The Revisit Rack: Grouping Web Search Thumbnails for Optimal Visual Recognition

Rhys Morgan, Max L. Wilson
Future Interaction Technologies Lab
Swansea University, UK
{438124, m.l.wilson}@swansea.ac.uk

ABSTRACT

Opinion and use of thumbnails in web search is still divided, despite agreement over their particular value during re-finding tasks. In this paper we introduce the idea of a Revisit Rack that, during re-finding tasks, presents thumbnails together at the top of the page rather than beside each result, so that users can more effectively utilize visual recognition without scrolling. The results of a pilot re-finding comparison with a traditional thumbnail and text-based layout, however, were mixed. Further investigation suggests that the spatial disconnect between thumbnail and result, when a desired target is not in the Revisit Rack, may be more costly than the benefits provided when the result can be found. The study did, however, highlight several ways in which the idea of a Revisit Rack could be more formally studied in future work.

Keywords

Web Search, Thumbnails, Re-finding,

INTRODUCTION

While search engines have varied in the use, or not, of thumbnails or representative images in web search, research continues to investigate both when and where such visual representations are useful (e.g. recent work by Teevan and colleagues (2009)). This paper presents the idea of a Revisit Rack, shown in **Figure 1**, which has been designed based on the following findings and assumptions:

1. Thumbnails, in search results, provide most value when users are re-finding (e.g. Teevan et al (2009) and Woodruff et al (2001))
2. 50%-80% of web search involves re-finding (e.g. Cockburn & McKenzie (2001) and Tauscher & Greenberg (1997))
3. Visual perception is faster than semantic perception (discussed by Woodruff et al (2001)).
4. Search engines can easily track when users are re-issuing queries.

ASIST 2010, October 22-27, 2010, Pittsburgh, PA, USA.

The Revisit Rack is optimized for visual recognition and quick comparison during re-finding tasks (or re-issued queries) by grouping thumbnails together above web search results, rather than beside them individually. In the following sections, we first describe related work on re-finding and revisitation, thumbnails, and visual perception, before describing the design of our prototype and pilot study.

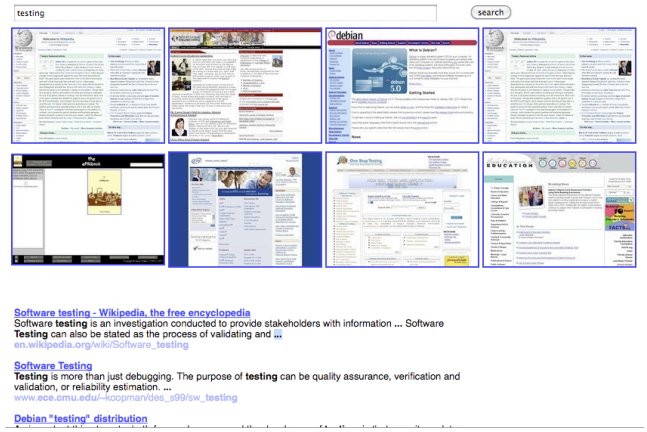

Figure 1: On detecting a re-issued query, the Revisit Rack is shown, presenting thumbnails together above the search results.

RELATED WORK

Unlike the almost consistent use of keywords-in-context textual representations of results (R. W. White, Ruthven, & Jose, 2002), the use of thumbnails and images in web results is still unstable. Thumbnails are small screenshots of a website, usually the view seen when it has finished loading, but before any user interaction. Most search engines have trialed the presentation of thumbnails by results at some stage. Currently, Google provides the option to present thumbnails, or example images, on the right hand side of each search result. Similarly, Ask Jeeves offers an on-hover display of a thumbnail for each result. Further, some browser extensions are available that automatically augment web search engines that do not provide thumbnails by default[1]. Academically, research continues to suggest alternatives to thumbnails, as discussed further below.

One of the first conclusions about thumbnails in web search was that they typically only provide value when users can

[1] https://addons.mozilla.org/en-US/firefox/addon/211

recognize results during re-finding tasks (e.g. (Teevan, et al., 2009; Woodruff, Faulring, Rosenholtz, Morrsion, & Pirolli, 2001)). While re-finding could be perceived as a minority of tasks, research suggest that between 50% (Tauscher & Greenberg, 1997) and 80% (Cockburn & McKenzie, 2001) of web search activity involves revisiting previously seen content. Further, Teevan and colleagues (2007) noted that users often rely on web search instead of proactively bookmarking results. Adar and colleagues (2009) have also shown how revisitation often varies based on the expectation of change within websites. Re-finding and revisitation are clearly significant forms of web search.

