Bias Issues and Solutions in Recommender System

Tutorial on the RecSys 2021

Jiawei Chen
cjwustc@ustc.edu.cn
University of Science and Technology of China
China

Xiang Wang
xiangwang@u.nus.edu
National University of Singapore
Singapore

Fuli Feng
fulifeng93@gmail.com
National University of Singapore
Singapore

Xiangnan He
xiangnanhe@gmail.com
University of Science and Technology of China
China

ABSTRACT

Recommender systems (RS) have demonstrated great success in information seeking. Recent years have witnessed a large number of work on inventing recommendation models to better fit user behavior data. However, user behavior data is observational rather than experimental. This makes various biases widely exist in the data, including but not limited to selection bias, position bias, exposure bias. Blindly fitting the data without considering the inherent biases will result in many serious issues, e.g., the discrepancy between offline evaluation and online metrics, hurting user satisfaction and trust on the recommendation service, etc. To transform the large volume of research models into practical improvements, it is highly urgent to explore the impacts of the biases and develop debiasing strategies when necessary. Therefore, bias issues and solutions in recommender systems have drawn great attention from both academic and industry.

In this tutorial, we aim to provide an systemic review of existing work on this topic. We will introduce six types of biases in recommender system, along with their definitions and characteristics; review existing debiasing solutions, along with their strengths and weaknesses; and identify some open challenges and future directions. We hope this tutorial could stimulate more ideas on this topic and facilitate the development of debiasing recommender systems.

ACM Reference Format:

Jiawei Chen, Xiang Wang, Fuli Feng, and Xiangnan He. 2021. Bias Issues and Solutions in Recommender System: Tutorial on the RecSys 2021. In *Fifthteenth ACM Conference on Recommender Systems (RecSys '21), September 27-October 1, 2021, Amsterdam, Netherlands.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3460231.3473321

RecSys '21, September 27-October 1, 2021, Amsterdam, Netherlands

ACM ISBN 978-1-4503-8458-2/21/09…$15.00
https://doi.org/10.1145/3460231.3473321

1 INTRODUCTION

Serious bias issues in RS. Recent years have witnessed flourishing publications on recommendation, most of which aim at inventing machine learning model to fit user behavior data. However, in real-world scenarios, RS usually faces many bias issues which are challenging to handle and may deteriorate the recommendation effectiveness. Blindly fitting the recommendation model without considering the inherent biases will result in skewed results. It can be seen from the following factors:

- User behavior data, which lays the foundation for recommendation model training, is observational rather than experimental. The main reason is that a user generates behaviors on the basis of the exposed items, making the observational data confounded by the exposure mechanism of the system and the self-selection of the user. Blindly fitting the model with biased data would capture skewed user preference.
- Items are not evenly presented in the data, e.g., some items are more popular than others and thus receive more user behaviors. As a result, these popular items would have a larger impact on the model training, making the recommendations biased towards them. The same situation applies to the user side.
- One nature of RS is the feedback loop — the exposure mechanism of the RS determines user behaviors, which are circled back as the training data for the RS. Such feedback loop not only creates biases but also intensifies biases over time, resulting in "the rich get richer" Matthew effect.

Increasing attention on recommendation bias. Recent years have seen a surge of research effort being devote to explore the impact of biases and correspondingly debaising solutions in RS. Figure 1 shows the number of related papers in top venues increases significantly since the year of 2015. The specific prestigious international conference on information retrieval, SIGIR, has organized specific sessions in 2018 and 2020 to discuss topics on bias elimination[1]. SIGIR even presents the Best Paper award to the paper on this topic in 2018 [3] and 2020 [24], respectively. Also, there are two tutorials on this topic conducted in the conference UMAP'20 [2] and Recsys'20 [13] respectively. These successful tutorials attract many audiences from the RS communities.

