

Evaluating the Effectiveness of Visual Summaries for Web Search

Hilal Al Maqbali Falk Scholer James A. Thom Mingfang Wu
School of Computer Science and Information Technology
RMIT University
GPO Box 2476, Melbourne 3001
Victoria, Australia

{hilal.almaqbali,falk.scholer,james.thom,mingfang.wu}@rmit.edu.au

Abstract

With ever-increasing amounts of information on the World Wide Web, an effective interface for displaying search results is required. Recent studies have developed various novel approaches for visual summaries, aiming to improve the effectiveness of search results. In this study we evaluate the effectiveness of four types of visual summary: thumbnails, salient images, visual snippets and visual tags. Fifty participants carried out five informational topics using five different interfaces. The results show that visual summaries significantly impact on the behavior of users, but not on their performance when predicting the relevance of answer resources. Users spend significantly less time looking at the textual components of summaries with the visual summary interfaces. Comparing the performance of users in predicting the relevance of answer pages with a text interface versus visual interfaces suggests that the tested visual summaries can mislead users to select non relevant items on informational search topics.

Keywords Information Retrieval, User Studies Involving Documents, Web Documents, Visual Summaries, Eye Tracking.

1 Introduction

The amount of information on the World Wide Web has been increasing exponentially; and search engines are the key tool for supporting users in finding information. Search results presentation and organisation are important components that impact on the overall search effectiveness [1, 2]. Existing search engines not only show textual summaries (such as a page title, a short textual snippet, and a URL), but also provide visual features. One of the most common types of feature is the visual summary, such as a thumbnail or a dominant image [14].

Although popular search engines such as Google historically focused more on improving textual summaries for each result page, they have started showing visual summaries for some of the top search results. Other search engines, such as Middlespot, Nex-

Proceedings of the 15th Australasian Document Computing Symposium, Melbourne, Australia, 10 December 2010. Copyright for this article remains with the authors.

plore and Viewzi, display a visual summary for every result in their answer list [1].

It has been said that one image is worth a thousand words. Humans can digest the meaning of an image quicker than text; in the time that a user spends to understand the gist of an image, a user can read only one to four words [7]. Visual summaries have been shown to significantly help users in the judgment of web search results in many studies [9, 14, 15, 19, 21, 22]. Visual summaries are helpful for refinding previously visited web pages and can also provide hints for users when the search tasks are confusing or when users are not proficient at them.

In this study, we investigate the impact of different approaches of visual summaries for informational tasks on user behavior and analyse the time spent looking at each component of the search results presentation (title, text snippet, URL and visual summary). We compare visual snippets, visual tags, excerpt images, and thumbnails. Each one of these visual summaries is presented on an interface together with a text summary. We investigate the following questions:

1. Does providing additional visual summaries with search results improve the ability of users to predict the relevance of search results?
2. How do visual summaries impact on user behavior, particularly on text summaries when additional visual summaries are presented?
3. How do users interact with different components of a results screen, for those search results that include visual summaries?
4. How does the presence of visual summaries affect task completion time?

Our analysis shows that users spend significantly less time looking at textual summaries when visual summaries were available. However, overall, the results suggest that visual summaries do little to increase user performance with informational topics.

This paper is organised as follows: in Section 2, related work is reviewed. In Section 3, we describe our experiment on design including the visual summaries, users and topics. Experimental results are analysed in

Section 4, and discussion and conclusion are presented in Section 5.

2 Related work

Dziadosz [9] described the interaction between user and information retrieval systems, and summarized these into three steps: query formulation, relevance prediction, and relevance evaluation. Visual summaries can help to improve the relevance prediction as shown in many studies [9, 14, 15, 19, 21, 22].

Several novel approaches have been developed for visual summaries to improve user performance in finding desired information. Some of these approaches use the snapshot of a web page, such as a thumbnail or an enhanced thumbnail [21], whereas others use a salient image within the retrieved web page such as a salient image [14] or visual snippet [19].

