

TheHotMap.com: Enabling Flexible Interaction in Next-Generation Web Search Interfaces

Orland Hoerber
Department of Computer Science
Memorial University of Newfoundland
St. John's, NL, Canada A1B 3X5
hoeber@cs.mun.ca

Michael Brooks, Daniel Schroeder, Xue Dong Yang
Department of Computer Science
University of Regina
Regina, SK, Canada S4S 0A2
{brooks1m, schroedd, yang}@cs.uregina.ca

Abstract

TheHotMap.com is an example of a next-generation Web search interface, wherein users are able to take an active role in the Web search process. Information that can assist searchers in their Web search tasks is presented in a visual manner. The system supports user interaction both during the query refinement process, as well as the search results exploration process. This paper describes the features of TheHotMap.com; usage log analysis illustrates the benefits of supporting flexible interaction within a single unified Web search interface.

1. Introduction

Although Web search has been an important part of the Web since its early days, very little advancement has been made in the interfaces used by the top Web search engines. Even though user interfaces have been an active research area of traditional information retrieval for many years [2, 5], few of these results have been applied to Web search. A query box continues to be the primary mechanism for capturing searcher intent, and a paged search results list continues to be the primary method for displaying the search results. With only a few exceptions, minimal support is provided to assist users in crafting queries and subsequently exploring the search results.

Our prior research has focused on enabling users to take an active role in their Web search processes, resulting in steps toward the next generation of Web search interfaces. The goal has been to allow searchers to interact with and affect the outcomes of their Web searches by making intelligent choices and selections as they specify and refine their queries, and as they evaluate and explore the search results.

In this paper, we present *TheHotMap.com*¹, a public

¹<http://www.TheHotMap.com/>

Web search interface developed as an extension of some of our previously published research prototypes for visual and interactive Web search [9, 8]. A screenshot of the system is shown in Figure 1. Lightweight interface extensions are added to the common list-based representation of the search results to support both interactive query refinement and interactive search results exploration in a single unified visual interface. The goal has been to add easy-to-use interface elements that allow users to visually interpret and make sense of the fundamental Web search information, and take an active role in the Web search process.

2. Related Work

This work is closely related to our previous research activities in the development of next-generation interfaces for Web search. In particular, the system is based on a combination of two of our previous research prototypes: HotMap [9] and Wordbars [8]. Both of these prototypes were originally developed with the purpose of exploring visual representations and the use of various types of information to support Web search activities. As research tools, the prototypes were useful for validating the potential utility of the proposed techniques [7, 11]. However, they were not designed for public use. *TheHotMap.com* is a complete reimplementation and extension of the methods described in these previous works.

Others have explored the use of visual interfaces to support the exploration of Web search results. Heimonen and Jhaveri [6] created an icon-based representation of the location of specific query terms within individual search results sets. Based on TileBars [4], this system allowed the searcher to see where in the resulting documents their search terms were being used together.

In VIEWER [3], the frequency of all combinations of the query terms were counted within the document surrogates. This information was presented in a histogram representa-




Figure 1. A screenshot of *TheHotMap.com*. Note the lightweight controls labeled in the figure.

tion. Selection within the histogram allowed the searcher to filter the search results set based on specific combinations of the query terms.

In terms of supporting users in crafting accurate Web search queries, very little work has been done. Joho et al. [13] developed a method for dynamically generating a hierarchy of potentially useful terms for query expansion, and presented these in a menu-like structure. We developed a method for visually representing a query space derived from an external knowledge base called VisiQ [12]. In addition, some of the top search engines and Web browsers have begun to address this issue by suggesting potential queries based on their usage logs. This feature appears in both Ask [1] and Yahoo [16], as well the Firefox [14] Web search bar.

One item of note is that we are unaware of any work beyond our own [10] that has attempted to provide an interactive interface for Web search, supporting both query refinement and search results exploration activities in a unified manner. *TheHotMap.com* is a novel interface that can allow searchers to readily switch back and forth between their tasks of crafting an accurate query and exploring the search results. This flexibility is one of the key benefits of the system, and is the focus of this paper.

3. TheHotMap.com

TheHotMap.com is implemented as a Web search interface layer overtop of the search results provided by the

Yahoo API [15]. There are three main features that enable flexible interaction within the system: the WordBars histogram, the re-sortable search results list using the HotMap query term headers, and the HotMap overview of the full search results set. These features were selected and refined based on the results of user studies conducted with the original HotMap and WordBars prototypes [7, 11].

The WordBars histogram provides a visual representation of the most frequently appearing terms within the search results set (using the title, snippet, and URL), allowing the relative frequency of these terms to be easily observed. Users can interactively re-sort the search results set by selecting the arrow icon beside any term of interest. Interactive query refinement is supported by clicking the plus icon beside any term users wishes to add to their queries, or the minus icon beside any term users wish to remove from the query.

In addition to the re-sorting supported via the WordBars histogram, searchers may also re-sort the search results based on the use of their specific query terms. Clicking on any of the query term headers above the search results list will cause the search results to be re-sorted based on the use of the selected term within the titles, snippets and URLs. Although the default sorting method is to perform single-term sorting, an advanced feature is available that supports nested sorting.

