

Using Top-Ranking Sentences for Web Search Result Presentation

Ryen W. White
Information Retrieval Group
Department of Computing Science
University of Glasgow
Glasgow, Scotland. G12 8QQ.
ryen@dcs.gla.ac.uk

Joemon M. Jose
Information Retrieval Group
Department of Computing Science
University of Glasgow
Glasgow, Scotland. G12 8QQ.
jj@dcs.gla.ac.uk

Ian Ruthven
Department of Computer and
Information Sciences
University of Strathclyde
Glasgow, Scotland. G1 1XH.
ir@cis.strath.ac.uk

ABSTRACT

In this poster we propose a granular approach for presenting web search results. Sentences, taken from the top documents, are used as fine-grained representations of document content and, when combined in a ranked list, to provide an overview of the set of retrieved documents. Current search engine interfaces assume users examine such results document-by-document. In contrast our approach groups, ranks and presents the *contents* of the top ranked document *set*. We evaluate our approach by a comparison with traditional forms of web search result presentation.

Keywords

Top-ranking Sentences, Web Search, Visualisation

1. INTRODUCTION

Web searchers can find the formulation of queries that adequately express their information need a demanding process. However, such searchers may face even more difficulty when interpreting and assessing the relevance of the documents returned in response to the query. Users of web search engines are typically unwilling to examine large sets of individual documents and base initial judgments on what documents to view on *surrogates* such as titles, abstracts and URLs. These can be manually created (e.g. titles or keywords), or automatically created (e.g. summaries).

Presenting lists of document surrogates has remained a popular method of presenting search results. Lists allow documents to be *ranked* in order of their estimated utility to the user. However, lists encourage users to read, interpret and assess documents and their surrogates *individually*.

In this poster we suggest techniques that encourage a deeper examination of documents at the results interface and blur inter-document boundaries. We shift the focus of interaction from the document surrogate to the document's content. We compare traditional methods of producing surrogates (such as text fragments and titles) against a new method of presenting search results; sentences taken from the retrieved documents, ranked on how closely they match the user's query. This set of *top-ranking sentences* can be used to form an overview of the returned document set.

2. MOTIVATION

Web search engines are intended to help people find information that is relevant to completing a task. It is important therefore to design interfaces that maximise the amount of useful information users can obtain within a search.

Searchers use textual queries to communicate their need with the search system. The query is only an approximate description of the information need [1]. Web documents are ranked algorithmically based on this query and returned in a list to the user. These may not be *entirely* relevant, and it is the relevant parts that contribute most to satisfying the user's information need. By ranking *documents* we assume that all of a document conforms to relevance/matching criteria. This assumption is often incorrect as documents can have irrelevant parts. Research into summarisation and visualisation have tackled this problem, but still return document lists to users.

In our approach, we use a technique known as *sentence extraction* to present whole sentences to users, taken from the top thirty documents in the retrieved document set. These sentences provide a high level of granularity, removing the restriction of document boundaries and shifting the focus from the document to the information it contains. This means that users are not forced to access information through documents but through the actual content of documents. Through ranking this information with respect to the query, the user is given a query-specific overview of the *content* of the returned set. A document list is biased towards the user's information need at the document level. Documents that are a close match to the user's query appear near the top of the list. In our approach we bias at the sentence level. Sentences that are a close match to the user's query are shown near the top of a ranked list of sentences.

In the next section we describe an interface that implements the concepts discussed in the poster so far.

3. INTERFACE

The interface uses the Google¹ commercial search engine to search the Internet. In response to a query submitted by the user, the system returns a ranked list of document titles, abstracts and URLs, thirty in total.

Each document in this list is then downloaded and all sentences from each document extracted. Each sentence is assigned a score, using an algorithm similar to that in [2]. This uses factors such as position of the sentence in document and the presence of any emphasised words. In addition sentences receive additional scores depending on the proportion of query terms contained within the sentence. This *query-biasing* component biases the scoring mechanism to sentences that use words contained within the user's query.

¹ <http://www.google.com>

It then pools the top four scoring sentences from each document, and ranks all sentences in a ‘global’ list. This list is a query-biased overview of the returned document set. The sentences are shown individually, with query terms highlighted.

Users are not shown a list of retrieved document titles and URLs; only the list of top-ranking sentences. Initially there is no direct association between the sentence and its source document, i.e. there is no indication to the user of which document supplied each sentence. To view the association, the user must move the mouse pointer over a sentence. When this occurs, the sentence is highlighted and a window pops up next to it. Displaying this window next to the sentence, instead of in a fixed position on the screen, is intended to make the sentence-document relationship more lucid. The interface is shown in Figure 1 (with the sentences and popup window marked ① and ② respectively).