A **thumbnail** is a miniature image of a web page. This is the most common type of visual summary and has been evaluated in many studies [3, 6, 8–10, 12, 16, 20–22]. Thumbnails help users to recognise the layout of the retrieved web page. They not only make it easy to recognise the web page if it has been seen before, but can also provide users with relevant visual hints for their queries in the form of a picture, table or website logo.

Woodruff et al. [21] develop a novel approach for a visual summary called **an enhanced thumbnail**, which highlights and enlarges the query terms within the thumbnail images. They compared enhanced thumbnails with plain thumbnails and with text summaries. Four different search tasks were used where users were asked to find a picture, homepage, e-commerce website, or side-effects of a given drug (informational query). The results show that visual summaries reduced the number of visited pages to find the answer for a given search task. Also, they found no significant difference on performance and time spent to answer the informational queries.

Teevan et al. [19] developed a **visual snippet** which combines the page title, a salient image and a logo of the website. Twelve topics were used in the study, four for each type (homepage, shopping, medical information). Analysis shows that visual snippets are more useful than plain thumbnails, particularly for revisitation (refinding a previously seen web page). The study also discussed the possibility of generating visual snippets automatically.

Li et al. [14] examined the effectiveness of presenting **excerpt (salient) images** with search results, by examining two interfaces, one with text summaries only and the other with both text summaries and excerpt images. Two types of queries were used in the study, informational and navigational. The results showed that excerpt images are helpful, and can be generated for almost all query types. For example, according to the experimental results, excerpt images decreased the time spent by the user on informational queries by 30.4%



Figure 1: Salient image interface: (A) Text summary region. (B) Visual summary region.

compared with the time spent on text summaries without excerpt images.

A study by Jiao et al. [11] was conducted to evaluate four types of visual summaries: internal image (the dominant image in a web page); external image (a representative image from an external page); visual snippets; and, thumbnails. The study evaluated these visual summaries in two phases. In the first phase, participants were given the visual summaries and asked to type descriptions about the expected content of the related web page. Several hours later, in the second phase, researchers investigated how the visual summaries affected the recall of the web pages visited in phase 1. The results show different types of visual summary work better on particular types of web pages and search tasks.

Tag clouds have become a popular method for visualising information on the web. Tag clouds visualise the most frequently used words in a document by showing the relative importance of terms using different font sizes, weight or color. Many studies [4, 17, 18] show that tag clouds can provide effective cues about the content of text search results for users. We developed a novel approach (called visual tags) for the visual summary which combines the tag clouds of a document and its snapshot. Our hypothesis is that these combined visual tags will provide useful hints for users about the content of the retrieved web page.

3 Experimental methodology

In order to evaluate the effectiveness of different approaches for visual summaries, and study the impact of these visual summaries on user seeking behavior and performance, we conducted a user study that involved a series of five informational search topics using different search interfaces where visual summaries are a primary component of the search results presentation.

3.1 Experimental setup

In our user study, the participants were mostly undergraduate and high school students with some interest in computer science visiting RMIT University at the 2010 Open Day. A plain language statement was given to the subjects to outline the purpose of the experiment, the procedure, the tasks to be performed, and the data to be collected. Based on this information, 65 participants chose to take part in the experiment. However, due to interruptions and difficulty with calibrating the eye tracking for some volunteers (we eliminated users with less than 80% capture accuracy), the collected data of only 50 participants is included in the analysis. A short oral presentation about the visual summaries was given to each participant, but no training was given on the interfaces to be used.

Each participant was asked to evaluate items in search result lists for five informational search topics, each with a different interface. For each topic, five answer items were shown on a single page. Participants used the mouse to select all items that they considered to be relevant to the given topic. Participants were not able to browse the actual web pages embedded in the text search results, relying solely on the search result page given. Note that users were presented with a fixed search results list for the topic, and did not engage in interactive searching.

Data was collected using a Tobii T60 eye tracker. This non-invasive device calculates the exact point of a user's gaze using a geometrical model. Since all the search results were presented on a single screen page, participants did not spend extra time or visual attention having to scroll the search results page.

