Axiomatic Thinking for Information Retrieval – And Related Tasks

Enrique Amigo
National Distance Education University
SPAIN
enrique@lsi.uned.es

Stefano Mizzaro
University of Udine
Italy
mizzaro@uniud.it

Hui Fang
University of Delaware
USA
hfang@udel.edu

ChengXiang Zhai
University of Illinois at Urbana-Champaign
USA
czhai@illinois.edu

1 MOTIVATION

The main task of an Information Retrieval (IR) system is to return relevant documents to users in response to a query. However, this task is inherently an empirical task since the definition of *relevance* of a document to a query is not well defined, and in general, can only be judged by users who issued the query. Yet, defining relevance as rigorously as we can is essential to the development of both effective IR systems and sound evaluation metrics. As a result, modeling relevance has always been a central challenge in IR research for both retrieval model development and evaluation. Indeed, all information retrieval models developed so far (which are the basis of the algorithms used in all search engine applications) has explicitly or implicitly adopted one way or another to formalize the vague concept of relevance; similarly, all evaluation metrics of IR are meant to quantify the utility of the retrieved results from a user's perspective, and thus must also accurately reflect a user's view of relevance.

For many years, however, the notion of relevance has been modeled mostly as a "black box" without paying attention to the specific properties that a retrieval function or evaluation metric must satisfy from the perspective of relevance modeling; as such, there is often no way to analytically assess the soundness of a retrieval model or evaluation metric, hindering the study of their optimality. Recently, axiomatic thinking has been adopted for the development of both retrieval models and evaluation metrics with great promise. The general idea of axiomatic thinking is to seek a set of desirable properties expressed mathematically as formal constraints to guide the search for an optimal solution; the explicit expression of desirable properties makes it possible to analytically address issues that would otherwise appear to be purely empirical, provide theoretical guidance on how we might be able to optimize a retrieval model or evaluation metric, and apply any identified constraints directly in many practical applications.

The growth of research on axiomatic thinking for IR can be easily seen from the increasing number of publications on this general

topic (over 40 papers published recently), mostly in two lines. The first line is the application of axiomatic analysis to retrieval model development, where relevance is modeled based on a set of retrieval constraints that any reasonable retrieval functions need to satisfy. These constraints are then used to diagnose deficiencies of existing retrieval functions and guide the search for more effective retrieval functions. Axiomatic analysis has been shown to be effective for diagnosing deficiencies of basic retrieval models and improving them [3, 5, 6], including particularly the development of BM25+, an improvement of a long-standing state of the art model (i.e., BM25) to fix its deficiency in document length normalization [9]. The second line is the successful application of axiomatic analysis to evaluation, in particular, to formalizing evaluation metrics. In general, these studies focus on setting particular situations in which metrics should behave in a certain manner, that are specified by means of constraints or axioms that metrics should satisfy. This theme recurs in the literature since the 70s, but it has received an increasing interest in the last five or ten years, when most of the papers have been published [1, 2, 4, 7, 8, 10, 12–14]. Moreover, the formalization attempts of evaluation metrics have gone well beyond the IR field: several results concern related areas like classification [13] or clustering [1, 11], or even textual similarity, opinion mining, and so on.

The goal of workshop is to bring together researchers and practitioners interested in applying axiomatic analysis to all kinds of IR and IR-related problems, including particularly both those interested in developing retrieval models and those interested in developing evaluation measures, and to enable them to share their findings (both positive or negative), to present their latest research results, and to discuss future directions.

2 THEME

The general theme of the workshop will be about all aspects of applications of axiomatic thinking to solve IR and IR-related problems. The basis of this general theme is the recent growth of work on applying axiomatic thinking to analyze and improve both retrieval models and evaluation metrics, which we expect to continue. The existing work has clearly demonstrated many advantages of axiomatic thinking, including particularly specific theoretical results in the form of novel constraints and improved retrieval models and evaluation metrics. However, much more research is still needed in multiple directions.

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'17, August 7–11, 2017, Shinjuku, Tokyo, Japan

© 2017 Copyright held by the owner/author(s). 978-1-4503-5022-8/17/08.

DOI: 10.1145/3077136.3084369

Take, for example, the study of retrieval models. Retrieval constraints are often formalized based on retrieval heuristics, and the heuristics are strategies motivated by empirical observations. Unfortunately, there has been little attention on identifying, formalizing and testing useful heuristics of retrieval models, and learning how to utilize them to develop effective retrieval models. Most heuristics lie in the minds of seasoned information retrieval researchers and exercising heuristics effectively is something of a black art. Clearly, it has come to a point where the constraints (i.e., heuristics) of retrieval models for various tasks should be revisited, organized, summarized and analyzed.

Also, it has been shown that, in most existing retrieval models, there exist some underlying heuristics that are closely related to the empirical performance. For example, many models include TF-IDF style term weighting. In addition, extensions to these models often also rely upon very similar techniques [15]. These sorts of extensions include term proximity and pseudo-relevance feedback. Ideally, as we discover these shared constraints, we would have more guidance on developing effective retrieval models. Since there exist some commonality among various retrieval tasks, discovering and summarizing the effectiveness of different constraints would deepen our understanding and open some new research directions.

