Axiomatic Result Re-Ranking

(Extended Abstract)*

Matthias Hagen Michael Völske Steve Göring Benno Stein

Bauhaus-Universität Weimar
99421 Weimar, Germany
<first name>.<last name>@uni-weimar.de

ABSTRACT

We consider the problem of re-ranking the top-k documents returned by a retrieval system given some search query. This setting is common to learning-to-rank scenarios, and it is often solved with machine learning and feature weighting based on user preferences such as clicks, dwell times, etc. In this paper, we combine the learning-to-rank paradigm with the recent developments on axioms for information retrieval. In particular, we suggest to re-rank the top-k documents of a retrieval system using carefully chosen axiom combinations.

**Original paper with all omitted details in CIKM 2016 [2]. A pre-print is available from http://www.uni-weimar.de/medien/webis/publications/papers/stein_2016o.pdf*

1. INTRODUCTION

Axioms for information retrieval play a rather "theoretical" role so far. Most of the respective studies focus on the question of whether the results of known retrieval models are in accordance with specific reasonable axioms that formalize ranking preferences. For instance, from two documents of the same length, the one containing the query terms more often should be favored. The notion of axiomatic retrieval in the modern sense first appeared in a study of Bruza and Huibers [1], which proposes a way of formalizing what should be expected from a good result ranking. Especially in the last decade, interest in axiomatic retrieval has increased substantially. Hui Fang's web page gives a good overview of the existing literature and axioms.[1] The goal of most of the recent studies is to propose reasonable axioms and to evaluate how well existing retrieval models match these assumptions. Researchers have analyzed a wide range of standard retrieval models for conformance to the proposed axioms and, at times, suggested certain adjustments to the models. We take up this axiomatic view—but, instead of adjusting the retrieval models themselves, we suggest the following innovation: to adopt the learning-to-rank idea and to re-rank the top-k results *directly* using promising axiom combinations. This way, we can turn every reasonable basic retrieval model into an axiom-based retrieval model.

This observation leads to our main research question: Is it possible—and how—to seamlessly integrate axioms for ranking preferences in order to improve the results of a basis retrieval model? Our proposed solution is inspired by the learning-to-rank framework: Given some basis retrieval model, a carefully weighted axiom combination re-ranks the top-k results and produces an axioms-compliant output.

[1] http://www.eecis.udel.edu/~hfang/AX.html

2. AXIOMATIC RE-RANKING

Most axioms in the information retrieval literature have a similar basic structure: for a pair or a triple of documents, ranking preferences are deduced from standard features such as document length, term frequency, or semantic similarity. When such an axiom is applied to all pairs or triples of documents in a retrieval model's result list, the matrix of the inferred preferences may induce a result re-ranking.

For example, consider a situation with an axiom A and three initially retrieved documents d_1, d_2, and d_3. After applying axiom A to all document pairs, one might end up with the preferences $d_2 >_A d_1$, $d_2 >_A d_3$, and $d_1 >_A d_3$, where $d_i >_A d_j$ means that document d_i should be ranked above d_j according to axiom A. Only the ranking $[d_2, d_1, d_3]$ matches these preferences and will thus become the re-ranked document list. In the general case there are many axioms of different importance and contradictory rank preferences become likely. As a solution and a way of combining the weighted axioms' matrices of rank preferences, we apply fusion algorithms that were developed in the field of computational social choice.

3. CONTRIBUTIONS

The effectiveness of our axiom-based retrieval system is studied in a large-scale evaluation with 17 basis retrieval models in the setting of the TREC Web tracks 2009–2014. The performance of almost all basis retrieval models is improved via axiomatic result re-ranking. It is thus possible to improve existing retrieval models in an "ex-post manner," considering the latest insights from the research on retrieval axioms. The main contributions of our work are: (1) We show how combinations of known axioms can be incorporated into a learning-to-rank inspired result re-ranking for any given basis retrieval model. The resulting axiom-based scheme is shown to significantly increase retrieval performance for many standard models. (2) We propose axioms to model term proximity preferences and show their effect in the axiomatic re-ranking.

References

[1] Peter Bruza and T. W. C. Huibers. Investigating aboutness axioms using information fields. In *Proceedings of SIGIR'94*, pages 112–121. ACM/Springer, 1994.

[2] Matthias Hagen, Michael Völske, Steve Göring, and Benno Stein. Axiomatic Result Re-Ranking. In *Proceedings of (CIKM'16) (to appear)*. ACM, October 2016.