

How People Recognize Previously Seen Web Pages from Titles, URLs and Thumbnails

Shaun Kaasten¹, Saul Greenberg¹ and Christopher Edwards²

¹*Department of Computer Science and* ²*Department of Psychology*
University of Calgary, Calgary, Alberta, Canada T2N 1N4

Tel: +1-403-220-6087

Email: *saul@cpsc.ucalgary.ca*

The selectable lists of pages offered by web browsers' history and bookmark facilities ostensibly make it easier for people to return to previously visited pages. These lists show the pages as abstractions, typically as truncated titles and URLs, and more rarely as small thumbnail images. Yet we have little knowledge of how recognizable these representations really are. Consequently, we carried out a study that compared the recognizability of thumbnails between various image sizes, and of titles and URLs between various string sizes. Our results quantify the tradeoff between the size of these representations and their recognizability. These findings directly contribute to how history and bookmark lists should be designed.

Keywords: History system, bookmarks, web browsers.

1 Introduction

Web browsers supply various features to help people revisit their previously seen pages. These are typically some variation of a Back button, history, and bookmark facility. Excepting Back, which draws the page directly in the browser window, all facilities represent the page by some abstraction in an ordered or hierarchical list: by its title, or its URL, or more rarely as a miniature thumbnail.

Titles and URLs, of course, differ from what people see on the rendered page, and consequently they may encounter difficulties finding and recognizing the exact page they want to revisit. Titles, usually extracted from the html <Title> tag, are fraught with complications. They are often missing, are inaccurate or completely wrong, or identify the site it came from but not the page, or do not match the text actually seen at the top of the web page (Cockburn & Greenberg, 1999). URLs are similarly problematic: while they do sometimes give a human-comprehensible descriptive path and label of the current page, they are often cryptic, especially when dynamically generated by a web server (Cockburn & Greenberg, 1999a).

Cite as:

Kaasten, S. and Greenberg, S. and Edwards, C. (2002) How People Recognize Previously Seen WWW Pages from Titles, URLs and Thumbnails. In X. Faulkner, J. Finlay, F. Detienne (Eds) *People and Computers XVI (Proceedings of Human Computer Interaction 2002)*, BCS Conference Series, 247-265, Springer Verlag.

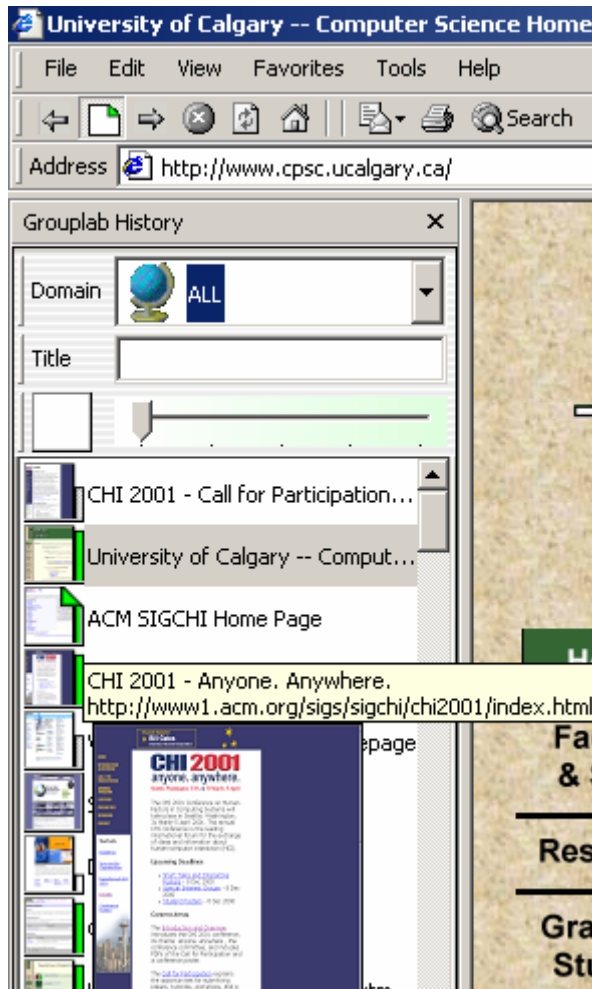


Figure 1. Our experimental system.

Additionally, because both URLs and titles can be long, most browsers will truncate them to fit within the (usually narrow) size constraints of the history or bookmark facility.

Some researchers (including ourselves) suggest that small image thumbnails of captured pages are better representations (Ayers & Stasko, 1995; Cockburn et al., 1999b, Hightower et al., 1998; Kaasten 2001; Kaasten & Greenberg 2001, Roberston et al., 1999; Woodruff et al., 2001, Suh et al, 2002). Even though they are small and of low image fidelity, they are direct representations of what the user actually saw. However, thumbnails suffer problems as well. For example, while small thumbnails allow for many to be presented in the list, this compromises their fidelity and thus their recognizability (Cockburn & Greenberg, 1999a).

We know that a revisitation system needs to display pages in a way that makes the pages easy for users to recognize. Yet we have little formal knowledge of how well people recognize web pages by title, or URL, or thumbnail. Thus designers

have created revisitation systems via hunches, guesswork, or by just copying what was done before. In our own work, for example, we are designing a system that combines Back, History and Bookmarks into a single model (Figure 1) (Kaasten 2001; Kaasten & Greenberg 2001). A sidebar lists pages as low-fidelity thumbnails and truncated titles, while a page's full title, URL, and high-fidelity thumbnail are shown in a popup window that appears when the user hovers the mouse pointer over an item in the list. Yet this seemingly simple system prompted several design uncertainties: how large should thumbnails be, should titles vs. URLs be used, how should titles or URLs be truncated to fit into the list, etc. Surprisingly, there are almost no studies excepting our own investigating thumbnail vs. title vs. URL use in web browsing (but see Czerwinski et al 1999 for their evaluation of thumbnails and text in a 3D environment).

Consequently, we decided to examine experimentally how well people recognize previously seen web pages from their titles, URLs and thumbnails. After stating four specific research questions, we describe our experimental design. We then present our results along with their implications to the design of revisitation systems.

2 Research Questions

This study investigates how well people recognize pages they have previously seen when shown representations of these pages as titles, URLs and thumbnails at various sizes. The study frames the following four research questions.