3.2 Interfaces

Five interfaces were designed for this experiment, each presenting exactly the same text summary, but with different visual browse features (visual summaries), except for the text interface which presents only text summaries. The interfaces present for each item: document title; a text snippet, that is a short text extract from the source document that closely relates the query terms; the URL; and (apart from text interface) a visual summary component.

In order to control the design of the interfaces, a template was created to enable data for the five interfaces to be consistently and uniformly added. On the template, the text summaries are shown on the right-hand side at the screen, and pictures are displayed at a maximum of 200x150 pixels using original ratio on the left side. For each particular topic, the same textual surrogates (titles, snippets and URLs) are shown in exactly the same place and using the same format on all of the interfaces. The visual summaries are also displayed in exactly the same place on each interface, except for the text-only interface where the visual summaries are replaced by white space. Figure 1 shows an example of the salient image interface for one of the topics involved

in this study. So, for all the interface that present the same topic, text summary region (A) was the same, and only the visual summary region (B) was changed.

Apart from the text-only interface, the remaining interfaces provide a specific type of visual summary on their search results presentation for each result page. The four visual summaries are: thumbnails, visual tags, visual snippets, and salient images. Some of these visual summaries have been evaluated by other researchers.

Thumbnail: A thumbnail is a miniature snapshot of a web page. An example is shown in Figure 2(a). A software tool called WebShot¹ was used to generate the screenshots for the test collection, to ensure the same properties for all the thumbnails.

Visual tags: The second visual summary is our approach, the visual tags summary, which is a combination of a thumbnail with a tag cloud of the retrieved web page. A tag cloud presents the most frequently used words in a document and shows the relative importance of terms using different font sizes. Our hypothesis is that this combined visual tags will provide an effective cue to the content of the retrieved web page. The construction of the visual tags includes two main stages. Firstly, a transparent image of each tag cloud was created using Wordle website². The next step was to combine this with the thumbnail of the related web page. Buscher et al. [5] have found that people focus more on the top left corner of a web page, because the logo and the main navigation bar are usually located in that area, so to preserve this information region the tag cloud was located on the right of the thumbnail as shown in Figure 2(b).

Visual snippet: Visual snippets were proposed by Teevan et al. [19], and consist of a logo, a salient image and the page title. In this experiment, the visual snippet is the integration of the salient image from the retrieved web page and the website logo as shown in Figure 2(c). The page title was not included in our visual snippet because it is already presented in the related textual surrogates. Salient images were collected using Google image search over the target URL, and selecting the top-ranked image.

Salient image: The fourth type of visual summary is a salient image extracted from the underlying web page, see Figure 2(d). The salient image is extracted using Google image search, as explained for visual snippets.

3.3 Topic selection

Web search tasks can be classified as informational, transactional or navigational [5]. In this study, we focus on informational search tasks which aim to find specific information for a given topic. Five informational topics on general knowledge were developed :

¹www.websitescreenshots.com

²www.wordle.net



Figure 2: Examples of the form types of visual summaries.

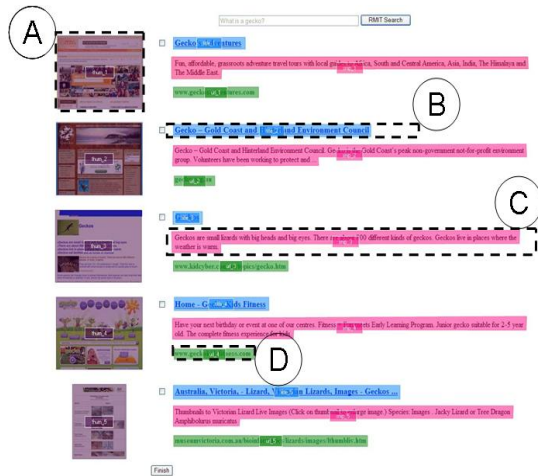


Figure 3: The mask used to collect time spent on the specific informative components (A) Exact visual summary. (B) Page title. (C) Text snippet. (D) URL.

