Advances in Recommender Systems: From Multi-stakeholder Marketplaces to Automated RecSys

Rishabh Mehrotra[1], Ben Carterette[2]
Spotify Research
[1]London, UK; [2]New York, US
{rishabhm,benjaminc}@spotify.com

Yong Li[3], Quanming Yao[4], Chen Gao[3]
[3]Tsinghua University. [4]4Paragidm Inc.
liyong07@tsinghua.edu.cn,yaoquanming@4paradigm.com,gc16@mails.tsinghua.edu.cn

James Kwok, Qiang Yang
Hong Kong University of Science and Technology
{jamesk,qyang}@cse.ust.hk

Isabelle Guyon
University Paris-Saclay
isabelle@clopinet.com

ABSTRACT

The tutorial focuses on two major themes of recent advances in recommender systems:

Part A: Recommendations in a Marketplace: Multi-sided marketplaces are steadily emerging as valuable ecosystems in many applications (e.g. Amazon, AirBnb, Uber), wherein the platforms have customers not only on the demand side (e.g. users), but also on the supply side (e.g. retailer). This tutorial focuses on designing search & recommendation frameworks that power such multi-stakeholder platforms. We discuss multi-objective ranking/recommendation techniques, discuss different ways in which stakeholders specify their objectives, highlight user specific characteristics (e.g. user receptivity) which could be leveraged when developing joint optimization modules and finally present a number of real world case-studies of such multi-stakeholder platforms.

Part B: Automated Recommendation System: As the recommendation tasks are getting more diverse and the recommending models are growing more complicated, it is increasingly challenging to develop a proper recommendation system that can adapt well to a new recommendation task. In this tutorial, we focus on how automated machine learning (AutoML) techniques can benefit the design and usage of recommendation systems. Specifically, we start from a full scope describing what can be automated for recommendation systems. Then, we elaborate more on three important topics under such a scope, i.e., feature engineering, hyperparameter optimization/neural architecture search, and algorithm selection. The core issues and recent works under these topics will be introduced, summarized, and discussed. Finally, we finalize the tutorial with conclusions and some future directions.

ACM Reference Format:
Rishabh Mehrotra[1], Ben Carterette[2], Yong Li[3], Quanming Yao[4], Chen Gao[3], James Kwok, Qiang Yang, and Isabelle Guyon. 2020. Advances in Recommender Systems: From Multi-stakeholder Marketplaces to Automated RecSys. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '20), August 23–27, 2020, Virtual Event, CA, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3394486.3406463

ACM ISBN 978-1-4503-7998-4/20/08.
https://doi.org/10.1145/3394486.3406463

PART A: RECOMMENDATIONS IN A MARKETPLACE

Multi-sided marketplaces involve interaction between multiple stakeholders among which there are different individuals with assorted needs. While traditional recommender systems focused specifically towards increasing consumer satisfaction by providing relevant content to the consumers, multi-sided marketplaces face an interesting problem of optimizing for multiple stakeholder objectives [3, 4]. Part A of the tutorial consider research problems which need to be addressed when developing a recommendation framework powering a multi-stakeholder marketplace.

Outline of the tutorial

We begin by contrasting traditional recommendations systems with those needed for marketplaces, and identify four key research areas which need to be addressed. First, we discuss algorithmic techniques for multi-objective rankings & recommendations [5] to jointly optimize the different objectives. Second, we discuss different ways in which stakeholders specify their objectives. Third, we discuss user & content specific characteristics which could be leveraged while jointly optimizing such models. Furthermore, we discuss evaluation of such systems and present numerous case industrial studies.

(1) Introduction to Marketplace
(2) Phase I: Multi-Objective Ranking
(3) Phase II: Optimization Objectives
(4) Phase III: Leveraging User & Content Understanding
(5) Phase IV: Multi-sided evaluation
(6) Phase V: Open Research problems

The tutorial is aimed at introducing practitioners to the methods that can be used to develop multi-stakeholder search and recommender systems. The main focus of the tutorial is (i) components which constitute a multi-stakeholder system, and (ii) imparting knowledge on how to design and develop each of these components in a scalable way. The tutorial builds upon recent RecSys 2019 tutorial [3] on similar topic, and presents additional multi-objective algorithms and recent case studies.

Presenter Bio

Rishabh Mehrotra is a Senior Research Scientist at Spotify Research in London. He obtained his PhD in the field of Machine Learning and Information Retrieval from University College London where he was partially supported by a Google Research Award.

His current research focuses on bandit based recommendations, counterfactual analysis and experimentation. Some of his recent work has been published at top conferences including WWW, KDD, SIGIR, NAACL, CIKM, RecSys and WSDM.

Ben Carterette is a Senior Research Manager at Spotify and an Associate Professor of Computer and Information Sciences at the University of Delaware. His research focuses on evaluation in Information Retrieval, including test collection construction, and statistical testing. He has published over 100 papers in venues such as ACM TOIS, SIGIR, CIKM, WSDM, ECIR, and ICTIR, winning three Best Paper Awards for his work. Dr Carterette was elected Chair of ACM SIGIR in 2019.