While research indicates that appropriately sized thumbnails permit 80% recognition success (Kaasten, Greenberg, & Edwards, 2002), other research has studied alternative representations. Woodruff and colleagues (2001; 2002) created enhanced thumbnails by increasing the size of key text snippets within thumbnails to make it more readable. They found that the enhanced text thumbnails improved both finding and re-finding tasks, making them more applicable for participants who were not dependent on recognition. Others have tried using representative images from web pages, or portions of the webpage (Ayers & Stasko, 1995) instead of thumbnails, which has shown to improve performance in normal search tasks. Teevan and colleagues (2009), however, noted that the combination of salient images and key text, called Visual Snippets, can also provide a strong balance of support for both finding and re-finding, while text alone provided best support for finding, and thumbnails for re-finding.

One of the key advantages of thumbnails, however, is that users in re-finding scenarios can use visual perception to recognize websites they have seen before. Visual perception is faster than semantic interpretation of text (discussed by Woodruff and colleagues (2001)). Zheng and colleagues (2009), for example, demonstrated that users make decisions using just the visual perception of pages with 150ms, of the quality and professionalism of web page design.

The similar principle of enabling rapid visual recognition of web sites has been applied in most current web browsers. As users open a new tab, browsers like Safari and Firefox can display thumbnails that represent a combination of the most-used bookmarks and other frequently or recently visited pages. Pages that are visited frequently in this way are often being monitored (Kellar, Watters, & Shepherd, 2006), which is a slightly different task to re-finding on the web.

THE REVISIT RACK

The idea being embodied by the Revisit Rack, shown in **Figure 1**, is to make optimal use of thumbnails. Thumbnails have been shown to be most valuable during re-finding tasks, so the Revisit Rack displays thumbnails when queries are re-issued. Users are able to recognize websites quickly, using visual features such as color and layout, and so the Revisit Rack places the thumbnails all together. The Revisit Rack is placed at the top of the page, above the results, so that users do not have to scroll to recognize. Research has shown that the majority of eye fixations occur in this space within the first few seconds of a page loading (Buscher, Cutrell, & Morris, 2009). Thumbnails have been shown to be most recognizable when at approximately 200 pixel wide (Kaasten, et al., 2002), and so this size is used in the Revisit Rack. Finally, grouping thumbnails together means that they can be more easily compared, which could otherwise involve scrolling in traditional thumbnail-beside-result interfaces.

The design of the Revisit Rack, like with previous thumbnail-based designs (Kaasten, et al., 2002), is not without design tradeoffs. The size of the thumbnails means that fewer results are visible 'above the fold' (visible without scrolling). Optimally sized thumbnails, however, are often taller than the approximately 80 pixels often used for textual representations. Consequently, the use of the thumbnails, regardless of position, reduces the number of results that are visible above the fold. Further, the size of thumbnails meant that 8 results, in two lines of 4, was a more optimal number to display per page.

Despite these design tradeoffs, our hypothesis was that the increased value of enabling rapid recognition during re-finding tasks would make the Revisit Rack a valuable idea.

Prototype Implementation

Our prototype implementation uses the Yahoo Boss API[2] to return results for each query. Results are paginated with 8 per page, but were otherwise left unaffected. Thumbnails, generated from a relative fast API[3] are then displayed, as working links, in the Revisit Rack above the results. Cookies and session variables are used to track users.

USER STUDY

To learn more about the value of the Revisit Rack, a user study was created that involved both finding and re-finding tasks. Three user interfaces were included in the study: 1) a baseline generic web search interface, 2) a traditional thumbnail-beside-result experimental condition, and 3) the Revisit Rack experimental condition. During the study, the Revisit Rack was included in both finding and re-finding conditions, despite being designed for re-finding tasks.

Participants took part in two stages: first finding, and then 5 days later, re-finding. Participants performed two types of task on their first visit: A) a simple lookup fact-finding task, and B) a more exploratory subjective-choice task. In both cases, participants were provided with a motivating scenario and specific aim. An example fact-finding task was to find the population of Brazil. An example subjective task was to choose a piece of jewelry for a family member. The use of both the scenarios and, for the second task, personal subjective choice, was designed to make the tasks less impersonal and more memorable. In the re-finding phase, participants were given two minutes to re-find the source page that provided their original answer. Partici-

[2] http://developer.yahoo.com/search/boss/
[3] http://www.shrinktheweb.com/

pants were not given a query history and were not reminded of their original answer. Participants were able to submit any query, including those used previously.

Procedure

Participants were first provided with their legal rights and provided with sufficient detail to give informed consent. Participants then filled out a demographic questionnaire before beginning to perform the tasks with each user interface (UI). Each participant performed all 6 tasks (one of each type for each UI); UI ordering was counter-balanced. After 5 days, participants returned and were asked to re-find the sources or chosen items for the experimenter, repeating the tasks on the 3 UIs in the same order. Participation was concluded with a debrief discussion of the three UIs. The first session lasted around 40 minutes, and second only 20 minutes. All tasks were timed, and all queries and viewed results were logged.

