

Learning to Rank in Theory and Practice

From Gradient Boosting to Neural Networks and Unbiased Learning

Claudio Lucchese
Franco Maria Nardini
Ca' Foscari University, Venice, Italy,
ISTI-CNR, Pisa, Italy
claudio.lucchese@unive.it
f.nardini@isti.cnr.it

Rama Kumar Pasumarthi
Sebastian Bruch
Michael Bendersky
Xuanhui Wang
Google Research
ramakumar,bruch,bemike,
xuanhui@google.com

Harrie Oosterhuis
Rolf Jagerman
Maarten de Rijke
University of Amsterdam
oosterhuis,rolf.jagerman,
derijke@uva.nl

ABSTRACT

This tutorial aims to weave together diverse strands of modern Learning to Rank (LtR) research, and present them in a unified full-day tutorial. First, we will introduce the fundamentals of LtR, and an overview of its various sub-fields. Then, we will discuss some recent advances in gradient boosting methods such as LambdaMART by focusing on their efficiency/effectiveness trade-offs and optimizations. Subsequently, we will then present TF-Ranking, a new open source TensorFlow package for neural LtR models, and how it can be used for modeling sparse textual features. Finally, we will conclude the tutorial by covering unbiased LtR – a new research field aiming at learning from biased implicit user feedback.

The tutorial will consist of three two-hour sessions, each focusing on one of the topics described above. It will provide a mix of theoretical and hands-on sessions, and should benefit both academics interested in learning more about the current state-of-the-art in LtR, as well as practitioners who want to use LtR techniques in their applications.

CCS CONCEPTS

• Information systems → Learning to rank.

KEYWORDS

Learning to rank; Efficiency/effectiveness trade-offs; Deep learning; Unbiased learning

ACM Reference Format:

Claudio Lucchese, Franco Maria Nardini, Rama Kumar Pasumarthi, Sebastian Bruch, Michael Bendersky, Xuanhui Wang, Harrie Oosterhuis, Rolf Jagerman and Maarten de Rijke. 2019. Learning to Rank in Theory and Practice: From Gradient Boosting to Neural Networks and Unbiased Learning. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '19)*, July 21–25, 2019, Paris, France. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3331184.3334824>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR '19, July 21–25, 2019, Paris, France

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6172-9/19/07.

<https://doi.org/10.1145/3331184.3334824>

1 OVERVIEW

This full-day tutorial is organized in three sessions lasting two hours each. Together these sessions provide a wide overview of recent advances in the field of Learning to Rank (LtR).

Session I: Efficiency/Effectiveness Trade-offs

We propose an analysis of the efficiency/effectiveness trade-offs in Learning to Rank. In the last years, LtR, had a significant influence in the Information Retrieval field, with large research efforts coming both from the academia and the industry. Indeed, efficiency requirements must be fulfilled in order to make an effective research product deployable within an industrial environment. The evaluation of a model can be too expensive due to its size, the features used and several other factors.

This session discusses the recent solutions that allow to build an effective ranking model that satisfies temporal budget constraints at evaluation time. We first introduce LtR solutions for a multi-stage ranking pipeline with a focus on decision tree ensembles. Then we present several complementary strategies for optimizing the efficiency of a *ranking forest* including: feature analysis [19], tree pruning [9], effectiveness optimization at training time [16], approximate computation [3] and efficient traversal [5].

This session will be presented by Claudio Lucchese from the Ca' Foscari University of Venice and Franco Maria Nardini from the National Research Council of Italy.

Session II: Neural Learning to Rank using TensorFlow

A number of open source packages harnessing the power of deep learning have emerged in recent years and are under active development, including TensorFlow [1], PyTorch [13], Caffe [7], and MXNet [4]. Supervised learning is one of the main use cases of deep learning packages. For example, one task in the ImageNet competitions [15] is to predict image categories, which can be formulated as a multi-class classification problem. However, compared with the comprehensive support for classification or regression in open-source deep learning packages, there is a paucity of support for ranking problems.

To address this gap, we developed TensorFlow Ranking¹: an open-source library for training large-scale LtR models using deep

¹<https://github.com/tensorflow/ranking>

learning in TensorFlow [12]. The library is flexible and highly configurable: it provides an easy-to-use API to support different scoring mechanisms, loss functions, example weights, and evaluation metrics. In this hands-on tutorial, we aim to cover how TensorFlow Ranking can be effectively employed in a variety of learning-to-rank scenarios.

First, we will present a brief overview of neural LtR, TensorFlow and Estimator frameworks. Then, we will introduce Tensorflow Ranking components and APIs, and demonstrate how it can handle advanced losses, scoring functions and sparse textual features. Finally, we will provide hands-on codelabs using two existing LtR datasets: MSLR-Web30k [14] and MS MARCO [10].

This session will be presented by Rama Kumar Pasumarthi, Sebastian Bruch, Michael Bendersky and Xuanhui Wang from Google Research.

Session III: Unbiased Learning to Rank

User interactions provide great potential for LtR: they give valuable implicit feedback and are easy to obtain in large amounts. However, user interactions contain biases such as position bias: documents displayed at higher ranks receive more attention. Naively learning from interactions while ignoring such biases can lead to detrimental performance. Consequently, the field of Unbiased LtR aims to learn the true user preferences from their interactions, thus avoiding the effect of biases. In this part we will cover and contrast the two main approaches to Unbiased LtR: *Counterfactual* LtR and *Online* LtR.

