

Fast and Reliable Online Learning to Rank for Information Retrieval

Katja Hofmann
University of Amsterdam, The Netherlands
k.hofmann@uva.nl

The amount of digital data we produce every day far surpasses our ability to process this data, and finding useful information in this constant flow of data has become one of the major challenges of the 21st century. Search engines are one way of accessing large data collections. Their algorithms have evolved far beyond simply matching search queries to sets of documents. Today's most sophisticated search engines combine hundreds of relevance signals to provide the best possible results for each searcher.

Current approaches for tuning the parameters of search engines can be highly effective. However, they typically require considerable expertise and manual effort. They rely on *supervised* learning to rank, meaning that they learn from manually annotated examples of relevant documents for given queries. Obtaining large quantities of sufficiently accurate manual annotations is becoming increasingly difficult, especially for personalized search, access to sensitive data, or search in settings that change over time.

In this thesis, I develop new *online* learning to rank techniques, based on insights from reinforcement learning. In contrast to supervised approaches, these methods allow search engines to learn directly from users' interactions. User interactions can typically be observed easily and cheaply, and reflect the preferences of real users. Interpreting user interactions and learning from them is challenging, because they can be biased and noisy. The contributions of this thesis include a novel interleaved comparison method, called probabilistic interleave, that allows unbiased comparisons of search engine result rankings, and methods for learning quickly and effectively from the resulting relative feedback.

The obtained analytical and experimental results show how search engines can effectively learn from user interactions. In the future, these and similar techniques can open up new ways of gaining useful information from ever larger amounts of data.

The thesis is available online at <http://katja-hofmann.de> and <http://dare.uva.nl/record/446342>. Software written for this thesis, including reference implementations of the developed interleaving and online learning methods, is available at <http://ilps.science.uva.nl/resources/online-learning-framework>.

Supervisors: Maarten de Rijke and Shimon Whiteson (University of Amsterdam, NL). **Committee:** Nick Craswell (Bing, Microsoft, US), Lynda Hardman (University of Amsterdam, NL), Thorsten Joachims (Cornell University, US), Arjen de Vries (Delft University of Technology, NL), Max Welling (University of Amsterdam, NL). **Defended:** May 28, 2013 in Amsterdam.