1. What are the side effects of energy drinks?
2. What is a gecko?
3. What is an appropriate sitting posture at a computer?
4. What is a solar eclipse?
5. What is a Vuvuzela?

To obtain realistic search engine results, the top ten search results from the Bing search engine were selected for each query. Wikipedia entries were excluded as they have obvious answers for the experimental tasks, then five items were randomly chosen from the remaining search results. To ensure balance in the result sets, the quality of the five selected items was restricted to include at least one relevant and one non-relevant answer, as judged by the authors.

3.4 Size of the visual summaries

Based on studies by Kaasten et al. [13] and Won et al. [20], the visual summaries in our study were set to a size of 200x150 pixels. A pilot test showed this thumbnail size to be appropriate. Search engines such as Google and Yahoo usually present ten items per page for query results, so presenting visual summaries at 200x150 pixels would not require more space than the size of a standard search result page.

3.5 Experiment design

After reading a topic from the screen, the participant clicked a “Start” button to load the search interface. Five items were displayed as search results, and participants were asked to select all items that they consider to be a relevant answer for the task. Then, the participant clicked on a “Finish” button to move to the next task.

The presentation of topics and interfaces were determined by a Latin square, giving 25 combinations of interfaces and topics, to control for presentation order effects.

3.6 Measures

User behavior is analysed by calculating the average time that users spent on specific informative components. A mask was built in the eye tracking data to collect these areas of attention, as shown in Figure 3.

The effectiveness of visual summaries was measured by click precision, click recall, and click F-measure. Click precision measures the correctly identified relevant answers as a proportion of all answers that the users selected, while click recall shows the number of relevant answer selected by users as a proportion of the total number of relevant answer available for that topic. The click F-measure gives the harmonic mean between click precision and click recall. Also, scripts were built into the HTML documents to record interaction events: the time it takes to identify a relevant item, the number of click events, and the total time taken to finish the task. In addition, t-test (p) and Chi squared test (χ^2) are applied to find the statistical difference between interfaces on each measure.

4 Results

We analyse user behavior when carrying out the five informational search topics, using a different interface for each, based on topic completion time and the relative attention paid to different summary features (page title, textual snippet, URL and visual summary).

4.1 Effectiveness of relevance prediction

In the user study, participants were asked to select all answer items that looked relevant for the given search topic. Table 2 shows the click precision, click recall and click F-measure for how effectively the users were able to identify relevant answers. We treat the text-only

Answer	Text	Thum	Tag	Image	Visual
Relevant	71	71	64	64	72
Non-relevant	18	26	34	21	27

Table 1: Distribution of the number of relevant and non-relevant answers selected by users, grouped by interface.

interface as a baseline, and compare the performance of each form of visual summary system against this using a t-test. Click precision shows no significant difference between the visual summary interfaces compared to the text interface, except for the visual tags interface which is significantly worse ($p = 0.005$). These observations indicate that the tag cloud can mislead users, causing non-relevant items to appear as potentially useful. Although the salient image interface achieves the highest average click recall (0.6), a statistical test shows no significant difference to the text-only interface ($p = 0.753$).

The number of relevant and non-relevant items that users selected are shown in Table 1, split by the interface used. The results show largely consistent rates of success with no significant differences between the interfaces ($\chi^2, p > 0.1$).

4.2 Interaction with textual summaries

User interaction with the search results was captured using eye tracking data. By using this dataset, we can investigate how users interact with textual summaries when additional visual summaries are presented. The amount of time spent looking at the overall text and visual summary regions is shown in Table 3. Users spent substantially more time looking at the text region for all interfaces. While users in general spent less time looking at the text region when using a visual summary interface, this difference was only significant for the interface that presented salient images ($p = 0.032$).