The HotMap overview provides a compact visual representation of the entire set of search results that are presented

in the list-based representation. In the current implementation, the system collects 50 search results per page. Colour coding is used to represent the frequency of the query terms within the search results overview. This same colour coding is also used in the search results list. The HotMap overview supports an interactive exploration of the search results. As users visually identify documents of interest, they may click on the abstract representation of the search result to cause the search results list to scroll to that location. The system temporarily highlights the corresponding search result that was selected in the HotMap overview, allowing users to easily relate their selection in the overview to the scrolled location in the search results list.

4. Flexible Interaction

The Web search interface of *TheHotMap.com* provides interactive support for different types of Web search activities. After submitting an initial query, the Web search results page will be loaded and dynamically populated with search results, appearing similar to Figure 1. This interface allows the searcher to perform tasks related to interactive query refinement and interactive search results exploration.

For example, suppose a user begins with an initial query of “interactive visual web search”. If the user is confident in the quality of their query, they may choose to explore the search results by re-sorting. This can occur in the WordBars histogram (e.g., selecting the arrow beside the term “interface”) or within the search results list (e.g., selecting “interactive” in the search results list header). Alternately, the user can inspect the HotMap overview and jump to a location of interest (e.g., identifying a search result that makes frequent use of “interactive” and “visual”, and jumping to that location to investigate the document further). Alternately, searchers may wish to use the WordBars histogram to add or remove terms from their query (e.g., adding “interface” and removing “interactive”), in order to specify more accurate descriptions of their information needs.

The benefit of this system is that users can easily switch back and forth between these Web search activities, since the controls that support them are always available and do not require the user to explicitly switch between operation modes. While exploring the search results, users may identify terms in the WordBars histogram that appears to be highly relevant. These term can be readily added to the query and new sets of search results generated. Similarly, while re-sorting the search results based individual query terms, users may realize that these query terms are misleading. Such query terms can be removed with a click of the mouse button, allowing users to continue exploring the search results as they seek replacement terms that are more relevant.

This free-form interaction allows searchers to use the

tools that are appropriate for their specific Web search needs. There is no need to switch between operation modes, or make a priori decisions about which tools or features to use. As their needs change, both within a single Web search and as they learn to use the system, they can choose to perform whichever activity best suits their current activities.

5. Usage Log Analysis

5.1. Methods

In order to assist in our understanding of how *TheHotMap.com* is being used, the system includes a feature that logs the interactive activities undertaken by searchers. Whenever users activate or select a feature of the system (e.g., re-order the search results, add a term to the query, etc.), this information is sent back to the server via an AJAX request, where it is processed and stored in an activity log.

The analysis in this paper is from a 15 week period during which an early beta-release of the system was publicly available. During this time, the system was not actively promoted; nor was it used by the authors of this paper. Our goal was to make the system available to the general public and measure how effectively they are able to use it. In order to reduce the learning curve associated with the introduction of a new Web search interface, a link to a tutorial was prominently displayed on the home page of the system.

5.2. Sessions

During the timeframe in which the data was logged, there were 205 unique visitors who submitted at least one query to the system. These searchers were individuals who found out about the system via word-of-mouth or who happened to discover the system and chose to try using it.

Of these 205 unique visitors of the system, nearly half chose to not make use of any of the new features provided by the system that are not seen in normal Web search engines. We characterize these users as people who were just trying the system out, but were not committed to actually conducting a search, refining their query, or exploring the search results. The analysis in this paper focuses on the remaining 104 visitors who made interactive selections using the system.

5.3. Queries

The histogram in Figure 2 represents the number of queries submitted in each unique session. The figure also shows the number of sessions in which the WordBars histogram was used to add words to the query and to re-sort




Figure 2. The frequency of query submissions and use of WordBars features.

the search results. The average number of queries per visitor was 1.95. In over half of these sessions (58 out of 104), users submitted multiple queries to the system in the same session.

It is interesting to note that not all the users who submitted multiple queries generated their subsequent queries using the features of the system. Indeed, nearly half of the additional queries were manually generated. This may be due to a significant change in the users' information needs, or a lack of understanding of the query refinement features in the interface. Further, nearly 24% of the sessions that contained only one query also included the users choosing to add a term to the query. However, these refined queries were not submitted, indicating that the users were either playing with the interface rather than conducting active searches, or did not notice that this feature was adding terms to their query.

It should be noted that while the system also supports removing terms from the query using the WordBars histogram, none of the searchers made use of this feature. Clearly, the users were much more inclined to add additional terms to their Web search queries than remove terms.

One of the design features of the system was to allow users to sort the search results using the WordBars list prior to committing to adding a potentially relevant term to the query. We can see from Figure 2 that in almost every case, the re-sorting feature was used at least as much as adding terms to the query. This provides positive evidence that the system was being used as designed with respect to this feature.

