Unbiased Learning to Rank: Theory and Practice

Qingyao Ai
CICS, UMass Amherst
Amherst, MA, USA
aiqy@cs.umass.edu

Jiaxin Mao
DCST, Tsinghua University
Beijing, China
maojiaxin@gmail.com

Yiqun Liu
DCST, Tsinghua University
Beijing, China
yiqunliu@tsinghua.edu.cn

W. Bruce Croft
CICS, UMass Amherst
Amherst, MA, USA
croft@cs.umass.edu

ABSTRACT

Implicit user feedback (such as clicks and dwell time) is an important source of data for modern search engines. While heavily biased [10, 11, 13, 27], it is cheap to collect and particularly useful for user-centric retrieval applications such as search ranking and query recommendation. Understanding the bias inherent in current systems and designing learning to rank algorithms that can effectively learn from implicit user feedback without bias is an important research direction that can significantly improve the quality of modern search engines. To develop such an *unbiased learning-to-rank* (ULTR) system, previous studies have focused on constructing probabilistic graphical models (e.g., click models) with user behavior hypothesis to extract and train ranking systems with unbiased relevance signals. Recently, a novel counterfactual learning framework that estimates and adopts examination propensity for unbiased learning to rank has attracted much attention, both in academia and industry. Despite its popularity, there is no systematic comparison and analysis of the unbiased learning-to-rank frameworks based on graphical models and counterfactual learning. In this tutorial, we provide an overview of the fundamental mechanism and algorithms for unbiased learning to rank. We describe and analyze the theory behind each learning framework, and give detailed instructions on how to conduct unbiased learning to rank in practice.

KEYWORDS

click model, counterfactual learning, unbiased learning to rank

ACM Reference Format:
Qingyao Ai, Jiaxin Mao, Yiqun Liu, and W. Bruce Croft. 2018. Unbiased Learning to Rank: Theory and Practice. In *Proceedings of 2018 ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR'18).* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3234944.3234980

1 INTRODUCTION

Machine learning techniques for Information Retrieval (IR) have become widely used in both academic research and commercial search engines. Although there have been studies that use unsupervised data or pseudo supervision for learning-to-rank models [2, 6], the best retrieval system is typically constructed based on supervised learning. Many of the state-of-the-art retrieval systems today make use of deep models [9, 15], which require large amounts of labeled data. Despite the development of crowdsourcing systems [7], obtaining large-scale and high quality human annotations (e.g. TREC-style relevance judgments) is still expensive, if not impossible. Therefore, implicit feedback such as clicks are still the most attractive data source for the training of ranking systems.

Directly training a ranking model to optimize click data, however, is infeasible because click data are heavily biased [10, 11, 13, 27]. For example, the order of documents in a search engine result page (SERP) has a strong influence on where users click [10]. Studies of position bias show that users tend to examine and click results on the top of a SERP while ignoring those on the bottom. A naive method that treats click/non-click signals as positive/negative feedback will lead to a ranking model that optimizes the order of a search result page but not the relevance of documents.

To leverage the full power of click data for learning to rank, IR researchers have attempted to remove the effect of user bias in the training of ranking models. One such effort is the development of click models. Click models [5, 8, 21, 22] make hypotheses about user browsing behaviors and estimate the true (unbiased) relevance feedback by optimizing the likelihood of the observed user clicks. Ranking models are then trained with the estimated relevance signals so that the overall system is unbiased [14]. Another effort to debias click data is result interleaving [4, 16, 18–20, 26]. By collecting clicks on swapped results, we can obtain unbiased pair preferences for documents in the same result list. These pairwise preference signals can then be used to train learning-to-rank models in an online manner.

More recently, a new research direction has emerged that focuses on directly training ranking models with biased click data using counterfactual learning [12, 23, 24]. This unbiased learning-to-rank framework treats click bias as a counterfactual effect and debiases user feedback by weighting each click with their Inverse Propensity Weights [17]. It uses a propensity model to quantify click biases and does not explicitly estimate the query-document relevance with training data. As theoretically proven by Joachims et al. [12], given the correct bias estimation, ranking models trained with click data under this framework will converge to the same model trained with true relevance signals.