Hypotheses

We had 3 motivating hypotheses for our pilot study:

1) In line with previous studies, we expected to see no significant variation between the UI conditions during finding tasks.
2) For re-finding we expected that users would more quickly re-find information online with both thumbnail conditions, and specifically with the Result Rack.
3) We expected that the improved times for the thumbnail conditions would hold for both fact-finding and subjective exploratory tasks.

Participants

12 participants were recruited, aged between 18-60 (mean: 30), with mixed educational backgrounds; 8 were undergraduates, 4 in computer science. 9 participants were male, and 3 female. All participants indicated that they searched the web at least daily, re-visited pages frequently, and 8 participants said they had used thumbnails in web searches.

RESULTS

Unfortunately, within this pilot study, we were so far unable to provide any consistent evidence for these hypotheses, but instead gained several insights into the factors that affected our mixed results. **Figure 2** shows the average times of the finding and re-finding tasks for both factual and subjective tasks in the 3 UI conditions. Failed and incorrect, as well as successful re-finds, are listed in **Table 1**; failed re-finds were counted as 2 minutes in **Figure 2**.

Figure 2 shows that the standard thumbnails-by-text condition was the fastest in all cases. With the Revisit Rack actually performing the slowest for fact re-finding. One difference, however, was actually statistically significant within this small pilot study, and that was the improvement that both thumbnail conditions had over the control condition for the initial subjective task ($F(2,22)=4.37$, $p<0.05$).

Discussion

We sought to understand these mixed results in more detail, using both the results of short debriefing interviews with participants and detailed timings from the logs. We discovered that the pagination of results had a large effect on timing. The loading of thumbnails in both experimental conditions, of course, also incurred longer loading times than the control condition. Consequently, if the result they chose in the finding phase was not in the first page of results, participants may have taken significantly longer to re-find results when they returned. It seems, therefore, that the Revisit Rack was fast and efficient *if* the result was easily recognizable on that page of results. We saw several short re-finding times of less than 10 seconds, especially in the Revisit Rack condition. However, if the result *was not* clearly recognizable in the Revisit Rack, then it took longer for people to decide that the target was not present and move on. From discussions, it appears that the spatial disconnect between thumbnails and results slowed participants in deciding to move onto the next page. Consequently, we also saw several longer re-finding task times in the Revisit Rack condition. We plan to first study the effect of this spatial disconnect to see if it does indeed take longer for participants to move on from a page of results during re-finding. Conceivably, this disconnect could also be mitigated by dual highlighting activated by hovering over either the result or the thumbnail. Dual highlighting, however, does not mitigate the potentially need to scroll to correlate thumbnails and results.

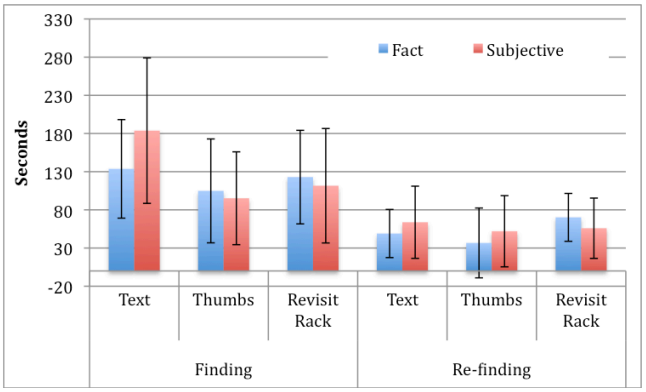

Figure 2: Finding and re-finding times for both fact and subjective retrieval for each condition. Error bars show Standard Deviations.

		Success	Incorrect	Failed
Fact	**Text**	9	3	0
	Thumbs	9	2	1
	Revisit Rack	8	1	3
Subjective	**Text**	6	4	2
	Thumbs	9	2	1
	Revisit Rack	10	0	2

Table 1: Rates of failed and in-correct re-finds

The potentially conditional benefit of the Revisit Rack does make some sense. The motivation of the Revisit Rack is to facilitate quick visual recognition, and so its real benefit over the traditional representation of thumbnails, is when the correct re-finding target *is* present in the current page of results. Consequently, another more macro-focused study

could measure the time taken to re-find a result that is definitely contained in a set of results, considering the position of the target in the set of results (similar to (Teevan, et al., 2007)). Combined with studying the layout-disconnect discussed above, this second study could tell us more concretely if and when the Revisit Rack provides value. This pilot study has provided, therefore, two more purposeful methods for studying the Revisit Rack in the future.