Question 1: Thumbnail Recognition and Size

Are thumbnails recognizable? What is the tradeoff between recognition vs. thumbnail size? What makes them recognizable?

We need to know how often people recognize a page by its visual appearance, and what parts of its visual features (text, fonts, layout, etc.) contribute to its recognition. However, thumbnail image size (i.e., pixels/image) is obviously an important factor (Cockburn & Greenberg, 1999). The larger the thumbnail, the more it will resemble the page the person actually saw, and the more likely the person will recognize it. Yet there is a tradeoff between thumbnail size and the number of thumbnails that can be displayed in the modest screen space typically allocated for a revisitation list. If the thumbnails are very large, the page may be recognizable but the person will be able to see only a few items in the list at a time. Finding off-screen items requires scrolling, which is tedious. If the thumbnails are too small, the person will see many of them at a glance but will also find it difficult to recognize the page from its tiny graphic: its text, embedded images and even its typographic structure may be illegible. Also, revisitation lists often collect thumbnails of pages from the same site; if these pages have a consistent visual look, people may not be able to discriminate between them because subtle differences between pages will not be discernable (Cockburn & Greenberg, 1999).

To make thumbnails useful, this first research question searches for a reasonable tradeoff that balances thumbnail recognition with space demands: at what size thresholds do thumbnails have a reasonable chance of being recognized?

Question 2: Title Recognition and Size

Are titles recognizable? If truncated, what is the tradeoff between recognition vs. title size per truncation method?

As mentioned in the introduction, titles have many problems (Cockburn & Greenberg, 1999). This second question asks how recognizable a page is from its title. We also need to know how recognition trades off with title size. In this context, size refers to how much of the title's text is visible to the user. The problem is that most practical revisitation systems are designed to occupy a conservative amount of screen width (e.g., Internet Explorer's history bar; or our own system displayed in Figure 1), and these cannot fit long titles within the narrow column. As a result, these systems truncate titles to fit within the narrow list width.

There are three different approaches to truncating: right, middle, and left. Different browsers often use different methods. Table 1a illustrates each method by example, where the title "University of Calgary -- Computer Science Home Page" is truncated to 30 letters. Notice how the title reads quite differently with each truncation method. Right truncation shows only the title's beginning, where one sees that the page is from the University of Calgary, and guesses that it has something to do with computers. Middle truncation shows only beginning and end portions, and one sees that it is from a university beginning with the letter 'C', and that it is some kind of homepage. Left truncation shows only the ending, so one sees that it is a Computer Science homepage, but not that it is from a university.

For titles to be useful, we need to know the threshold title size per truncation method that offers a reasonable chance of page recognition.

Question 3: URL Recognition and Size

Are URLs recognizable? If truncated, what is the tradeoff of recognition vs. URL size per truncation method?

The same tradeoff we see in title sizes also applies to URL sizes, and Table 1b illustrates how the left, middle, and right truncation methods are applied to a URL truncated to 30 characters. Again, each method hints at different aspects about the page—that it is from the University of Calgary in Canada (right truncation), that it refers to software (middle truncation), and that it is the software portion of the GroupLab research group (left truncation).

Question 4: Distribution of Title and URL Sizes

What is the distribution of title and URL sizes from pages typically found on the Web?

We would expect both titles and URLs between random web pages to vary greatly in their size. Some will be short, and should easily fit in even a narrow revisitation list, while others will be very long. If most titles/URLs are short then truncation is not that important. If they are long, we can expect much truncation.

Method	a) Example Title	b) Example URL
Right	University of Calgary – Comput...	http://www.cpsc.ucalgary.ca/gr...
Middle	University of C...ience Home Page	http://www.cpsc...ouplab/software
Left	...y – Computer Science Home Page	...ucalgary.ca/grouplab/software

Table 1. Truncation examples showing only 30 letters of a title and url.

What we need to know is the distribution of title and URL sizes if we are to place answers from questions 2 and 3 in context.

3 Method

We designed a controlled study to answer the above questions. For each subject, we collected a list of their recently visited pages. We analyzed the title and URL length distribution of these pages (Question 4), and then showed the subject a succession of representations of his/her previously visited pages. We first displayed each representation at a tiny size (i.e., image size for thumbnails, string length for titles/URLs). We then gradually increased this size until the subject could describe the site from which the represented page came from, and then which specific page it represented. Thus we probed for the size threshold at which representations became recognizable (Questions 1, 2 & 3). Finally, we asked people to evaluate the correctness of their responses.

3.1 Variables

Independent variables were the *representation type* shown to the subject (thumbnail, title, URL), and for titles and URLs the *truncation method* (right, middle, left). The main dependent variables were two *size thresholds*: first where the subject could identify the web site and second where they could identify the exact page. Another dependent variable was the *correctness of the identifications*, as rated by the subject. This number of correct pages yields an overall indication of each representation's recognizability. Qualitative data were the subjects' *written descriptions* about how they were able to identify the page.

3.2 Subjects

We recruited 20 paid 2nd year or higher computer science students, all practiced Internet users. As this group is likely proficient at recognizing pages by URL vs. other groups, they provide a 'best-case' scenario for URL effectiveness.

3.3 Stimuli

We wanted to test subjects' recognition of their previously visited pages. This implies two study phases: *priming* where the subject looked at a chosen set of web pages, and *testing* where the subject attempted to recognize selected pages from their representations. To reduce variability and increase repeatability, we should prime subjects with identical pages. However, finding a good set of candidate pages introduces three problems with serious implications on how we could generalize our results to browser design.

- a. **Artificiality of page interest.** Subjects may have no personal connection with the set of pages we give them. This could profoundly affect how well (or how poorly) they remember these pages. In real use, we expect people will attend to various pages quite differently due to their immediate interest or page appeal.
- b. **Artificiality of learning.** The way we ask subjects to 'learn' pages could profoundly affect how well they are remembered. We could present pages for a timed duration, or insist they read each page, or have them search the page. In

real use, we would expect people to learn pages differently: as a function of their interest, how they read them, etc.