ICTIR'18, September 14–17, 2018, Tianjin, China

ACM ISBN 978-1-4503-5656-5/18/09.
https://doi.org/10.1145/3234944.3234980

In this tutorial, we introduce the theory behind each technique and talk about their differences in detail. We provide hands-on instructions on how to customize and apply these unbiased learning-to-rank techniques to different retrieval tasks. The goal is to help more people understand the concepts of unbiased learning to rank and the existing methods in the field.

2 FORMAT AND SCHEDULE

This tutorial consists of a series of talks on different unbiased learning-to-rank techniques and their applications. Specifically, we will start from describing the concept of user feedback in search scenarios and introducing the definition of examination bias. We will discuss the use of implicit feedback (e.g. clicks and dwell times) in the design of real-word systems and briefly go through several cases where user bias (e.g. position bias [10], trust bias [11], etc.) can affect the performance of retrieval models. These motivate the studies of unbiased learning to rank and explain why they are important for the design of modern retrieval systems.

The technical content of this tutorial will be organized in two parts. In the first part, we will talk about previous studies on click models [5, 8, 21, 22, 25]. The idea of click models is to extract unbiased relevance signals from biased user feedback. They construct hypotheses on user behaviors and build machine learning models (e.g. probabilistic graphic models) to debias user feedback so that we can train a learning-to-rank algorithm using unbiased relevance signals. In this topic, we will first introduce the concept of examination hypothesis and representative user behavior models. Then, we will describe how to derive a click model based on each examination hypothesis and how to estimate the unbiased relevance signals step by step. Finally, we will empirically compare different click models in a joint retrieval framework and discuss the advantages and limitations of unbiased learning to rank with click models. In the second part, we will talk about recent developments on counterfactual learning for unbiased learning to rank [12, 23, 24]. In contrast to click models, unbiased learning to rank with counterfactual learning focuses on the estimation of user examination propensity and uses an inverse perpensity weighting schema to create a learning framework in which a ranking model trained with biased user feedback can converge to the same model trained with unbiased relevance signals. In this topic, we will first introduce the idea of counterfactual learning and its underlining theory. After that, we will describe how to build an unbiased learning-to-rank framework with inverse propensity weighting and how to estimate examination propensity in online systems.

At the end of this tutorial, we will discuss the connections and differences between existing unbiased learning-to-rank techniques. We will show some demos and introduce some cutting-edge work in the field [1, 3].

3 SUPPLEMENTAL MATERIALS

The supplemental materials of this tutorial can be found on homepage of the authors[1].

[1]http://www.cs.umass.edu/~aiqy/

REFERENCES

[1] Qingyao Ai, Keping Bi, Cheng Luo, Jiafeng Guo, and W. Bruce Croft. 2018. Unbiased Learning to Rank with Unbiased Propensity Estimation. In *Proceedings of the 41th ACM SIGIR*. ACM.

[2] Qingyao Ai, Liu Yang, Jiafeng Guo, and W Bruce Croft. 2016. Analysis of the paragraph vector model for information retrieval. In *Proceedings of the 2rd ACM ICTIR*. ACM, 133–142.

[3] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. In *Proceedings of the 41st ACM SIGIR*. ACM.

[4] Olivier Chapelle, Thorsten Joachims, Filip Radlinski, and Yisong Yue. 2012. Large-scale validation and analysis of interleaved search evaluation. *ACM Transactions on Information Systems* 30, 1 (2012), 6.

[5] Olivier Chapelle and Ya Zhang. 2009. A dynamic bayesian network click model for web search ranking. In *Proceedings of the 18th WWW*. ACM, 1–10.

[6] Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn, Jaap Kamps, and W. Bruce Croft. 2017. Neural Ranking Models with Weak Supervision. In *Proceedings of the 40th ACM SIGIR (SIGIR '17)*. ACM, 65–74.

