Adapting Rankers Online

Katja Hofmann, Shimon Whiteson, and Maarten de Rijke

ISLA, University of Amsterdam, Science Park 904, 1098 XH Amsterdam
{K.Hofmann,S.A.Whiteson,deRijke}@uva.nl

At the heart of many effective approaches to the core information retrieval problem—identifying relevant content—lies the following three-fold strategy: obtaining content-based matches, inferring additional ranking criteria and constraints, and combining all of the above so as to arrive at a single ranking of retrieval units.

Over the years, many models have been proposed for content-based matching, with particular attention being paid to estimations of query models and document models. Different task and user scenarios have given rise to the study and use of priors and non-content-based ranking criteria such as freshness, authoritativeness, and credibility. The issue of search result combinations, whether ranked-based, score-based or both, has been a recurring theme for many years. As retrieval systems become more complex, *learning to rank* approaches are being developed to automatically tune the parameters for integrating multiple ways of ranking documents. This is the issue on which we will focus in the talk.

Search engines are typically tuned offline; they are tuned manually or using machine learning methods to fit a specific search environment. These efforts require substantial human resources and are therefore only economical for relatively large groups of users and search environments. More importantly, they are inherently static and disregard the dynamic nature of search environments, where collections change and users acquire knowledge and adapt their search behaviors. Using *online* learning to rank approaches, retrieval systems can learn directly from implicit feedback, while they are running.

The talk will discuss three issues around online learning to rank: balancing exploitation and exploration, gathering data using one pair of rankers and using it to compare another pair of rankers, and the use of rich contextual data.

Balancing exploitation and exploration. In an online setting, algorithms need to both explore new solutions to obtain feedback for effective learning, and exploit what has already been learned to produce results that are acceptable to users. In recent work [1], we have formulated this challenge as an *exploration-exploitation dilemma* and present the first online learning to rank algorithm that works with implicit feedback and balances exploration and exploitation. We leverage existing learning to rank data sets and recently developed click models to evaluate the proposed algorithm. Our results show that finding a balance between exploration and exploitation can substantially improve online retrieval performance, bringing us one step closer to making online learning to rank work in practice.

Generalizing to novel rankers. Implicit feedback, such as users' clicks on documents in a result list, is increasingly being considered as an alternative to explicit relevance judgments. For example, previous work has shown that click data can be used to detect

A. Hanbury, A. Rauber, and A.P. de Vries (Eds.): IRFC 2011, LNCS 6653, pp. 1–2, 2011.

even small differences between rankers, and that it can be used for online learning to rank. Previous methods can identify the better of two rankers with high confidence, but currently the data collected for comparing one pair of rankers cannot be reused for other comparisons. As a result, the number of rankers that can be compared is limited by the amount of use of a search engine. In the talk, ongoing work will be presented on re-using previously collected, historical data by applying importance sampling to compensate for mismatches between the collected data and distributions under the target data. We will show that in this way, rankers can be compared effectively using historical data.

Contextual data. The last part of the talk will be forward-thinking and focus on future work and challenges. Most retrieval systems are integrated in a larger contextual setting, where no item is an island. Events in one document stream are correlated with events in another. Increasingly rich declarative models describe the task, workflow, interaction and organisational structure. Structured knowledge, for instance in the form of linked open data, is available in large quantities to help us inform our retrieval algorithms. How can we use these sources of information in an online setting?

Reference

[1] Hofmann, K., Whiteson, S., de Rijke, M.: Balancing exploration and exploitation in learning to rank online. In: Clough, P., Foley, C., Gurrin, C., Jones, G.J.F., Kraaij, W., Lee, H., Mudoch, V. (eds.) ECIR 2011. LNCS, vol. 6611, pp. 251–263. Springer, Heidelberg (2011)