- c. **Page composition.** Pages on the web are remarkably inconsistent. In visual terms, they vary greatly in their typographic structure (use of proximity, white space, fonts, contrast), and their graphical elements (image type, quantity, size and noise such as advertisements). Similarly, pages vary greatly in how titles and URLs are composed. There are virtually no statistics that describe common page attributes. If we ‘make up’ our own pages, the ability of people to recognize them may have little bearing on how they recognize the perhaps quite different pages on the web.

Consequently, we decided to use the actual pages viewed by subjects during normal browsing activity as stimuli, which means our study is a quasi-experimental design.

3.4 Materials

We used a high-end computer running Windows 2000 with a 17” monitor at 1152x864 resolution with 32-bit color. We used Internet Explorer 5 (IE-5) to display pages to the subject during a verification phase. All subjects had previously browsed the web using IE-5 or Netscape Navigator 4.x.

For the priming phase, our *history extraction software* extracted into a file the URL, title and last-visit date of pages that an IE-5 user had visited on their machine. Our *stimuli preparation software* read in this history file, displayed each page, let the experimenters select 30 pages, and generated a high-quality thumbnail of these pages. Our *stimuli presentation software* presented the stimuli to the subjects and recorded their responses (Figure 2), while our *stimuli verification software* let them verify the correctness of their responses (Figure 3).

3.5 Procedure

Step 1 Stimuli Preparation. First, the subject submitted their history record to us (most included pages they visited within the last 3 weeks). IE users used our history extraction software, while Netscape users invoked its history list ‘save’ option. Second, using the stimuli preparation software, we pseudo-randomly selected 30 pages and captured high-quality smoothed thumbnail images of them. We manually filtered out pages that would not load properly (e.g., slow and password-protected pages), ‘frames’ pages (which contain multiple history entries for a single page), and pages without titles. We also excluded a page if several others from the same site had already been selected.

Step 2 Stimuli Presentation. We ran a subject about 1–3 days after receiving his/her history file. The procedure for each trial began by showing the subject (using our stimuli presentation software) one of the page representations at a tiny, probably unrecognizable size. For thumbnails, this was 16^2 pixels. For titles or URLs, the initial string size was two letters. Depending on the truncation method used, this meant the subject saw the first two letters, the first and last letter, or the last two letters. Truncated URLs included their ‘http://’ prefixes, and we will discuss issues

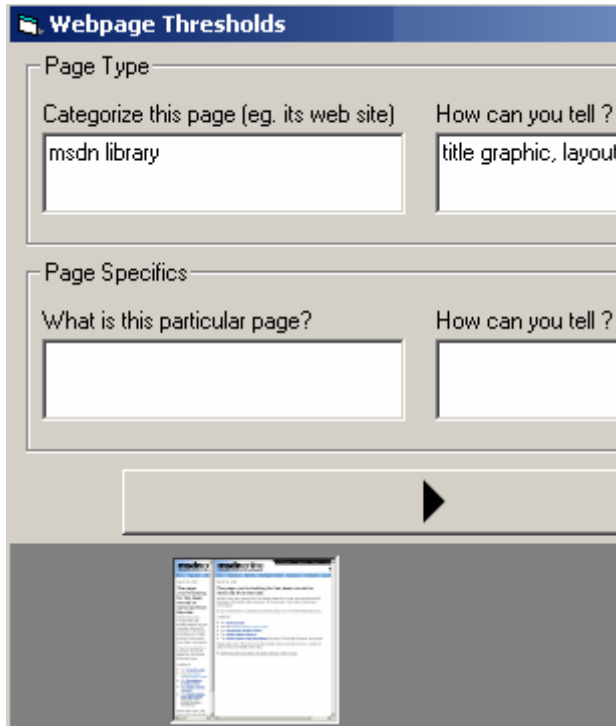


Figure 2. Subject pauses thumbnail growth at 36^2 and identifies web site.

related to this later. We then gradually increased the representation size until the subject could just recognize the web site the page came from. The subject would then continue until he or she could identify the specific page.

Figure 2 illustrates a moment in the thumbnail trial sequence. Previously, the subject had seen a 16^2 thumbnail image but did not recognize it. The thumbnail dimensions then grew automatically at a rate of 16 pixels every 3 seconds. This subject watched the thumbnail as it increased in size until she just recognized what web site it came from, in this case at size 36^2 pixels. At this point, she clicked the 'play/pause' button (Figure 2 middle) to pause the thumbnail's growth. She then typed a description of the web site (in this case 'msdn library') in the top left field labeled 'Page type', and how she recognized it in the top right field labeled 'How can you tell?' (title, graphic, and layout). Figure 2 was taken at this point. She was still uncertain about which particular MSDN page the thumbnail represented, so she pressed the 'play' button and the thumbnail continued growing. Finally, she recognized the page at size 108^2 pixels as a 'redirection page' because she could read its textual contents at this size. She clicked 'pause' and filled in this page-specific information (fields in the middle of Figure 2). She then clicked a 'next' button (not shown) to proceed to the next trial page.

The sequence for the textual representations was similar, except the text replaced the thumbnail image in Figure 2. The title or URL was initially truncated to display only two letters using one of the left, middle or right methods, and this text size then grew by two letters every 3 seconds.

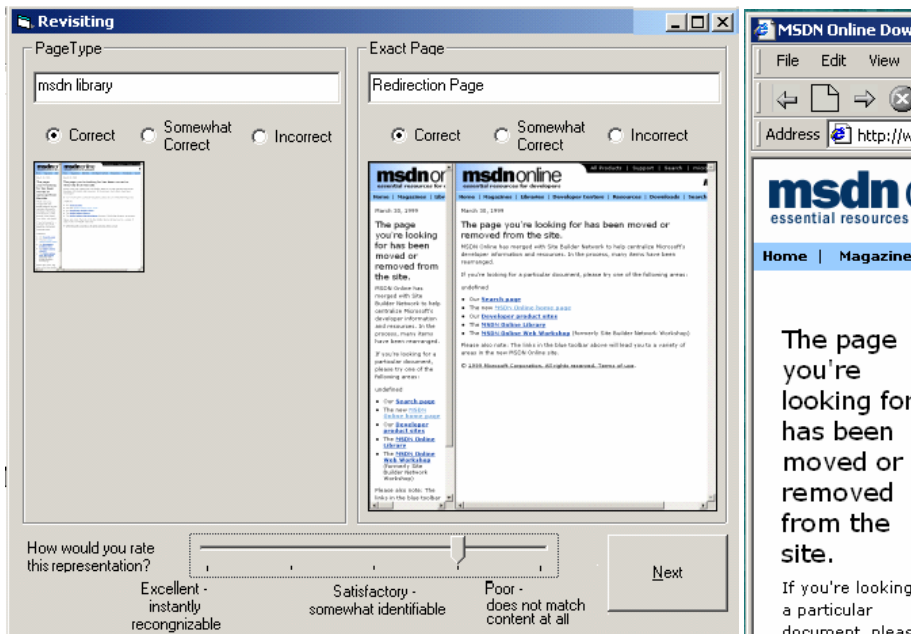


Figure 3. Subject verifies their choice and rates the representation.