PART B: AUTOMATED RECOMMENDATION SYSTEMS

With diverse task settings and complicated models, it is a severe problem to design recommender systems that can adapt well to new tasks. Recently, automated machine learning [7], which targets at easing the usage of machine learning tools and designing task-dependent learning models, has become an important and popular area with both practical needs and research values. In this tutorial, we discuss leveraging AutoML to help solve the problem.

Outline of the tutorial

The outline of the tutorial is as follows. After introducing the background and preliminaries for AutoML and RecSys, we will discuss using AutoML to design and improve recommendation models from four aspects: model design [6], hyper-parameter optimization [1], feature engineering [2], and exploitation of rich side information. Lastly, we will summarize the tutorial and discuss future directions.

(1) Introduction to AutoML
(2) Phase I: Automated Model Design / Neural Architecture Search
    - Efficient Neural Interaction Functions Search
    - Automated Model Search for Collaborative Filtering
(3) Phase II: Hyper-parameter Optimization for Recommendation
    - Regularization Automatic: Framework and Method
    - Embedding Size Automatic: Methods and Directions
    - Learning Rate and Other Parameters Optimization
(4) Phase III: Feature Engineering for Recommendation
    - AutoCross: Automatic Feature Crossing for Tabular Data
    - AutoFM: Automatic Feature Selection for FM
(5) Phase IV: Automated Exploitation of Rich Side Information
    - Automated Knowledge Graph for Recommendation
    - Automated Graph Neural Networks for Recommendation
(6) Phase V: Conclusion and open discussions

There were two related tutorials: Alexandros Karatzoglou and Balázs Hidasi, Deep Learning for Recommender Systems, at RecSys 2017; Parashar Shah and Krishna Anumalasetty, Democratizing & Accelerating AI through Automated Machine Learning, at KDD 2019. This tutorial is significantly different from them as it focuses on how to leverage automated machine learning in recommender system.

Presenter Bio

Yong Li is currently a Tenured Associate Professor of the Department of Electronic Engineering, Tsinghua University. He received the Ph.D. degree in electronic engineering from Tsinghua University in 2012. His research interests include machine learning and big data mining, particularly, automatic machine learning and spatial-temporal data mining for urban computing, recommender systems, and knowledge graphs. He served/is serving as SPC or PC of major Data Mining and AI conferences, including KDD, WWW, IJCAI, AAAI, SIGIR, and UbiComp. He has published over 100 papers on first-tier international conferences and journals, including KDD, WWW, UbiComp, SIGIR, AAAI, TKDE, TMC, etc.

Quanming Yao is a senior scientist in 4Paradigm (Hong Kong), who has established and currently is the leader of the company’s machine learning research team. He obtained his Ph.D. degree at the Department of Computer Science and Engineering of Hong Kong University of Science and Technology (HKUST). His research interests are in machine learning, optimization, and automated machine learning. He has 30 top-tier journal and conference papers, including ICML, NeurIPS, JMLR and TPAMI.

Chen Gao is currently a Ph.D. candidate in the Department of Electronic Engineering, Tsinghua University. His research focuses on recommender systems and data mining.

James T. Kwok is a Professor in the Department of Computer Science and Engineering, Hong Kong University of Science and Technology. He is an IEEE Fellow. He received his B.Sc. degree in Electrical and Electronic Engineering from the University of Hong Kong and his Ph.D. degree in computer science from the Hong Kong University of Science and Technology.

Isabelle Guyon is chaired professor in big data at the University Paris-Saclay, specialized in statistical data analysis, pattern recognition and machine learning. She is one of the co-founders of the ChaLearn Looking at People (LAP) challenge series and she pioneered applications of the Microsoft Kinect to gesture recognition. Her areas of expertise include computer vision and and bioinformatics.

Qiang Yang is a New Bright Endowed Chair Professor of Engineering in Computer Science and Engineering Department at Hong Kong University of Science and Technology (HKUST). His research interests are artificial intelligence including machine learning and data mining. He is a fellow of AAAI, IEEE, IAPR and AAAS.

REFERENCES

[1] Yihong Chen, Bei Chen, Xiangnan He, Chen Gao, Yong Li, Jian-Guang Lou, and Yue Wang. 2019. λOpt: Learn to Regularize Recommender Models in Finer Levels. In *KDD*. 978–986.
[2] Yuanfei Luo, Mengshuo Wang, Hao Zhou, Quanming Yao, Wei-Wei Tu, Yuqiang Chen, Wenyuan Dai, and Qiang Yang. 2019. Autocross: Automatic feature crossing for tabular data in real-world applications. In *KDD*. 1936–1945.
[3] Rishabh Mehrotra and Benjamin Carterette. 2019. Recommendations in a marketplace. In *RecSys*. 580–581.
[4] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *CIKM*. 2243–2251.
[5] Rishabh Mehrotra, Niannan Xue, and Mounia Lalmas. 2020. Bandit based Optimization of Multiple Objectives on a Music Streaming Platform. In *KDD*.
[6] Quanming Yao, Xiangning Chen, James T Kwok, Yong Li, and Cho-Jui Hsieh. 2020. Efficient neural interaction function search for collaborative filtering. In *WWW*.
[7] Quanming Yao, Ju Xu, Wei-Wei Tu, and Zhanxing Zhu. 2020. Efficient Neural Architecture Search via Proximal Iterations. In *AAAI*. 6664–6671.