Opportunities of applying axiomatic thinking also go beyond analyzing the basic retrieval functions; in fact, understanding constraints is also beneficial to many IR tasks that use machine learning techniques. Instead of having a designer carefully choose a set of assumptions to make when designing a formal model, these approaches use machine learning to weight items in a pool of features derived from many retrieval heuristics. However, this potentially results in a bloated backend which computes many features irrelevant to the task or collection. Having knowledge about relevant features would help slim down backends and speed learning and ranking. An important strength of the axiomatic methodology is that evaluation data sets become resources used to check motivated hypotheses instead of optimization mechanisms, which are at risk of overfitting.

There are even more opportunities for new research on applying axiomatic thinking to evaluation as has already been happening where researchers have done axiomatic analysis of metrics for tasks such as text categorization. In general, an understanding of how to apply axiomatic thinking to IR problems may become increasingly important as information retrieval continues to broaden into new areas. New tasks often require new constraints, and an understanding of these constraints can provide guidance on how to adapt existing methods or how to develop new methods for the new tasks. For example, domain-specific IR tasks such as medical record search might require new retrieval constraints that can capture the domain knowledge.

3 WORKSHOP AGENDA

The workshop aims to bring together researchers and practitioners from a broader community to exchange research ideas and results and foster collaborations across subcommunities. Three synergistic activities are planned: invited talks, technical presentations, and a panel discussion.

The expected outcomes are summarized as follows.

- We want to engage a broader discussion of “axiomatic thinking” in general in terms of its potential broad impact on optimizing evaluation metrics for any tasks (not just IR tasks) and optimizing ranking for many other tasks.
- We hope to discuss and summarize the general guidances on how to apply axiomatic thinking to a new problem, including particularly, (1) how to identify potential constraints; (2) how to verify those constraints; (3) how to apply these constraints to search for a solution; and (4) how to refine the constraints based on the evaluation results.
- We aim to form a community for researchers and practitioners who are interested in axiomatic thinking, build shared resources and infrastructure, and identify important future research directions, especially promising interdisciplinary topic areas.

REFERENCES

- [1] Enrique Amigó, Julio Gonzalo, Javier Artiles, and Felisa Verdejo. 2009. A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Information Retrieval* 12, 4 (2009), 461–486.
- [2] Enrique Amigó, Julio Gonzalo, and Felisa Verdejo. 2013. A General Evaluation Measure for Document Organization Tasks. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '13)*. ACM, New York, NY, USA, 643–652.
- [3] P. D. Bruza and T. W. C. Huibers. 1994. Investigating aboutness axioms using information fields. In *Proceedings of the 17th ACM SIGIR*. Springer-Verlag New York, Inc., New York, NY, USA, 112–121. <http://dl.acm.org/citation.cfm?id=188490.188521>
- [4] Luca Busin and S. Mizzaro. 2013. Axiometrics: An Axiomatic Approach to Information Retrieval Effectiveness Metrics. In *Proceedings of ICTIR 2013: 4th International Conference on the Theory of Information Retrieval*. ACM, New York – USA, 22–29.
- [5] Hui Fang, Tao Tao, and Chengxiang Zhai. 2011. Diagnostic Evaluation of Information Retrieval Models. *ACM Trans. Inf. Syst.* 29, 2, Article 7 (April 2011), 42 pages. DOI: <http://dx.doi.org/10.1145/1961209.1961210>
- [6] Hui Fang and Chengxiang Zhai. 2005. An exploration of axiomatic approaches to information retrieval. In *SIGIR '05*. 480–487. DOI: <http://dx.doi.org/10.1145/1076034.1076116>
- [7] Marco Ferrante, Nicola Ferro, and Maria Maistro. 2015. Towards a Formal Framework for Utility-oriented Measurements of Retrieval Effectiveness. In *Proceedings of ICTIR 2015*. ACM, New York, NY, USA, 21–30. DOI: <http://dx.doi.org/10.1145/2808194.2809452>
- [8] César Ferri, José Hernández-Orallo, and R. Modroiu. 2009. An experimental comparison of performance measures for classification. *Pattern Recognition Letters* 30, 1 (2009), 27–38.
- [9] Yuanhua Lv and Chengxiang Zhai. 2011. Lower-bounding Term Frequency Normalization. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management (CIKM '11)*.
- [10] Eddy Maddalena and S. Mizzaro. 2014. Axiometrics: Axioms of Information Retrieval Effectiveness Metrics. In *Proceedings of the Sixth EVIA Workshop*. National Institute of Informatics, Tokyo, Japan, 17–24. ISBN: 978-4-86049-066-9.
- [11] Marina Meila. 2003. Comparing clusterings. In *Proceedings of COLT 03*. <http://www.stat.washington.edu/mmp/www.stat.washington.edu/mmp/Papers/compare-clust.pdf>
- [12] Alistair Moffat. 2013. Seven Numeric Properties of Effectiveness Metrics. In *AIRS '13*. 1–12.
- [13] Fabrizio Sebastiani. 2015. An Axiomatically Derived Measure for the Evaluation of Classification Algorithms. In *Proceedings of ICTIR 2015*. ACM, 11–20. DOI: <http://dx.doi.org/10.1145/2808194.2809449>
- [14] Marina Sokolova. 2006. Assessing Invariance Properties of Evaluation Measures. In *Proceedings of NIPS'06 Workshop on Testing Deployable Learning and Decision Systems*.
- [15] Chengxiang Zhai. 2008. Statistical Language Models for Information Retrieval A Critical Review. *Found. Trends Inf. Retr.* 2, 3 (March 2008), 137–213.

Acknowledgments. We thank the support and efforts of our keynote speaker, panelists and the program committee members.

We acknowledge the support by the U.S. National Science Foundation under IIS-1423002.