The subjects saw 30 different pages; thus they had 30 trials. Each trial used only a single page and representation. Trials alternated through thumbnail, title, and URL. Titles and URLs alternated between right, middle and left truncation.

Step 3 Stimuli Verification. After completing all 30 trials, the verification process began. Using the stimulus verification software, subjects went through their responses to see if they correctly identified the pages. For each page, we showed the subject the form in Figure 3 (left window), as well as the actual Web page at full size in the IE-5 browser (right window). This form displayed that page in the representation they saw at the two sizes he/she indicated as just being able to recognize the web site (left side) and exact page (right side). The subject used this form to indicate if his/her answers were correct, somewhat correct, or incorrect (top). The subjects used the same form to rate how well the representation ‘captured’ the page (bottom). For titles and URLs, this rated the non-truncated title/URL, not just the portion that he or she saw before answering. For the thumbnail, the question referred to the page concept rather than its size—does the ‘look’ of the full sized page, as seen in the browser, give a good indication about its content?

4 Results and Discussion

4.1 Thumbnail Recognition and Size

Research question 1 asked about the tradeoff between recognition vs. thumbnail size.

Results. Figure 4 plots the threshold thumbnail sizes where people were able to just recognize the web site (top 2 lines) and the exact page (bottom 2 lines) as cumulative distributions. We call these recognition points the ‘stop sizes’. Each point gives the




Figure 4. Running sums of thumbnail stop sizes

running sum of all previous points. The dotted lines plot all responses, which we include for comparative purposes. The solid lines exclude incorrect responses—this measure shows only responses that were rated ‘Correct’ or ‘Somewhat correct’ in Step 3—and are thus a reasonable measure of thumbnail recognition. This measure will form the basis for our discussion.

Discussion. All subjects tried to recognize the thumbnail images by the time they reached 304². In total, they rated only ~10% of their guesses as completely incorrect. This 90% success rate means that (not surprisingly) people are fairly good at recognizing pages from their visual images if they are appropriately sized.

More specifically, Figure 4 provides a cost-benefit guide of the recognizability of a given thumbnail size. We premise this on the (reasonable) assumption that showing a person a thumbnail at one particular size is equivalent to the cumulative effects of seeing the thumbnail at all of its smaller sizes. That is, a larger thumbnail will be at least as recognizable as all of its smaller versions. For example, if we wanted at least 60% recognition of web sites by thumbnails, we would need a thumbnail sized at least 96² pixels or less. Choosing this size also means that people will recognize the exact page only ~35% of the time.

To ease comparisons, we establish benchmarks for recognition that will allow us to directly compare and make recommendations for thumbnail, title and URL size. We will set the benchmarks as 15%, 30%, 60% and 80% for minimum, low, medium, and high recognition levels respectively. We also include a ‘maximum’ which indicates the percentage of pages correctly recognized. Of course, developers

can choose their own benchmarks, where they can look up particular sizes directly in Figure 4.

Table 2 (last page) tabulates these benchmarks. For example, the table (and Figure 4) suggests that if space is very tight, the minimum useful size for a thumbnail is $\sim 32^2$ pixels for identifying web sites, and $\sim 48^2$ pixels for identifying exact pages. If space demands are somewhat less stringent, low recognition ($\sim 30\%$) is achieved with $\sim 48^2$ pixels for web sites and $\sim 80^2$ pixels for exact pages. For medium recognition (60%), we need 96^2 pixels for web sites and 144^2 pixels for exact pages. Finally, for high recognition (80%), we need 144^2 pixels for web sites and 208^2 pixels for exact pages. The maximum recognition we could achieve with even larger thumbnails is $\sim 90\%$.

4.2 Thumbnail Features

Research question 1 also asked what makes thumbnails recognizable at given sizes.

Results. Each time the subjects made a guess at the web site or exact page, they specified the predominant features that influenced their guess (Figure 2, top). These comments invariably dealt with the following visual attributes:

- **Colors:** background and font colors for the page;
- **Text-related:** legible text from the title or secondary titles on the form. This can be graphical (as in a banner);
- **Image-related:** a distinctive image on the page;
- **Layout-related:** distinctive formatting of page elements.

We categorized the subjects' answers into these attributes and counted how often they occurred. When subjects mentioned more than one attribute, we counted them in multiple categories. When identifying web sites, we found that 'early' identifications ($< 64^2$ pixels) were primarily due to gross page features such as colors or layout. Between sizes 64^2 – 96^2 , all four attributes were roughly mentioned at par. From 100^2 pixels up, text and to a lesser extent images predominated. This makes sense: while text and image fidelity increased at these larger sizes, gross features such as layout and colors would not change much.

The importance of text is even more evident in the exact page identifications. While layout and color are somewhat important at small thumbnail sizes, nearly all 'late' identifications (100^2 pixels and larger) mentioned text-related attributes. Yet we see from Figure 4 that only a modest number of exact page identifications took place with thumbnails smaller than 100^2 . Thus subjects needed larger thumbnails, and the vast majority of identifications were based on reading text-related cues. In fact, subjects mentioned text 90% of the time whenever they identified the exact page, compared to $\sim 30\%$ for layout and image-related attributes, and 12% for color (these do not sum to 100% because people can list multiple attributes). Thus we conclude that subjects relied heavily on reading text inside the thumbnail. Of course, this implies that the thumbnails were large enough for the subjects to read the text.