[7] Anhai Doan, Raghu Ramakrishnan, and Alon Y Halevy. 2011. Crowdsourcing systems on the world-wide web. *Commun. ACM* 54, 4 (2011), 86–96.

[8] Georges E Dupret and Benjamin Piwowarski. 2008. A user browsing model to predict search engine click data from past observations.. In *Proceedings of the 31st ACM SIGIR*. ACM, 331–338.

[9] Jiafeng Guo, Yixing Fan, Qingyao Ai, and W Bruce Croft. [n. d.]. A deep relevance matching model for ad-hoc retrieval. In *Proceedings of the 25th ACM CIKM*.

[10] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2005. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the 28th annual ACM SIGIR*. Acm, 154–161.

[11] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, Filip Radlinski, and Geri Gay. 2007. Evaluating the accuracy of implicit feedback from clicks and query reformulations in web search. *ACM Transactions on Information Systems (TOIS)* 25, 2 (2007), 7.

[12] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2017. Unbiased learning-to-rank with biased feedback. In *Proceedings of the 10th ACM WSDM*. ACM, 781–789.

[13] Mark T Keane and Maeve O'Brien. 2006. Modeling Result-List Searching in the World Wide Web: The Role of Relevance Topologies and Trust Bias. In *Proceedings of the Cognitive Science Society*, Vol. 28.

[14] Cheng Luo, Yukun Zheng, Jiaxin Mao, Yiqun Liu, Min Zhang, and Shaoping Ma. 2017. Training deep ranking model with weak relevance labels. In *Australasian Database Conference*. Springer, 205–216.

[15] Bhaskar Mitra, Fernando Diaz, and Nick Craswell. 2017. Learning to Match Using Local and Distributed Representations of Text for Web Search. In *Proceedings of the 26th WWW (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 1291–1299. https://doi.org/10.1145/3038912.3052579

[16] Karthik Raman and Thorsten Joachims. 2013. Learning socially optimal information systems from egoistic users. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 128–144.

[17] Paul R Rosenbaum and Donald B Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70, 1 (1983), 41–55.

[18] Anne Schuth, Harrie Oosterhuis, Shimon Whiteson, and Maarten de Rijke. 2016. Multileave gradient descent for fast online learning to rank. In *Proceedings of the 9th ACM WSDM*. ACM, 457–466.

[19] Adith Swaminathan and Thorsten Joachims. 2015. Batch learning from logged bandit feedback through counterfactual risk minimization. *Journal of Machine Learning Research* 16 (2015), 1731–1755.

[20] Adith Swaminathan and Thorsten Joachims. 2015. Counterfactual risk minimization: Learning from logged bandit feedback. In *ICML*. 814–823.

[21] Chao Wang, Yiqun Liu, Meng Wang, Ke Zhou, Jian-yun Nie, and Shaoping Ma. 2015. Incorporating non-sequential behavior into click models. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 283–292.

[22] Hongning Wang, ChengXiang Zhai, Anlei Dong, and Yi Chang. 2013. Content-aware click modeling. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 1365–1376.

[23] Xuanhui Wang, Michael Bendersky, Donald Metzler, and Marc Najork. 2016. Learning to rank with selection bias in personal search. In *Proceedings of the 39th ACM SIGIR*. ACM, 115–124.

[24] Xuanhui Wang, Nadav Golbandi, Michael Bendersky, Donald Metzler, and Marc Najork. 2018. Position Bias Estimation for Unbiased Learning to Rank in Personal Search. In *Proceedings of the 11th ACM WSDM (WSDM '18)*. ACM, New York, NY, USA, 610–618. https://doi.org/10.1145/3159652.3159732

[25] Wanhong Xu, Eren Manavoglu, and Erick Cantu-Paz. 2010. Temporal click model for sponsored search. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*. ACM, 106–113.

[26] Yisong Yue and Thorsten Joachims. [n. d.]. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th ICML*.

[27] Yisong Yue, Rajan Patel, and Hein Roehrig. 2010. Beyond position bias: Examining result attractiveness as a source of presentation bias in clickthrough data. In *Proceedings of the 19th WWW*. ACM, 1011–1018.