

Diversified Relevance Feedback

Matt Crane
Department of Computer Science
University of Otago
Dunedin, New Zealand
mcrane@cs.otago.ac.nz

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval – Search process

General Terms

Experimentation

Keywords

Diversification; Relevance feedback

ABSTRACT

The need for a search engine to deal with ambiguous queries has been known for a long time (diversification). However, it is only recently that this need has become a focus within information retrieval research. How to respond to indications that a result is relevant to a query (relevance feedback) has also been a long focus of research. When thinking about the results for a query as being clustered by topic, these two areas of information retrieval research appear to be opposed to each other. Interestingly though, they both appear to improve the performance of search engines, raising the question: they can be combined or made to work with each other?

When presented with an ambiguous query there are a number of techniques that can be employed to better select results. The primary technique being researched now is diversification, which aims to populate the results with a set of documents that cover different possible interpretations for the query, while maintaining a degree of relevance, as determined by the search engine. For example, given a query of “java” it is unclear whether the user, without any other information, means the programming language, the coffee, the island of Indonesia or a multitude of other meanings.

In order to do this the assumption that documents are independent of each other when assessing potential relevance has to be broken. That is, a documents relevance, as calculated by the search engine, is no longer dependent only on the query, but also the other documents that have been selected. How a document is identified as being similar to previously selected documents, and the trade off

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR '13, July 28–August 1, 2013, Dublin, Ireland.
ACM 978-1-4503-2034-4/13/07.

between estimated relevance and topic coverage are current areas for information retrieval research.

For unambiguous queries, or for search engines that do not perform diversification, it is possible to improve the results selected by reacting to information identifying a given result as truly relevant or not. This mechanism is known as relevance feedback. The most common response to relevance feedback is to investigate the documents for their most content-bearing terms, and either add, or subtract, their influence to a newly formed query which is then re-run on the remaining documents to re-order them.

There has been a scant amount of research into the combination of these methods. However, Carbonell *et al.* [1] show that an initially diverse result set can provide a better approach for identifying the topic a user is interested in for a relevance feedback style approach. This approach was further extended by Raman *et al.* [4].

An important aspect of relevance feedback is the selection of documents to use. In the 2008 TREC relevance feedback track, Meij *et al.* [3] generated a diversified result set which outperformed other rankings as a source of feedback documents.

The use of pseudo-relevance feedback (assuming the top ranked documents are relevant) to extract sub-topics for use in diversification was explored by Santos *et al.* [5]. These previous approaches suggest that these two ideas are more linked than expected.

The ATIRE search engine [6] will be used to further explore the relationship between diversification and relevance feedback. ATIRE was selected because it is developed locally, and is designed to be small and fast. ATIRE also produces a competitive baseline, which would have placed 6th in the 2011 TREC diversity task while performing no diversification and index-time spam filtering [2], although we concede this is not equivalent to submitting a run.

1. REFERENCES

- [1] J. Carbonell and J. Goldstein. The use of MMR, diversity-based reranking for reordering documents and producing summaries. In *SIGIR 1998*, pages 335–336.
- [2] M. Crane and A. Trotman. Effects of spam removal on search engine efficiency and effectiveness. In *ADCS 2012*, pages 1–8.
- [3] E. Meij, J. He, W. Weerkamp, and M. De Rijke. Topical diversity and relevance feedback. In *TREC 2009*.
- [4] K. Raman, T. Joachims, and P. Shivaswamy. Structured learning of two-level dynamic rankings. In *CIKM 2011*, pages 291–296.
- [5] R. L. T. Santos, J. Peng, C. Macdonald, and I. Ounis. Explicit search result diversification through sub-queries. *Advances in information retrieval*, pages 87–99, 2010.
- [6] A. Trotman, X. F. Jia, and M. Crane. Towards an efficient and effective search engine. In *SIGIR 2012 Workshop on Open Source Information Retrieval*, pages 40–47.