Discussion. These results suggest what thumbnail cues enable recognition. Very often, subjects identified web sites by small thumbnails (less than 96^2 pixels) through its color and layout rather than details. This is likely because many web sites have a

distinctive ‘look’ that can be recognized in a small image icon. Yet for identifying exact page, being able to read some of the page’s text was clearly important¹. What is likely is that a page’s surrounding colors, page layout and images provide the context and redundancy to make the site recognizable, while dominant text pinpoints the exact page.

These results have implications for web site and page design. First, they re-enforce the value of repeating color/layout/images across pages, for pages become recognizable as coming from a particular site. Second, if thumbnails become an important interface feature then page designers should be encouraged to use large title and banner font sizes that are visible in small thumbnails (Woodruff et al., 2001).

4.3 Title Recognition and Size

Research question 2 asked about the tradeoff between recognition vs. title size per truncation method.

Results. Similar to Figure 4, Figure 5 plots the stop-size distributions as a running sum for each of the title truncation methods when identifying web sites (5a) and exact pages (5b). For clarity, we only graph and discuss data that excludes incorrect responses. Table 2 (center) tabulates this data using our benchmarks.

Discussion. Ignoring size, people managed to correctly identify the web site between 87-93% of the time, and the exact page between 75-83%. When size is taken into account, we see that right truncation stood out as best for web site identification (Figure 5a). This is not surprising; many titles begin with the web site name, as in “University of Calgary, Department of Computer Science-Research” and the right-truncation method reveals this beginning portion.

For identifying the exact page, the discerning portion appears at the end of the title, as revealed by both the left and middle truncation methods. Thus, right truncation fairs poorly compared to the other two methods (Figure 5b). Except at very low sizes, middle truncation is slightly favored over left truncation for identifying the exact page, as seen in the running sums of Figure 5b. This suggests that both prefix and suffix slightly re-enforce recognition. Invariably, people need to see more letters of the title for identifying the exact page than the web site. For example, comparing the best-performing truncation methods between Figures 5a and 5b at 26 letters, we see that right-truncation gives us 82% recognition for web sites, while middle truncation gives only 54% recognition for exact pages; indeed we have to double the title length to 52 to bring the exact page recognition rate to 82%.

As before, these distributions allow us to make recommendations for designing a revisitation list based on titles, as tabulated in Table 2. For example, for medium (60%) recognition, we need 15–20 letters (depending on the truncation method) for web sites, and 30–39 letters for exact pages. Unfortunately, no truncation method stands out as best for both web sites and exact page identification.

¹ We could argue that that using a thumbnail to display text defeats its purpose of using graphics, for instead we could simply display the text, at a font size that is much easier to read than in a shrunken graphic. However, the text that dominates a page and therefore its thumbnail is often different from the page’s technical title.

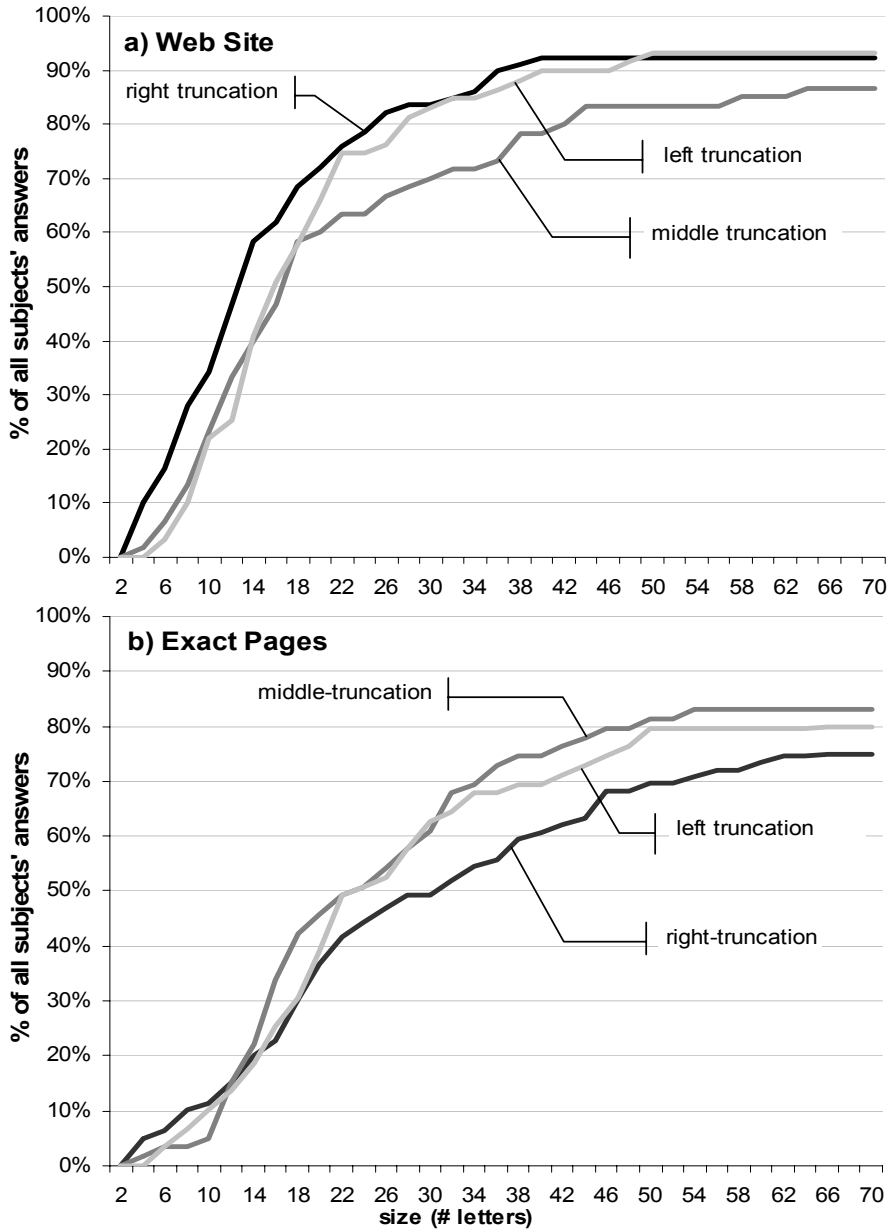


Figure 5. Running sums of title stop sizes.

4.4 URL Recognition and Size

Research question 3 asked about the tradeoff between recognition vs. URL size per truncation method.