[1]http://www.sigir.org/sigir2020/schedule/; http://sigir.org/sigir2018/program/program-at-a-glance/

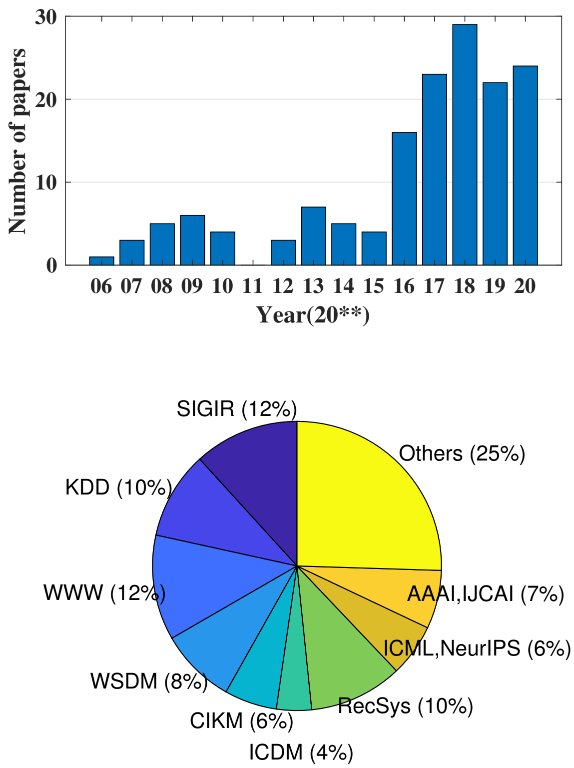

Figure 1: The statistics of publications related to biases in RS with the publication year and venue.

Necessity of this tutorial. Given the importance of debiasing in RS, the increasing attention of this topic, the rapid development of debiasing techniques, and the flourishing recent work, we believe it is the right time to conduct a tutorial of this area, so as to benefit the researchers and practitioners to understand current progress and further work on this topic. Especially we find recent works on biases are rather fragmented — despite the wide usage of the terminology "bias" in the literature, its definition is usually vague and even inconsistent across papers. Our systemic tutorial on recent work could help the beginners to fast step into this area and to keep up with state-of-the-art debiasing technologies in RS. Note that this area is not matured and has many open problems. We also would like to arise discussions on these core problems, with the ambition of inspiring more new idea and facilitating the development of this area.

2 OUTLINE

This tutorial focuses on bias and debias in recommender system. Here we present an outline of the topics to be covered, with timing:

- *Introduction.* (20 Min)
  - Introduction of recommender system.
  - Collaborative filtering fundamentals: data, models and evaluations. [16, 26]

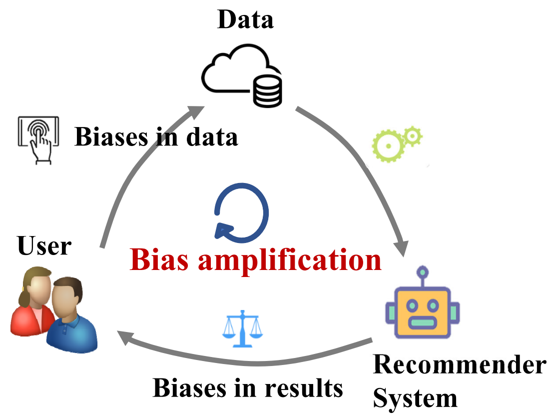

Figure 2: Feedback loop in recommendation, where biases occur in different stages. [5]

  - Lifecycle of RS — Feedback loop, as shown in Figure 2. [22]
  - Ubiquity of biases in RS and the motivation of debiasing.
  - Organization of the tutorial.
- *Biases in Recommendation Data.* (50 Min):
  - Definition of data bias.
  - Categories: including selection bias [23], conformity bias [21], exposure bias [19] and position bias [11].
  - Recent solutions for data biases including: propensity score [4, 28], data imputation [29], sampling [9], generative model [12], etc.
  - Open problems and future directions on data bias.
- *Biases in Recommendation Results.* (100 Min)
  - Popularity bias: definition, characteristic, solutions and future directions[1, 10, 30, 35].
  - Unfairness: definition, characteristic, solutions and future directions [14, 18, 25, 27, 31, 32].
- *Bias Amplification in Loop and its Solutions.* (10 Min)
  - Bias Amplification over time along the loop [22].
  - Solutions: using uniform data [17, 20], reinforcement learning [34].

The content of this tutorial is based on our recent survey paper "Bias and Debias in Recommender System: A Survey and Future Directions" [5].