5.4. Re-Sorting

Figure 3 illustrates how frequently the searchers made use of the re-sorting features provided by the WordBars histogram and the HotMap query term headers. In 44% of the




Figure 3. The frequency of use of the re-sorting features of the system.

sessions, there was no WordBars-based re-sorting, and in 52% of the sessions, there was no HotMap-based re-sorting. Upon further analysis, it was found that 23% of the sessions included both types of re-sorting, 19% of the sessions performed neither type of re-sorting, and the remaining 58% performed one type of re-sorting or the other.

This indicates that while some searchers were willing to use multiple methods for re-sorting and exploring the search results, more than half the searchers chose to use one method or the other. This result highlights the need for providing an interface that can simultaneously support multiple methods for re-sorting the search results. As the users become more accustomed to the features of *TheHotMap.com*, we may find that they are willing to perform both types of search activities as required by their specific search goals.

5.5. Visual Inspection and Scrolling

Figure 4 represents how frequently the scroll and jump features of the HotMap overview were used in the sessions. While many of the sessions did not include the use of these features, others made moderate use of the ability to visually identify potentially relevant documents and navigate to the corresponding location in the search results list.

Although we believe that there is a great potential for using this feature to explore the search results set, it is possible that the visual cues that indicate the interactive nature of the HotMap overview are too subtle. It may not be apparent that the box representing which documents are being viewed can be dragged to scroll the search results, and that clicking on a location of interest in the HotMap overview will jump to the corresponding location in the search results. Again, as searchers become more accustomed to the interface, they may be able to discover and make effective use of this feature.




Figure 4. The frequency of use of the scroll and jump features of the HotMap overview.

6. Conclusions & Future Work

TheHotMap.com is a next-generation Web search interface that supports flexible interaction, allowing the searcher to choose the features and tools that best support their current Web search goals. The system is built upon a foundation of prior research, and is available for public use.

The key benefit of the system is the ability to support both interactive query refinement and interactive search results exploration using lightweight extensions to the common list-based representation of Web search results. The flexibility of the interaction is fundamentally important, allowing the searchers to use whichever features of the system best supports their current tasks.

The results from the usage log analysis indicate that the small set of searchers that found and used the beta-version of the system were able to readily make use of the query refinement features and the re-sorting features of the system. Further, they were able to use these features in multiple different ways, as was the goal in the design of the system. The visual inspection and scrolling features were used somewhat infrequently indicating a potential inability to identify the utility of these features.

In addition to continuing to explore methods for supporting users' Web search activities with visual and interactive tools, a longitudinal user study of the system under real-world Web search conditions is currently in the planning stages. In this study, participants will be recruited to use *TheHotMap.com* as their exclusive Web search engine for a period of three weeks. In addition to logging their activities, questionnaires will be periodically administered in order to measure their changing attitudes and impressions of the system as they learn to use the interactive query refinement and search results exploration features.

7. Acknowledgements

The authors wish to thank the University-Industry Liaison Office at the University of Regina for their support in the development of *TheHotMap.com*.

References

- [1] Ask. Ask.com search engine. <http://www.ask.com/>, 2008.
- [2] R. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley, 1999.
- [3] E. Berenci, C. Carpineto, V. Giannini, and S. Mizzaro. Effectiveness of keyword-based display and selection of retrieval results for interactive searches. *International Journal on Digital Libraries*, 3(3):249–260, 2000.
- [4] M. Hearst. Tilebars: Visualization of term distribution information in full text information access. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 1995.
- [5] M. Hearst. User interfaces and visualization. In R. Baeza-Yates and B. Ribeiro-Neto, editors, *Modern Information Retrieval*. Addison-Wesley, 1999.
- [6] T. Heimonen and N. Jhaveri. Visualizing query occurrence in search result lists. In *Proceedings of the International Conference on Information Visualization*, 2005.
- [7] O. Hoerber and X. D. Yang. A comparative user study of Web search interfaces: HotMap, Concept Highlighter, and Google. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence*, 2006.
- [8] O. Hoerber and X. D. Yang. Interactive Web information retrieval using WordBars. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence*, 2006.
- [9] O. Hoerber and X. D. Yang. The visual exploration of web search results using HotMap. In *Proceedings of the International Conference on Information Visualization*, 2006.
- [10] O. Hoerber and X. D. Yang. A unified interface for visual and interactive Web search. In *Proceedings of the IASTED International Conference on Communications, Internet, and Information Technology*, 2007.
- [11] O. Hoerber and X. D. Yang. Evaluating WordBars in exploratory Web search scenarios. *Information Processing and Management*, 44(2):485–510, 2008.
- [12] O. Hoerber, X. D. Yang, and Y. Yao. VisiQ: Supporting visual and interactive query refinement. *Web Intelligence and Agent Systems: An International Journal*, 5(3):311–329, 2007.
- [13] H. Joho, C. Coverson, M. Sanderson, and M. Beaulieu. Hierarchical presentation of expansion terms. In *Proceedings of the ACM Symposium on Applied Computing*, 2002.
- [14] Mozilla. Firefox Web browser. <http://www.mozilla.com/en-US/firefox/>, 2008.
- [15] Yahoo. Yahoo search Web services. <http://developer.yahoo.com/search>.
- [16] Yahoo. Yahoo Web search. <http://www.yahoo.com/>.