Before answering this question, we should mention that we included the standard ‘http://’ prefix in the URLs we presented to subjects as done in several existing history systems. Unfortunately, this meant that subjects shown the right and middle truncated URL did not see any useful portion of the beginning URL until after the 7 letters in ‘http://’. In hindsight, we should have filtered off this prefix e.g., by showing ‘www.ucalgary.ca...’ instead of ‘http://www.ucal...’. (The same argument is not true for the ‘www.’ extension as it often differentiates intranet from internet pages). Consequently, we corrected our results. First, we subtracted 8 letters

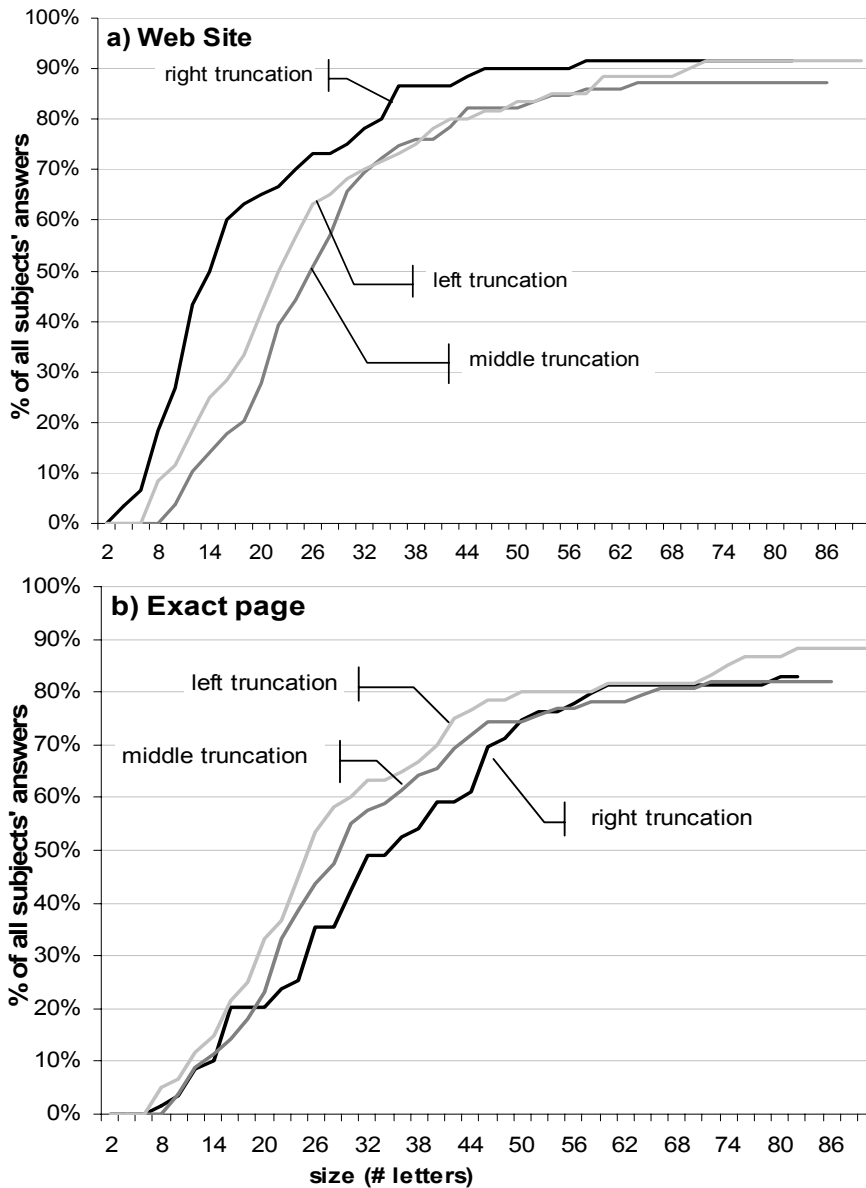


Figure 6. Running sums of (corrected) URL stop sizes

from the right truncation size (because we increased the size in multiples of two, we could not subtract exactly 7 letters). Next, we subtracted 4 letters from the middle truncation size as this method reveals both the suffix and prefix: while not a great solution, it is close enough for comparative purposes. We use and discuss only these corrected data in this paper.

Results. Similar to Figure 5, Figure 6 illustrates the stop size distributions as a running sum for each of URL truncation method when identifying web site (6a) and exact page (6b). Table 2 (bottom) tabulates these results.

Discussion. Ignoring size, people managed to correctly identify the web site from its URL between 87-92% of the time (which is comparable to titles), and the exact page between 82-88% (which is better than titles). When size is taken into account, we see that the right truncation method stood out as best for web site identification (Figure 6a). This is expected, as the web site name is often reflected within its URL prefix e.g., www.ucalgary.ca for the University of Calgary. For identifying exact page, the left truncation method proved best. This too makes sense as the suffix is often a meaningful label for the exact page.

4.5 Distribution of Titles and URLs on the Web

Research question 4 asked about the distribution of title and URL sizes of pages typically found on the Web.

Results. We analyzed the 9200 pages that comprised all submitted history records.

