

# Don't Hurt Them: Learning to Rank from Historical Interaction Data (Keynote)

Maarten de Rijke

University of Amsterdam, The Netherlands

One of the main advantages of online evaluation schemes is that they are user-based and, as a result, often assumed to give us more realistic insights into the real system quality than off-line methods. This is also one of their main disadvantages: comparing two rankers online requires presenting users with result lists based on those rankers and observing how users interact with them. New rankers may perform sub-optimal and hence hurt the user experience. Can we use or re-use historical data, collected from user interactions with a production system, to assess or optimize new alternative rankers? This question has increasingly gained interest in the past few years. In the talk I will contrast several proposals for learning from historical interaction, based on importance sampling, random buckets, and a Bayesian approach based on explicit user models.

This is based on joint work with Artem Grotov, Katja Hofman, Damien Lefortier, Anne Schuth, Shimon Whiteson.