At the component level, user attention to specific informative components was evaluated by collecting the amount of time that the user’s gaze rested on each component. The gaze regions were closely bounded on the interface component, leaving regions of white-space between them, as shown in Figure 3.

Table 3 shows the proportion of viewing time that users spent looking at the specific informative components (title, textual snippet, URL, and visual summary). Using the text interface as a baseline, there is no significant difference to the amount of time spent looking at titles or URLs when presented with a visual summary. However, on all the interfaces, users spent more time looking at page title components compared to URLs. Interestingly, the amount of time spent viewing textual snippets was significantly less when any visual summaries were displayed (for all the visual interfaces $p < 0.05$). That is, presenting the visual summary feature decreased the attention that users gave to the text-based summary information.

Figure 4 shows the percentage of time spent on the four specific informative components. On all the interfaces, users spent more time looking at page title and textual snippets than visual summaries and URLs. Comparing the four visual summaries with each other indicates that users spent more time looking at visual tags than other visual summaries, presumably because they spent time on reading the text of the visual tags.

We also analysed how users scanned the search results with each interface by collecting the time that users spent looking at each informative component for each search result item. The results show that users were more influenced by the vertical list of the search results when they used the text interface, but this behavior was less apparent on the interfaces that present visual summaries as shown in Figure 5. This supports our observation that users spend significantly less time looking at textual snippets when using the visual interfaces. This behavior appears to be one of the main reasons that users are misled when trying to identify relevant answers for the informational tasks.

4.3 Interaction with the visual search interfaces

Next we study user attention in relation to the different informative components of the result screen. A broader comparison was conducted, by pooling the data for the four interfaces that include visual summaries, and then comparing the four attention areas (informative components) for those. The results show that users spent significantly more time looking at the textual snippets than the other informative components ($p < 0.001$). However, there is no statistically significant difference between time spent on page titles and visual summaries, based on the aggregated attention areas.

4.4 Overall task completion time

The performance of users to complete a task was evaluated by collecting: the time users spent on each task; the time taken to first selection; and, the time taken to select first relevant item. However, no statistically significant differences were found between the five interfaces. Table 4 shows the average time spent to answer the search tasks for each interface. Although users required the least amount of time to finish their search tasks with the salient image interface, there was no statistically significant improvement compared to using the text-only interface ($p = 0.503$). Also, we calculated the average time that a user spent to answer each search task for each interface. However, no significant difference was observed between the interfaces on the measures of time completion.

5 Discussion and Conclusion

In this study, we evaluated the impact of different types of visual summaries on user behavior and performance. Fifty participants carried out five informational topics

Measures		Text	Thum	Tag	Visual	Image
Click Precision	Average	0.865	0.794	0.689	0.800	0.820
	Stdev	0.237	0.270	0.365	0.322	0.298
	t-test		0.168	0.005	0.253	0.412
Click Recall	Average	0.582	0.600	0.505	0.592	0.535
	Stdev	0.263	0.316	0.320	0.335	0.303
	t-test		0.753	0.193	0.868	0.413
Click F-measure	Average	0.645	0.625	0.545	0.624	0.569
	Stdev	0.190	0.232	0.290	0.271	0.250
	t-test		0.639	0.043	0.653	0.267

Table 2: Click precision, click recall and click F-measure for user selection of search result items.

Measures		Text	Thum	Tag	Visual	Image
Text summary region	Average	23.853	19.684	19.205	19.150	16.923
	Stdev	16.504	14.161	13.976	14.424	15.318
	t-test		0.178	0.132	0.132	0.032
Visual summary region	Average	0.113	3.742	5.771	4.436	3.919
	Stdev	0.244	4.038	6.068	4.650	4.294
	t-test		<0.001	<0.001	<0.001	<0.001
Titles	Average	4.263	3.514	3.846	3.791	3.615
	Stdev	2.176	2.369	3.032	3.292	2.944
	t-test		0.103	0.432	0.400	0.214
Textual snippets	Average	9.512	6.245	5.866	5.931	5.777
	Stdev	8.585	5.313	6.250	5.692	7.516
	t-test		0.024	0.017	0.016	0.023
URLs	Average	0.943	0.700	0.632	0.676	0.606
	Stdev	1.493	0.898	0.880	0.887	0.797
	t-test		0.327	0.209	0.280	0.163
Exact visual summaries	Average	-	0.659	1.136	0.744	0.645
	Stdev	-	0.754	1.186	0.801	0.728