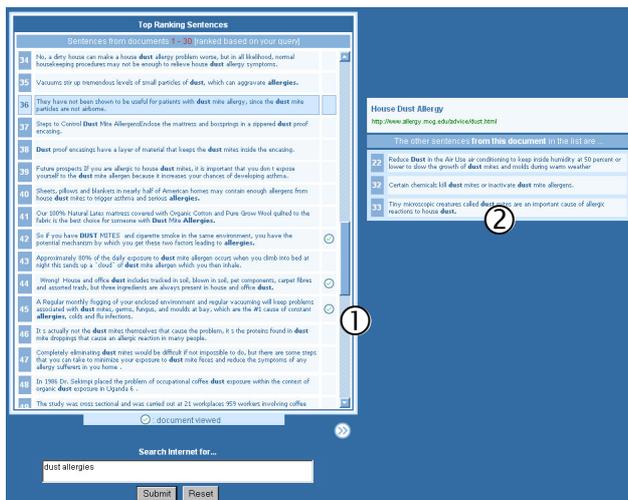


Figure 1: Top-ranking sentences results interface

In the popup window the user is shown the document title, URL and the rank position and content of any other sentences from that document that occur in the list of top-ranking sentences. If no other sentences appear an appropriate message is shown.

To visit a document the user must click the highlighted sentence, or any sentences in the pop-up window. It is the sentences (content) that drive the interaction. When the user has clicked a sentence and visited that document, all sentences from that document are marked to reflect this.

We conducted an evaluation, where we compared this granular top-ranking sentences approach (*experimental* system) with two systems using document lists; one is a traditional web search engine and the other is similar, but with substantially longer query-biased summaries (up to four sentences) for each document. These two systems are referred to as *baselines* in the remainder of this poster. This evaluation is described in the next section.

4. EVALUATION

A total of 18 subjects took part in the evaluation, each completed 3 tasks (*fact* search, *decision* search, *background* search), one on each of the 3 search systems. Tasks and systems were allocated according to a Greco-Latin square design. Each subject was given 10 minutes to complete each task, although the subjects could terminate the search early if they felt they had completed the task. The time restriction was imposed to ensure consistency between subjects. We elicited subject opinion using informal interviews

and questionnaires (before the first task, after each task and after all tasks). We used semantic differentials, Likert scales and open-ended questions to collect this data. Background system logging recorded subject interaction with each search interface.

Subjects were divided into two groups: inexperienced and experienced. This classification was made on the basis of the subjects’ responses to questions on the level of their computing, Internet and web searching experience.

5. RESULTS & DISCUSSION

We investigated whether using sentences, extracted from top ranked documents and presented in a query-biased ranked list would be more effective than traditional forms of web search result presentation. Our results show that our approach is liked by users and increases the effectiveness of their searching. For example, top-ranking sentences reduced task completion time and increased the number of tasks that were completed. Overall, 16 of the 18 subjects (9 experienced and 7 inexperienced) preferred the top-ranking sentences approach and found it more helpful, useful and effective than both baselines.

To make sound judgments on the effectiveness of a submitted query, searchers should be able to assess the actual content of the document set, not simply document surrogates. In the experimental system subjects were more aware of this content and document titles became less useful as a result.

Presenting a list of top-ranking sentences encouraged subjects to view documents outwith the first page of results. On the baseline systems subjects would rather reformulate and resubmit their queries than deeply peruse the documents returned to them or click the ‘next’ button. By doing so they discard potentially relevant documents without giving them due consideration. The document list returned is only an algorithmic match to the user’s query. Unless the information need is very specific the system may struggle to provide a ranking that is a match for the user’s information need. This problem is amplified if the system only ranks whole documents as small highly relevant sections may reside in documents with a low overall ranking.

In this poster we present an approach for presenting web search results that shifts the focus of perusal and interaction away from the document surrogates, such as document titles, abstracts and URLs, to the actual content of the document. The results of our evaluation have shown, with statistical significance, that ranking the content of the retrieved document set rather than the documents themselves leads to increased searcher efficiency, effectiveness and personal preference.

ACKNOWLEDGEMENTS

We would like to thank the experimental subjects, their comments and enthusiasm were very much appreciated. The work reported in this poster is funded by the UK Engineering and Physical Sciences Research Council grant number GR/R74642/01.

REFERENCES

[1] Taylor, R.S. ‘Question-negotiation and information seeking in libraries’. College and Research Libraries. **29**. pp. 178-194. 1968.

[2] White, R.W., Jose, J.M. and Ruthven, I. ‘A task-oriented study on the influencing effects of query-biased summarisation in web searching’. Information Processing and Management. 2003. *in press*.