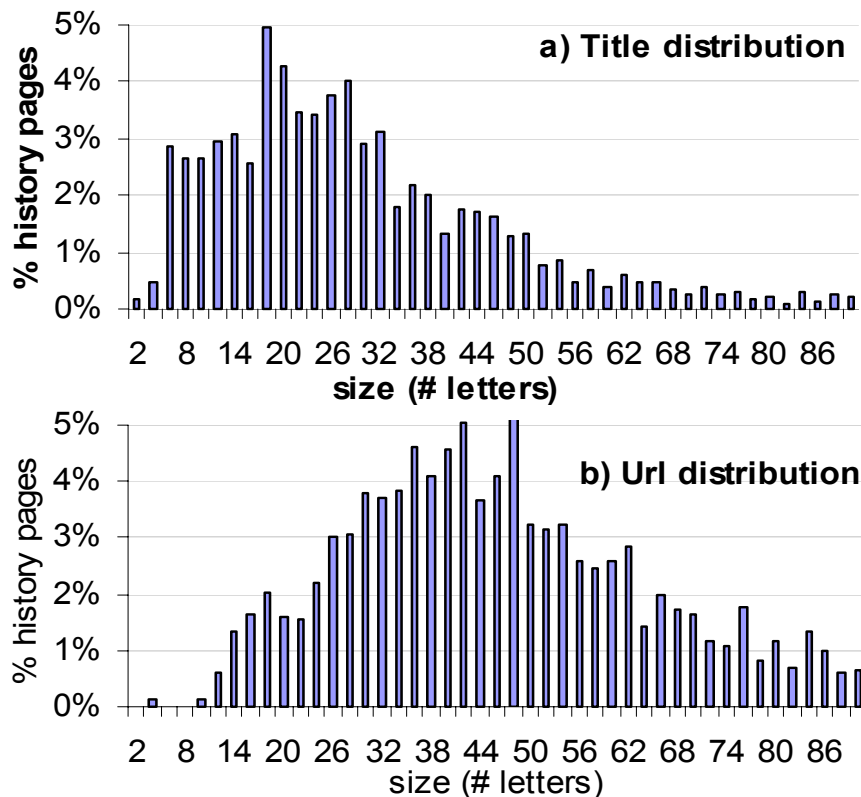


Figure 7. Distribution of title and URL lengths.

First, 30% of all pages did not have titles (remember that we excluded these from the set shown to subjects). We then plotted the title size frequency distribution of the remaining pages (Figure 7a) and the URL size distribution of all pages (Figure 7b). For pages with titles, the mean title length is 31. For all pages, the mean URL length is 40 (standard deviation=22 for both).

Discussion. 30% of these pages did not have titles. This is much higher than a recent finding that only 5% of pages lack titles (Cockburn & McKenzie, 2000). It could be that our logs included title-less popup advertisement windows often raised as a side effect of visiting a page. Clearly, this needs more study. Still, we can conclude that the overall recognition of pages by its title is between 5-30% worse than shown in Figure 5 and Table 2, as a title cannot be displayed if it is missing.

Ignoring pages without titles for now, the distributions in Figures 7a+b inform us about how many titles and URLs would be truncated if the revisitation could only display a certain number of letters. For example, if the system could only display the first 20 letters, then only 55% of the titles and 8% of URLs will fit completely (these are calculated as the running sum of all size frequencies). Thus, users will have to make decisions based on incomplete information for almost half of the items for titles, and for almost all the items for URLs. This suggests that our concerns in this study about the effects of size and truncation are valid.

4.6 Correctness / Error Rates

As a side note, we elaborate how often subjects correctly identified pages by thumbnails, titles, and URLs. Recall that we asked subjects to validate their web site and exact page guesses by scoring them as incorrect, somewhat correct, or correct (Figure 3 top). We aggregated these responses for the three different title and URL truncation methods into 'Titles' and 'URLs' categories.

Results. Statistically, there was no difference between the different error rates of the various representations (ANOVA $p > .05$, $F = 1.21$) for identifying web sites. However, subjects had fewer errors with thumbnails than with titles or URLs when identifying exact pages (ANOVA $p \leq .05$, $F = 12.21$).

Subjects proved quite accurate at identifying the web site for all representations, where they rated fewer than 10% of their answers being incorrect, another 10% as partly correct, and fully 80% as completely correct. For exact page identifications, subjects had a similar accuracy of 80% correct when using thumbnails. However, they were less accurate when using titles and URLs to identify the exact page: only 60% were completely correct.

Discussion. Thumbnails of web pages prove to be a reasonably accurate way for people to identify both the web site and the exact page. This supports the hypothesis that thumbnails are a useful representation for revisitation systems. It also suggests that our standard history lists that use only titles and URL representations are not as effective as one would like. If users only have a 60% chance of recognizing the exact content from a title or URL in the history list, they may not be motivated to invest the extra work it takes to operate the list (opening, scrolling, closing) in order to track down a page. Simply put, it is not worth the effort to switch to a list containing items where there is only a 60% chance that the desired page will be recognized.

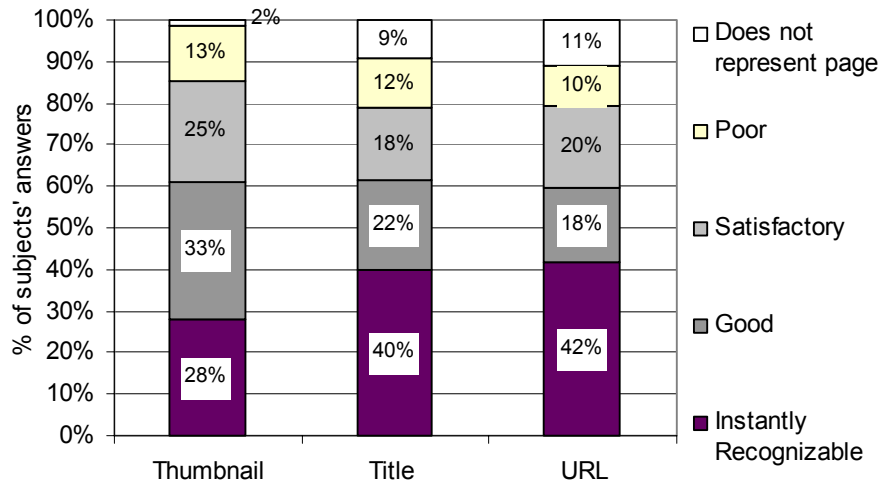


Figure 8. Subjects' ratings of thumbnails, titles and URLs

4.7 Subjects' Representation Ratings

Each subject used the 5 point scale in Figure 3 (bottom) to rate how well the thumbnail, title or URL representation 'captured' the content of the page.

Results. Figure 8 shows the results, converted to percentages. Subjects rated 15% of the thumbnails and 21% for both titles and URLs as poor or worse representations. Only 2% of thumbnails were in the 'does not represent the page' category when compared to 10% for titles/URLs. On the other side, all representations had ratings of ~60% in the good or higher category, but there were fewer thumbnails in the instantly recognizable category.

Discussion. These results suggest that thumbnails are a slightly better representation than titles or URLs. While thumbnails have marginally less instant recognition, they received generally better overall ratings. Titles are even worse than shown here because this data only includes pages with titles. We could safely assume that the extra 5-30% of pages without titles would be in the worst category.