Table 3: The Average, Standard deviation and *t*-test for the time in seconds spent viewing summary regions and informative components. The *t*-test evaluates the difference between the text interface and each of the interfaces that includes a visual component.




Figure 4: The percentage of time spent on the specific informative components.

Measures		Text	Thum	Tag	Visual	Image
Time to finish	Average	27.507	26.839	29.312	27.629	25.117
	Stdev	17.542	17.236	17.537	18.394	18.021
	t-test	N/A	0.848	0.608	0.973	0.503

Table 4: The Average time spent to finish search tasks for each interface. The t-test compares the text-only interface with each of the visual interfaces.




Figure 5: Average time spent on the textual surrogates for the five search result items.

using five different interfaces. Our study primarily focused on evaluating the ability of users to predict the relevance of answers when visual summaries are provided. Other studies [11, 19] focus on evaluating visual summaries in terms of finding and refinding issues, whereas in our study we considered informational search tasks.

Providing additional visual summary information with the text search results did not significantly improve the ability of users to predict the relevance of a result page for an informational search task. Although the salient image interface achieved the highest average click recall, a statistical test showed no significant difference compared to the text-only interface baseline. Also, no significant difference was found for the number of relevant result pages that users selected for each interface. Further, the results show that adding visual summaries may mislead users to select non-relevant results pages for the search topics. A possible reason for explaining this behavior is that users are not as familiar with these novel approaches as with standard text-only result lists.

We studied user behavior when presented with an additional visual summary, and the results show that visual summaries significantly affect user behavior. Although the informational search tasks seem to require reading text more than looking at pictures, users in general spent less time looking at the text region when us-

ing a visual summary interface. This may also explain the lower performance when predicting the relevance of answers items when visual information is displayed.

Furthermore, we analysed how user attention is devoted to specific informative interface components. Users spend more time looking at tag clouds, but there is no significant difference in attention between the four interfaces that include visual information. Also, the results show that users scan the search results exhaustively for the text interface, but economically for the visual interfaces. With the text interface, users spent more time looking at the top items and this amount gradually decreased as they move down the ranked list, while for the visual interfaces, the amount of time per item shows less variation.

In addition, we collected the time users spent on each task, time taken to first selection and time taken to select first relevant item. However, given our sample size no statistically significant differences were found. This suggests that visual summaries do not provide enough information for informational search tasks, since the answers for this type of search are more likely to be located in the text rather than visual summaries.

In future work, we plan to use the collected data to improve the use of an eye tracker for the evaluation of web search interfaces. Also, we plan to conduct further user studies over a wider range of tasks.

