

Enhanced Web Retrieval Task

M. S. Ali
University of Toronto
sali@cs.toronto.edu

Mariano P. Consens
University of Toronto
consens@cs.toronto.edu

ABSTRACT

This paper presents the Enhanced Web Retrieval Task to model how enhanced web search engines serve the information needs of users. To evaluate the task, we model enhanced results as trees that users navigate to locate relevant information and we propose suitable measures.

1. INTRODUCTION

State-of-the-art commercial web search engines retrieve links to web pages annotated with information facets such as a text summary of key phrases in the page [5], folksonomic tags that categorize the page or site [8], links to relevant related pages [7], semantic web relationships to retrieve reviews and ratings [10], and other information to inform or entice users to review sponsored content [4]. These are aggregated search results [9] where the search engine retrieves a main link which is annotated with information facets from other sources. We refer to these type of results as *enhanced results*. Enhanced results have been shown to improve the accuracy of search results [7, 3], and improve user satisfaction of systems [6, 4].

Enhanced results are typically composed of information retrieved from across pages and sites on the web. In this paper, we propose that this retrieval paradigm can be represented as the retrieval of trees of information from the web. In the next section, we present an example and show how trees provide a basis for this paradigm. In Section 3, we propose the Enhanced Web Retrieval Task and conclude in Section 4.

2. ENHANCED RESULTS

Figure 1 shows an enhanced result from an example online movie search application based on the Yahoo! Search Monkey service [2] (the recently announced Google Rich Snippets provides a similar service). The presentation of the example result includes the main link to a retrieved movie, a summary description with details of the movie, embedded reviews of it (**hReview**'s), supporting links to provide the user with show times and ticketing information, and opinion ratings of the movie from other people on the web.

This example demonstrates how an enhanced result can satisfy the information need of users who pose the same query but have very different needs. In this work, we can



Figure 1: Enhanced web search result

model the enhanced result as a set of web links. The example includes 8 links to pages on the web; (1) more details about the movie, (2) show times and ticketing information, (3) trailers and video clips for the movie, (4,5) links to two different sites where the movie was reviewed, (6) a link to see the cast and crew who made the movie, (7) a link to recommendations to see other similar movies, and (8) the main link to the web page of the movie. It should be noted that enhanced results (such as in Figure 1) may not be optimal for all users.

The effectiveness of the search engine can be measured via inferring classical precision-recall based on the click-through rates mined from weblogs of the main link to the movie [3, 10], inferred relevance of the different information facets from click-through rates on them mined from weblogs [4], and user studies to determine user satisfaction of the retrieved information and its presentation [6]. Researchers have also considered how search results help users locate relevant information on the web via navigation. This has led to the need to also evaluate issues such as redundancy and the effort that users expend to navigate [3, 4, 7].

It is challenging to evaluate enhanced results because each facet of a result can be assessed as to whether it represents relevant information for the user. In addition, the amalgam of the facets can be assessed to determine whether they together represent a relevant answer to the user. Moreover, the web is a vast, non-homogeneous collection that spans the gamut of human knowledge in a format that is not neatly organized. The number of possible combinations of facets in a result makes it impractical to utilize pooling without introducing system bias into assessments. For instance, if two search engines retrieve the same answer but use different facets to enhance the primary part of the answer (i.e., the main link), then should this affect the relative performance

measure of these systems? We contend that it should affect performance based on how users navigate to locate relevant information from enhanced results.

We propose to model the retrieval of enhanced search results as trees of information from the web that are used to form a single answer that is structured analogously to a sitemap of the relevant links across the web. A sitemap is typically a single web page in a web site that contains a set of links to the pertinent pages of general interest to the audience of the website.

3. ENHANCED WEB RETRIEVAL TASK

We define the Enhanced Web Retrieval Task as the retrieval of a ranked list of trees of information where each contains a main link and ancillary links that answer a priori known facets of the users' information need(s). An effective system for this task helps the user to navigate to different parts of an answer that are interspersed across the web.

