UvA-DARE (Digital Academic Repository)

Search engines that learn online

Hofmann, K.

DOI

[10.1145/2009916.2010176](https://doi.org/10.1145/2009916.2010176)

Publication date

2011

Document Version

Final published version

Published in

SIGIR'11

[Link to publication](#)

Citation for published version (APA):

Hofmann, K. (2011). Search engines that learn online. In *SIGIR'11: proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval ; July 24-28, 2011, Beijing, China* (pp. 1313). ACM. <https://doi.org/10.1145/2009916.2010176>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Search Engines that Learn Online

Katja Hofmann
ISLA, University of Amsterdam
K.Hofmann@uva.nl

ABSTRACT

The goal of my research is to develop self-learning search engines, that can learn *online*, i.e., directly from interactions with actual users. Such systems can continuously adapt to user preferences throughout their lifetime, leading to better search performance in settings where expensive manual tuning is infeasible. Challenges that are addressed in my work include the development of effective *online learning to rank algorithms* for IR, *user aspects*, and *evaluation*.

Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]: H.3.3 Information Search and Retrieval

General Terms

Algorithms, Experimentation

Keywords

Online learning to rank, Implicit feedback, Evaluation

1. INTRODUCTION

As search engines are becoming increasingly complex, and combine more and more sources of information, the question of how to tune these systems to optimize search results becomes essential. Current methods for optimizing ranking functions typically work *offline*, meaning that systems are tuned to a given data set before deployment. While offline methods are appropriate for many settings and can result in excellent search performance, there are limitations. Data needed for tuning is often created manually, and therefore expensive and error-prone.

We address the limitations of offline methods by working towards self-learning search engines, i.e., search engines that learn ranking functions *online*. Learning online means that there is no distinction between training and operating. Instead, the search engine observes users' natural interactions with the search interface (so-called *implicit feedback*, e.g., clicks on result documents), infers information from them, and improves its ranking function automatically. Because implicit feedback is used for learning, expensive data collection is not required, and the collected data matches the target users and the target setting.

2. RESEARCH QUESTIONS

Several steps have been taken to make online learning to rank for IR a reality, but many open challenges remain. My thesis work addresses the following questions.

RQ 1 How can retrieval systems *learn effectively* from implicit feedback in an online settings?

Online learning to rank for IR differs from other learning to rank tasks, because of the characteristics of the implicit feedback available for learning. While correlated with relevance, it varies widely across tasks, users, and queries. As a result, only noisy, relative preferences can be inferred. Furthermore, the feedback that can be obtained is biased towards the top-ranked documents displayed.

Our first results to address these issues focused on the problem of balancing *exploration* and *exploitation*, meaning that result lists to be displayed to the user need to both satisfy the user's information need as much as possible (i.e., exploit what has already been learned), while also allowing to collect feedback that can provide as much new information as possible (i.e., explore new solutions). We found that such a balance can significantly improve the performance of online learning to rank systems [1].

RQ 2 What *user aspects* need to be considered?

Previous work in online learning to rank for IR has been grounded in user studies, but many open questions remain. Here, we investigate the relation between explicit and implicit feedback, and develop methods for inferring information from implicit feedback more reliably. Our first results include a new, probabilistic method for comparing rankers using click data, that is more accurate and more robust to noise than previous methods [2].

RQ 3 How can online learning to rank methods be *evaluated*?

Online learning to rank approaches for IR can be evaluated using live search engine traffic, lab studies, or log data. Each of these has advantages and limitations. Moving forward, we need to analyze the benefits and limitations of existing evaluation methods for assessing online learning to rank approaches, and to develop new approaches to address any identified limitations.

Some online learning to rank approaches have been evaluated on life traffic of large-scale web search engines. However, such a setup is not always accessible or feasible, especially when potentially risky new methods need to be evaluated, or when many comparisons under a wide range of settings are required. As an alternative to real-life experiments, we developed an evaluation setup that simulates user behavior at a sufficient level of detail, using learning to rank data sets and a model of click behavior [1]. In the future, we need to continue to assess and refine this methodology.

References

- [1] K. Hofmann, S. Whiteson, and M. de Rijke. Balancing exploration and exploitation in learning to rank online. In *ECIR'11*, pages 251–263, 2011.
- [2] K. Hofmann, S. Whiteson, and M. de Rijke. A probabilistic method for inferring preferences from clicks. Under review, 2011.