References

- [1] Hilal Al Maqbali, Falk Scholer, James A. Thom and Mingfang Wu. Do users find looking at text more useful than visual representations? a comparison of three search result interfaces. In *The Fourteenth Australasian Document Computing Symposium (ADCS)*, pages 35–42, Australia, 2009.
- [2] Hilal Ali[Al Maqbali], Falk Scholer, James A. Thom and Mingfang Wu. User interaction with novel web search interfaces. In *OZCHI '09: Proceedings of the 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group*, pages 301–304. ACM, 2009.
- [3] Anne Aula, Rehan M. Khan, Zhiwei Guan, Paul Fontes and Peter Hong. A comparison of visual and textual page previews in judging the helpfulness of web pages. In *WWW '10: Proceedings of the 19th international conference on World wide web*, pages 51–60. ACM, 2010.
- [4] Scott Bateman, Carl Gutwin and Miguel Nacenta. Seeing things in the clouds: the effect of visual features on tag cloud selections. In *HT '08: Proceedings of the Nineteenth ACM conference on Hypertext and hypermedia*, pages 193–202. ACM, 2008.
- [5] Andrei Broder. A taxonomy of web search. *SIGIR Forum*, Volume 36, Number 2, pages 3–10, 2002.
- [6] Andy Cockburn, Carl Gutwin and Jason Alexander. Faster document navigation with space-filling thumbnails. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 1–10. ACM, 2006.
- [7] Veronika Coltheart. *Fleeting Memories: Cognition of Brief Visual Stimuli*. MIT Press: Cambridge, 1999.
- [8] Anna Divoli, Michael A. Wooldridge and Marti A. Hearst. Full text and figure display improves bioscience literature search. *PLoS ONE*, Volume 5, Number 4, pages e9619, 04 2010.
- [9] Susan Dziadosz and Raman Chandrasekar. Do thumbnail previews help users make better relevance decisions about web search results? In *SIGIR '02: Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 365–366. ACM, 2002.
- [10] May Eric, Eric Z. Ayers and John T. Stasko. Using graphic history in browsing the world wide web. In *The International World Wide Web Conference*, pages 11–14, 1995.
- [11] Binxing Jiao, Linjun Yang, Jizheng Xu and Feng Wu. Visual summarization of web pages. In *SIGIR '10: Proceeding of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 499–506. ACM, 2010.
- [12] Hideo Joho and Joemon Jose. A comparative study of the effectiveness of search result presentation on the web. In *Advances in Information Retrieval*, Volume 3936 of *Lecture Notes in Computer Science*, pages 302–313. Springer Berlin / Heidelberg, 2006.
- [13] Shaun Kaasten, Saul Greenberg and Christopher Edwards. How people recognize previously seen web pages from titles, urls and thumbnails. In *People and Computers XVI (Proceedings of Human Computer Interaction)*, pages 247–265, 2001.
- [14] Zhiwei Li, Shuming Shi and Lei Zhang. Improving relevance judgment of web search results with image excerpts. In *WWW '08: Proceeding of the 17th international conference on World Wide Web*, pages 21–30. ACM, 2008.
- [15] Maarten, Mary P. Czerwinski, Maarten Van Dantzich, George Robertson and Hunter Hoffman. The contribution of thumbnail image, mouse-over text and spatial location memory to web page retrieval in 3d. In *Human-Computer Interaction – INTERACT'99*, pages 163–170. IFIP, 1999.
- [16] William Ogden, Mark Davis and Sean Rice. Document thumbnail visualizations for rapid relevance judgments: When do they pay off? In *The Seventh Text REtrieval Conference (TREC7)*. NIST, pages 528–534, 1998.
- [17] Johann Schrammel, Michael Leitner and Manfred Tscheligi. Semantically structured tag clouds: an empirical evaluation of clustered presentation approaches. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 2037–2040. ACM, 2009.
- [18] James Sinclair and Michael Cardew-Hall. The folksonomy tag cloud: when is it useful? *J. Inf. Sci.*, Volume 34, Number 1, pages 15–29, 2008.
- [19] Jaime Teevan, Edward Cutrell, Danyel Fisher, Steven M. Drucker, Gonzalo Ramos, Paul André and Chang Hu. Visual snippets: summarizing web pages for search and revisitation. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 2023–2032. ACM, 2009.
- [20] Sungjoon Steve Won, Jing Jin and Jason I. Hong. Contextual web history: using visual and contextual cues to improve web browser history. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 1457–1466. ACM, 2009.
- [21] Allison Woodruff, Andrew Faulring, Ruth Rosenholtz, Julie Morrision and Peter Pirulli. Using thumbnails to search the web. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 198–205. ACM, 2001.
- [22] Songhua Xu, Tao Jin and Francis C. M. Lau. A new visual search interface for web browsing. In *WSDM '09: Proceedings of the Second ACM International Conference on Web Search and Data Mining*, pages 152–161. ACM, 2009.