews seemed to work well in different situations, suggesting that multiple information sources could provide better performance. To test this possibility, we recorded the minimum times for any of the thumbnail, URL, and title previews. This hypothetical combined preview would have ranked second in performance. Combined previews have been previously explored in the context of search tasks and were found to be effective [7, 26].
- *Is automatic genre categorization feasible?* The performance of genre as a preview is dependent on being able to correctly categorize pages. Our study categorized pages by hand, but research on automatic classification is underway [8,16]. Assigning a genre is hard for some types of pages, but is easy for others (such as login or search pages) that give clear cues. It is also possible that authors could supply the page genre (although with the same manipulation risks discussed above), or by a community effort such as that for subject categories (the Open Directory).

6 Conclusions and Directions for Future Work

During task-based navigation on the World Wide Web, links with low information scent can cause navigation problems for users. Page previews can help address these problems by improving navigation time and accuracy, but little is known about the effectiveness of various proposed preview techniques. We evaluated six different

preview techniques with two page load time conditions and with two presentation methods. Our study is the first to move beyond search results pages to investigate preview performance in realistic Web sites.

We found that participant performance was best using genre-related preview techniques, and that this technique was preferred by the majority of users. Performance was slowest using thumbnail previews, and most participants disliked this technique. Our results suggest that while any preview information is valuable in low-scent environments, previews that consider the user's activity can be particularly useful. In addition, the study showed that popup presentations have advantages over inline previews, and that page load times do not change performance or preferences.

There are several possible directions for future work in this area. First, we plan to explore other preview techniques and combinations of the methods investigated here, and test them with a larger pool of participants. Second, we will continue to develop the genre preview technique. Preliminary efforts have been made on automated genre classification [8,16], and we are currently working towards low-cost techniques for classifying Web pages with genre information, through both automated and social approaches. Third, we will investigate the longitudinal effects that page previews have on browsing behaviour, through an experiment that logs navigation activity with and without page previews. Fourth, it is worth determining whether the presence of previews might actually hinder navigation when the information scent contained in links is adequate: we will investigate the relationship between information scent and page previews in navigation tasks. Last, we plan to carry out further studies with eye tracking systems to obtain more detailed information about the way that users interact with page previews.

7 References

1. Berkun, S.: The Explorer Bar: Unifying and Improving Web Navigation. In: Proc. Interact, pp.156--162. IOS Press (1999)
2. Catledge, L.D. and Pitkow, J.E.: Characterizing Browsing Strategies in the World-Wide Web. Computer Networks and ISDN Systems. 27, pp.1065--1073. (1995)
3. Chen, H., Fan, H., Chau, M., Zeng, D.: MetaSpider: Meta-Searching and Categorization on the Web. J. Am. Soc. Information Science and Technology. 52(13), pp.1134--1147. (2001)
4. Chi, E. H., Pirolli, P., Chen, K., and Pitkow, J.: Using information scent to model user information needs and actions and the Web. In Proc. CHI, pp.490--497. ACM Press (2001)
5. Cockburn, A., Greenberg, S., McKenzie, B., Jasonsmith, M., and Kaasten, S.: WebView: A Graphical Aid for Revisiting Web Pages. In: Proc. OZCHI. (1999)
6. Drori, O.: Display of Search Results in Google-based Yahoo! vs. LCC&K Interfaces: A Comparison Study. In: Proc. Informing Science, pp. 309--320. (2003)
7. Drori, O.: How to Display Search Results in Digital Libraries-User Study. Proc. Workshop on New Developments in Digital Libraries, pp.13--28 (2003)
8. Dumais, S. and Chen, H.: Hierarchical Classification of Web Content. Proc. SIGIR, pp.256--63. ACM Press (2000)
9. Faaborg, A. and Lieberman, H.: A Goal-Oriented Web Browser. Proc. CHI 2006, 751-760.
10. Harper, S., Goble, C., Stevens, R. A Pilot Study to Examine the Mobility Problems of Visually Impaired Users Traveling the Web. In: Computers and the Physically Handicapped. 68, pp.10--19. (2000)
11. Huck, S. & Cormier, W.: Reading Statistics and Research. HarperCollins, New York (1996)

12. Hurst, M. The Page Paradigm. In: www.goodexperience.com/blog/archives/000028.php, (2004) (accessed 8/15/2008)
13. Kaasten, S., Greenberg, S., Edwards, C.: How People Recognize Previously Seen Web Pages from Titles, URLs and Thumbnails. In: Proc. HCI, pp.247--265. (2002)
14. Kellar, M., Watters, C., Shepherd, M.: The Impact of task on the Usage of Web Browser Navigation Mechanisms. In: Proc. Graphics Interface, pp.235--242. (2006)
15. Lee, C.H., Kan, M.Y., Lai, S.: Stylistic and Lexical Co-training for Web Block Classification. In: Proc. Web Information and Data Management, pp.136--143. (2004)
16. Matsuda, K. and Fukushima, T.: Task-oriented World Wide Web Retrieval by Document Type Classification. In: Proc. Information and Knowledge Management, pp.109--113. (1999)
17. Miura, M., Shizuki, B., Tanaka, J.: inlineLink: Inline Expansion Link Methods in Hypertext Browsing. In: Proc. Internet Computing, pp.653--659. (2001)
18. Nielsen, N.: Alertbox: Information foraging: Why Google makes people leave your site faster. In: <http://www.useit.com/alertbox/20030630.html>. (2003) (accessed 8/31/08)
19. Obendorf, H., Weinreich, H., Herder, E., Mayer, M.: Web Page Revisitation Revisited: Implications of a Long-term Click-stream Study of Browser Usage. In: CHI, pp.597--606. (2007)
20. Pirolli, P. and Card, S.K.: Information foraging. *Psychological Review*. 106(4), pp.643--675. (1999)
21. Plaisant, C., Shneiderman, B., Doan, K., Bruns, T.: Interface and Data Architecture for Query Preview in Networked Information Systems. *Transactions on Information Systems*, 17(3), pp.320--341. (1999)
22. Shih, L.K. and Karger, D.R.: Using URLs and Table Layout for Web Classification Tasks. *Proc. World Wide Web 2004*, 193--202.
23. Tauscher, L., Greenberg, S.: Revisitation Patterns in World Wide Web Navigation. In: CHI, pp.399--406. (1997)
24. Weinreich, H. and Lamersdorf, W.: Concepts for Improved Visualization of Web Link Attributes. *Computer Networks*. 33, pp.403--416. (2000)
25. Withrow, J.: Do your links stink? Techniques for good Web information scent. *Bulletin of the American Society for Information Science and Technology*. 28(5). (2002)
26. Woodruff, A., Faulring, A., Rosenholtz, R., Morrision, J., and Pirolli, P.: Using Thumbnails to Search the Web. In: CHI, pp.198--205. (2001)
27. Zamir, O. and Etzioni, O.: Grouper: A Dynamic Clustering Interface to Web Search Results. *Computer Networks*. 31(11-16), pp.1361--74. (1999)