UvA-DARE (Digital Academic Repository)

Don't hurt them: Learning to rank from historical interaction data

(Keynote)

de Rijke, M.

DOI

[10.1007/978-3-319-12979-2](https://doi.org/10.1007/978-3-319-12979-2)

Publication date

2014

Document Version

Final published version

Published in

Multidisciplinary Information Retrieval

[Link to publication](#)

Citation for published version (APA):

de Rijke, M. (2014). Don't hurt them: Learning to rank from historical interaction data: (Keynote). In D. Lamas, & P. Buitelaar (Eds.), *Multidisciplinary Information Retrieval: 7th Information Retrieval Facility Conference, IRFC 2014, Copenhagen, Denmark, November 10-12, 2014: proceedings* (pp. xi). (Lecture Notes in Computer Science; Vol. 8849). Springer. <https://doi.org/10.1007/978-3-319-12979-2>

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.

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.