Tree retrieval has been proposed in [1] as a search task for retrieving trees of information from structured documents (such as XML). A key differentiator of tree retrieval from other ad-hoc structured retrieval paradigms (such as passage or element retrieval) is that the purpose of the tree is meant to improve how users navigate to relevant information and to improve how complex information (such as, in this case, enhanced results) can be encoded. Specifically, in [1], it is noted that *the task of returning trees to satisfy an information need builds on a more complex notion of relevance that extends beyond the classical content-based criterion. The relevance of a tree depends on both its content and its context. Tree retrieval involves not only finding relevant information, but also finding trees that afford users access to this information.*

For instance, the result shown in Figure 1 can be represented as a tree as shown in Figure 2. The representation of a movie in Figure 2 suggests that the user seeks a single answer to combine information facets such as whether the movie is highly rated, how to go and see the movie, and to find details that might further entice the user to go see the movie such as who are the stars in the cast. In general, any movie retrieved from the web could be encoded in this way. Enhanced Web Retrieval provides a general way to consider the retrieval of enhanced results, particularly, as in this case, where search is embedded into a focused task (such as searching the web for movies).

Tree retrieval provides a basis for representing this search task, but important questions remain. The most significant is the question of the user's information need given enhanced results. Preliminary work in aggregated search [9] addresses issues such as defining the user's core information need, aggregating information from multiple sources, presenting enhanced results, and exploring how users will interact with systems that retrieve enhanced results. In short, the key challenge will be to assess the relevance of complex enhanced results in a way that is practical and effective.

We propose the following steps to evaluate Enhanced Web Retrieval Task. First, determine a suitable way to infer relevance from web query logs [3, 10]. Second, adapt evaluation measures that consider relevance and user navigation such as structural relevance [1]. Third, utilize appropriate user navigation models, such as user browsing graphs [7].

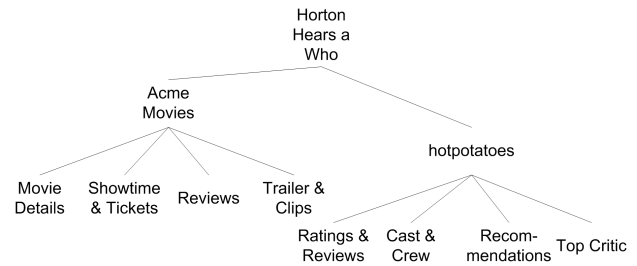


Figure 2: Tree model for enhanced result

4. CONCLUSION

To our knowledge, the Enhanced Web Retrieval Task outlined above is the first proposal in the literature for modelling enhanced search results. It can be applied to numerous, active areas in web IR including semantic relationships, opinions, sponsored content (i.e., advertising), geo-spatially localized results, personalization of search, and multilingual support in search results. A user study should be conducted to determine users' information needs and to validate whether users consider enhanced results as trees.

5. REFERENCES

- [1] M. S. Ali, M. P. Consens, G. Kazai, and M. Lalmas. Structural relevance: a common basis for the evaluation of structured document retrieval. In *CIKM '08*, pages 1153–1162, 2008.
- [2] R. Baeza-Yates, M. Ciaramita, P. Mika, and H. Zaragoza. Towards semantic search. In *NLDB '08*, pages 4–11. 2008. (See also <http://developer.yahoo.com/searchmonkey>).
- [3] M. Bilenko and R. White. Mining the search trails of surfing crowds: identifying relevant websites from user activity. In *WWW '08*, pages 51–60, 2008.
- [4] A. Fuxman, P. Tsaparas, K. Achan, and R. Agrawal. Using the wisdom of the crowds for keyword generation. In *WWW '08*, pages 61–70, 2008.
- [5] J. Goldstein, M. Kantrowitz, V. Mittal, and J. Carbonell. Summarizing text documents: sentence selection and evaluation metrics. In *SIGIR '99*, pages 121–128, 1999.
- [6] M.-Y. Kan, K. McKeown, and J. Klavans. Domain-specific informative and indicative summarization for information retrieval. In *DUC '01*, pages 19–26, 2001.
- [7] Y. Liu, M. Zhang, S. Ma, and L. Ru. User browsing graph: Structure, evolution and application. In *WSDM 2009 (Late Breaking-Results)*, 2009.
- [8] M. Melenhorst, M. Grootveld, M. van Setten, and M. Veenstra. Tag-based information retrieval of video content. In *UXTV '08*, pages 31–40, 2008.
- [9] V. Murdock and M. Lalmas. Workshop on aggregated search. *SIGIR Forum*, 42(2):80–83, 2008.
- [10] U. Shah, T. Finin, A. Joshi, R. Cost, and J. Matfield. Information retrieval on the semantic web. In *CIKM '02*, pages 461–468, 2002.