Finally, aside from discussing the potential intricacies of the Revisit Rack, our pilot study did also provide some other more general insights. Most notably, we saw an almost consistent 25% failure rate for re-finding within a given 2 minute period, after only 5 days from the original search time. This highlights the fact, as proposed in many prior publications (Cockburn & McKenzie, 2001; Tauscher & Greenberg, 1997; Teevan, et al., 2007), that re-finding should not be taken lightly. We also saw a more notable failure rate for subjective task re-finding without thumbnails. The subjective choices involved, for example, jewelry for a family birthday. Further, we saw that it took people significantly (our only significant result) longer to make initial subjective decisions without thumbnails. Participants may have been making quality-of-website judgments as per the evaluation of visual-perception judgments by Buscher and colleagues (2009). Although, given the size of the study, these variations in the subjective tasks may have been by chance, the Exploratory Search (R. White & Roth, 2009) community may wish to further investigate the role of thumbnails in decision making.

CONCLUSION

In this paper, we have introduced the idea of a Revisit Rack for re-finding on the web, which presents thumbnails of results together at the top of a results page. Thumbnails leverage visual recognition, rather than semantic interpretation of text, and so placing them together could optimize both recognition and comparison during re-finding. We compared the Revisit Rack to a traditional thumbnail and text-based result review in a pilot study, which provided mixed results that did not support our hypotheses. The investigation, however, did highlight a) that there is a potentially significant cost of presenting thumbnails separately from the results, b) that the Revisit Rack may have more subtle benefits in certain conditions, and c) how the Revisit Rack may therefore be studied more formally in the future.

REFERENCES

Adar, E., Teevan, J., & Dumais, S. T. (2009). Resonance on the web: web dynamics and revisitation patterns. In *Proceedings of the 27th international conference on Human factors in computing systems (CHI'09)* (pp. 1381-1390). Boston, MA, USA: ACM.

Ayers, E., & Stasko, J. (1995). Using graphic history in browsing the World Wide Web. In *Proceedings of the 4th International World Wide Web Conference* (WWW'95).

Buscher, G., Cutrell, E., & Morris, M. R. (2009). What do you see when you're surfing?: using eye tracking to predict salient regions of web pages. In *Proceedings of the 27th international conference on Human factors in computing systems (CHI'09)* (pp. 21-30). Boston, MA, USA: ACM.

Cockburn, A., & McKenzie, B. (2001). What do Web users do? An empirical analysis of Web use. *International Journal of Human-Computer Studies, 54*(6), 903-922.

Kaasten, S., Greenberg, S., & Edwards, C. (2002). How people recognize previously seen Web pages from titles, URLs and thumbnails. In *Proceedings of HCI'02*.

Kellar, M., Watters, C., & Shepherd, M. (2006). A goal-based classification of Web information tasks. In *Proceedings of the 2006 Annual Meeting of the American Society for Information Science and Technology (ASIST'06*) (pp. 1-22). Austin, TX, USA.

Tauscher, L., & Greenberg, S. (1997). How people revisit web pages: empirical findings and implications for the design of history systems. *International Journal of Human-Computer Studies, 47*(1), 97-137.

Teevan, J., Adar, E., Jones, R., & Potts, M. A. S. (2007). Information re-retrieval: repeat queries in Yahoo's logs. In *Proceedings of the 30th international conference on Research and development in information retrieval (SIGIR'07)* (pp. 151-158) Amsterdam: ACM.

Teevan, J., Cutrell, E., Fisher, D., Drucker, S. M., Ramos, G., Andre, P., et al. (2009). Visual snippets: summarizing web pages for search and revisitation. In *Proceedings of the 27th international conference on Human factors in computing systems (CHI'09)* (pp. 2023-2032). Boston, MA, USA: ACM.

White, R., & Roth, R. (2009). *Exploratory Search: Beyond the Query-Response Paradigm* (Vol. 1): Morgan & Claypool.

White, R. W., Ruthven, I., & Jose, J. M. (2002). Finding relevant documents using top ranking sentences: an evaluation of two alternative schemes. In *Proceedings of the 25th international conference on Research and development in information retrieval (SIGIR'02)* (pp.57-64). Tempere, Finland: ACM.

Woodruff, A., Faulring, A., Rosenholtz, R., Morrsion, J., & Pirolli, P. (2001). Using thumbnails to search the Web. In *Proceedings of the international conference on Human factors in computing systems (CHI'01)* (pp. 198-205). Seattle, WA, USA: ACM.

Woodruff, A., Rosenholtz, R., Morrison, J. B., Faulring, A., & Pirolli, P. (2002). A comparison of the use of text summaries, plain thumbnails, and enhanced thumbnails for Web search tasks. *Journal of the American Society for Information Science and Technology, 53*(2), 172-185.

Zheng, X. S., Chakraborty, I., Lin, J. J.-W., & Rauschenberger, R. (2009). Correlating low-level image statistics with users - rapid aesthetic and affective judgments of web pages. In *Proceedings of the 27th international conference on Human factors in computing systems (CHI'09)* (pp. 1-10). Boston, MA, USA: ACM.