5 CONCLUSION

In this paper, we contribute hard data valuable for the design of revisitation systems. The data directly compares URLs and titles, and shows that both have reasonable recognition rates. Because we cannot realistically show long titles and URLs in our revisitation system, the data tabulates the tradeoffs between truncation methods on page recognition e.g., Figures 4-6 and Table 2. Designers can use these to predict the consequences of using particular space-conservative displays. We also examine the recognizability of thumbnails. While they currently appear only in research systems, we have shown that people feel they are good representations of pages, and that they can accurately recognize the page they represent at particular sizes. We also showed that people mainly use a thumbnail's color and layout to identify its web site, and the

Recognition rate	size required for identification							
	Thumbnails		Titles					
	web site	exact page	web site			exact page		
			right	middle	left	right	middle	left
Minimal: 15%	32 ²	48 ²	6	8	9	12	12	12
Low: 30%	48 ²	80 ²	8	12	12	18	16	18
Medium: 60%	96 ²	144 ²	15	20	18	39	30	28
High: 80%	160 ²	208 ²	25	42	28	–	46	50
<i>Maximum</i>	92%	90%	92%	87%	93%	75%	83%	80%

Recognition rate	Urls					
	web site			exact page		
	right	middle	left	right	middle	left
Minimal: 15%	8	14	11	15	16	14
Low: 30%	11	20	17	25	22	19
Medium: 60%	16	29	25	43	34	30
High: 80%	34	43	42	58	65	50
<i>Maximum</i>	92%	87%	92%	83%	82%	88%

Table 2. Recommendations for sizes of thumbnails and titles (top), and URLs (bottom) at various benchmark recognition rates.

thumbnail's dominant text (legible only at larger sizes) to identify its exact page. However, text becomes legible only at larger sizes.

These results are also important for web page designers, for it is in their best interest to design pages that can be effectively recognized and therefore revisited easily. Web pages should have short but well named titles. The URL site and file name should be descriptive yet not long. For thumbnail images, pages throughout a single web site should have a consistent layout and color scheme. As well, a text banner should be large enough to be visible in a modest-sized thumbnail e.g., as in Woodruff et al., (2001).

Of course, more work needs to be done. First, this study looked at thumbnail recognition *in isolation* from one another. Yet in actual practice, a history list (as in Figure 1) will comprise many thumbnails, where similar-looking thumbnails for a site are likely located near each other. We believe these clusters will make thumbnails even more recognizable, meaning that our results likely suggest the 'worse case' of recognition. A next study should examine this. Second, although our study used actual pages visited, we did not separate peoples' recognition of familiar vs unfamiliar pages. In practice, we expect frequently visited pages will be more recognizable and we should test this. Third, we need to investigate the interplay and thus recognition between combined thumbnail/title/URL representations, as in the integrated history system shown in Figure 1. We expect the redundancy between representations will likely improve recognizability even further.

Finally, we need to redesign our system in Figure 1 to use these recommendations, deploy it to end-users, and evaluate its effectiveness in actual use. In our own experiences using this system, we have found the thumbnail-based integrated history system incredibly helpful, to the point where we find it quite painful to switch back to the normal history system provided by Internet Explorer.

References

1. Ayers, E. & Stasko, J. "Using Graphical History in Browsing the World Wide Web", in Goldstein, I. and Vezza, A. (eds.) *Proceedings of the 4th International World Wide Web Conference*, 1995.
2. Cockburn, A. & Greenberg, S. "Issues of Page Representation and Organisation in Web Browser's Revisitation Tools", in Scott, J. (ed.) *Proceedings of OZCHI'99, The 9th Australian Conference on Computer Human Interaction*, IEEE Computer Society Press, Australia, 1999.
3. Cockburn, A., Greenberg, S., McKenzie, B., Smith, M. & Kaasten, S. "WebView: A Graphical Aid for Revisiting Web Pages", in Scott, J. (ed.) *Proceedings of OZCHI'99, The 9th Australian Conference on Computer Human Interaction*, IEEE Computer Society Press, Australia, 1999.
4. Cockburn, A., & McKenzie, B. "What Do Web Users Do? An Empirical Analysis of Web Use", *International Journal of Human-Computer Studies* 54(6), 903-922, 2000.
5. Czerwinski, M., van Dantzich, M., Robertson, G. & Hoffman, H. "The Contribution of Thumbnail Image, Mouse-over Text and Spatial Location Memory to Web Page Retrieval in 3D", in Sasse, A. and Johnson, C. (eds.), *Human-Computer Interaction - INTERACT'99: Proceedings of the Seventh IFIP Conference on Human-Computer Interaction, Vol. 1*, IOS Press, 163-170, 1999
6. Hightower, R., Ring, L., Helfman, J., Bederson, B. & Hollan, J. "Graphical Multiscale Web Histories: A Study of PadPrints", in Gronbaek, K., Mylonas, E. and Shipman, F. (eds.) *Proceedings of the 9th ACM Conference on Hypertext – Hypertext'98*, 58-65, 1998.
7. Kaasten, S. & Greenberg, S. "Integrating Back, History and Bookmarks in Web browsers", *Extended Abstracts of ACM CHI'01*, 379-380, 2001.
8. Kaasten, S. *Designing an Integrated History/Bookmark System for Web Browsing*, MSc thesis, Dept Computer Science, University of Calgary, Canada, 2001.
9. Robertson, G., Czerwinski, M., Larson, K., Robbins, D., Thiel, D. & van Dantzich, M. "Data Mountain: Using Spatial Memory for Document Management", in Beaudouin-Lafon, M. (ed.) *Proceedings of the 12th Annual ACM Symposium on Use Interface Software and Technology, UIST'99*, 153-162, 1999.
10. Suh, B., Woodruff, A., Rosenholtz, R. & Glass, A. "Popout Prism: Adding Perceptual Principles to Overview+Detail Document Interfaces." in Terveen, L. and Wixon, D. (eds.) *Proceedings of the CHI 2002 Conference on Human Factors in Computing Systems*, ACM Press, 251-258, 2002.
11. Woodruff, A., Faulring, A., Rosenholtz, R., Morrison, J. & Pirolli, P. "Using Thumbnails to Search the Web", in Jacko, J. and Sear, A. (eds) *Proceedings of the CHI 2001 Conference on Human Factors in Computing Systems*, 198-205, 2001.