Counterfactual LtR [8, 17] uses an explicit model position bias, and through an inverse propensity weighing approach optimizes LtR metrics without bias. In addition to the learning method, we will discuss how models of position bias can be inferred [2, 18] and other practical considerations.

Online LtR [20] methods directly interact with users, and perform randomizations allowing them to deal with several biases. We will discuss the important *Dueling Bandit* approach [20], as well as the recent *Pairwise* approach [11].

Finally, we compare and contrast both approaches: on a theoretical level and by looking at empirical comparisons [6]. We discuss the situations for which each approach was designed, and the places where they are applicable. This helps LtR practitioners to choose between the two approaches.

This third session will be presented by Harrie Oosterhuis, Rolf Jagerman, and Maarten de Rijke from the University of Amsterdam.

2 SUPPORTING MATERIALS

You can find more materials related to this tutorial on our website <http://ltr-tutorial-sigir19.isti.cnr.it/>.

ACKNOWLEDGMENTS

The University of Amsterdam team was partially supported by Ahold Delhaize, the Association of Universities in the Netherlands (VSNU), the Innovation Center for Artificial Intelligence (ICAI), and the Netherlands Organisation for Scientific Research (NWO) under project nr. 612.001.551. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, and others. 2016. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation*. 265–283.
- [2] Aman Agarwal, Ivan Zaitsev, and Thorsten Joachims. 2018. Consistent position bias estimation without online interventions for learning-to-rank. *arXiv preprint arXiv:1806.03555* (2018).
- [3] B. Barla Cambazoglu, Hugo Zaragoza, Olivier Chapelle, Jiang Chen, Ciya Liao, Zhaohui Zheng, and Jon Degenhardt. 2010. Early exit optimizations for additive machine learned ranking systems. In *3rd ACM International Conference on Web Search and Data Mining*. ACM, 411–420.
- [4] Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. 2015. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *Neural Information Processing Systems, Workshop on Machine Learning Systems* (2015).
- [5] Domenico Dato, Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, Nicola Tonellotto, and Rossano Venturini. 2016. Fast ranking with additive ensembles of oblivious and non-oblivious regression trees. *ACM Transactions on Information Systems* 35, 2 (2016), Article 15.
- [6] Rolf Jagerman, Harrie Oosterhuis, and Maarten de Rijke. 2019. To model or to intervene: A comparison of counterfactual and online learning to rank from user interactions. In *42nd International ACM SIGIR Conference on Research & Development in Information Retrieval*. ACM, (to appear).
- [7] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. In *22nd ACM International Conference on Multimedia*. ACM, 675–678.
- [8] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2017. Unbiased learning-to-rank with biased feedback. In *10th ACM International Conference on Web Search and Data Mining*. ACM, 781–789.
- [9] Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, Fabrizio Silvestri, and Salvatore Trani. 2018. X-CLEaVER: Learning ranking ensembles by growing and pruning trees. *ACM Transactions on Intelligent Systems and Technology* 9, 6 (2018), Article 62.
- [10] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268* (2016).
- [11] Harrie Oosterhuis and Maarten de Rijke. 2018. Differentiable unbiased online learning to rank. In *27th ACM International Conference on Information and Knowledge Management*. ACM, 1293–1302.
- [12] Rama Kumar Pasumarthi, Sebastian Bruch, Xuanhui Wang, Cheng Li, Michael Bendersky, Marc Najork, Jan Pfeifer, Nadav Golbandi, Rohan Anil, and Stephan Wolf. 2019. TF-Ranking: Scalable TensorFlow library for learning-to-rank. In *25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, (to appear).
- [13] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. In *Advances in Neural Information Processing Systems, AutoDiff Workshop: The Future of Gradient-Based Machine Learning Software and Techniques*.
- [14] Tao Qin and Tie-Yan Liu. 2013. Introducing LETOR 4.0 Datasets. *arXiv preprint arXiv:1306.2597* (2013).
- [15] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, and others. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
- [16] Lidan Wang, Jimmy J. Lin, and Donald Metzler. 2010. Learning to efficiently rank. In *33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 138–145.
- [17] Xuanhui Wang, Michael Bendersky, Donald Metzler, and Marc Najork. 2016. Learning to rank with selection bias in personal search. In *41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. ACM, 115–124.
- [18] Xuanhui Wang, Nadav Golbandi, Michael Bendersky, Donald Metzler, and Marc Najork. 2018. Position bias estimation for unbiased learning to rank in personal search. In *11th ACM International Conference on Web Search and Data Mining*. ACM, 610–618.
- [19] Zhixiang Xu, Olivier Chapelle, and Kilian Q Weinberger. 2012. The greedy miser: Learning under test-time budgets. In *29th International Conference on Machine Learning*. 1175–1182.
- [20] Yisong Yue and Thorsten Joachims. 2009. Interactively optimizing information retrieval systems as a dueling bandits problem. In *26th Annual International Conference on Machine Learning*. ACM, 1201–1208.