3 INTENDED AUDIENCE AND LEVEL

This tutorial is intended for the following researchers and practitioners in RS:

- who are new to the bias issues and look for a tutorial to fast step into this area;
- who are confused by different bias definitions in the literature and need a tutorial to understand the biases;
- who face bias issues in building recommender systems and look for suitable solutions.

Only elementary knowledge on RS and basic linear algebra are required. Attendees are expected to gain a global picture of this area, high-level understanding of six different types of biases, state-of-the-art debiasing strategies and some promising future directions.

4 QUALIFICATION OF PRESENTERS

Our team has rich research experience on this field. We have published over 60 papers on recommendation, and over 15 papers including a survey papers on the topics of recommendation debiasing. Specifically, our paper on popularity bias won the Best Paper Honorable Mention on SIGIR 2021. We also gave a tutorial with the same topic on WWW 2021[2].

4.1 Bios of Presenters

Jiawei Chen is a Postdoc Research Fellow in School of Information Science and Technology, University of Science and Technology of China. His research interests include information retrieval, data mining, and causal reasoning. He received Ph.D. in Computer Science from Zhejiang University in 2020. He has published over 10 academic papers on international conferences such as WWW, SIGIR, AAAI, CIKM, KDD and ICDM.

Xiangnan He is a professor at the University of Science and Technology of China (USTC). His research interests span information retrieval, data mining, and multi-media analytics. He has over 90 publications in top conferences such as SIGIR, WWW, and MM, KDD, and journals including TKDE, TOIS, and TMM.

Xiang Wang is now a research fellow at National University of Singapore. He received his Ph.D. degree from National University of Singapore in 2019. His research interests include recommender systems, graph learning, and explainable deep learning techniques. He has published some academic papers on international conferences such as KDD, WWW, SIGIR, and AAAI.

Fuli Feng is a Research Fellow in the School of Computing, National University of Singapore (NUS). He received Ph.D. in Computer Science from NUS in 2019. His research interests include information retrieval, data mining, and multi-media processing. He has over 30 publications appeared in several top conferences such as SIGIR, WWW, and MM, and journals including TKDE and TOIS.

4.2 Relevant Publications by presenters

- A survey of recommendation debiasing [5].
- Basic recommendation models, *e.g.,* NCF[16], LightGCN[15].
- Debiasing strategies, *e.g.,* [7] for selection bias, [6, 8, 9] for exposure bias, [4] for debiasing framework, [30, 33, 35] for popularity bias.

REFERENCES

[1] Himan Abdollahpouri and Masoud Mansoury. 2020. Multi-sided exposure bias in recommendation. *arXiv preprint arXiv:2006.15772* (2020).
[2] Ludovico Boratto and Mirko Marras. 2020. Hands on Data and Algorithmic Bias in Recommender Systems. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*. 388–389.
[3] Rocío Cañamares and Pablo Castells. 2018. Should i follow the crowd?: A probabilistic analysis of the effectiveness of popularity in recommender systems. In *SIGIR*. ACM, 415–424.
[4] Jiawei Chen, Hande Dong, Yang Qiu, Xiangnan He, Xin Xin, Liang Chen, Guli Lin, and Keping Yang. [n.d.]. AutoDebias: Learning to Debias for Recommendation.
[5] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2020. Bias and Debias in Recommender System: A Survey and Future Directions. *arXiv preprint arXiv:2010.03240* (2020).

[2]https://lds4bias.github.io/

[6] Jiawei Chen, Yan Feng, Martin Ester, Sheng Zhou, Chun Chen, and Can Wang. 2018. Modeling Users' Exposure with Social Knowledge Influence and Consumption Influence for Recommendation. In *CIKM*. 953–962.
[7] Jiawei Chen, Can Wang, Martin Ester, Qihao Shi, Yan Feng, and Chun Chen. 2018. Social recommendation with missing not at random data. In *ICDM*. IEEE, 29–38.
[8] Jiawei Chen, Can Wang, Sheng Zhou, Qihao Shi, Jingbang Chen, Yan Feng, and Chun Chen. 2020. Fast Adaptively Weighted Matrix Factorization for Recommendation with Implicit Feedback.. In *AAAI*. 3470–3477.
[9] Jiawei Chen, Can Wang, Sheng Zhou, Qihao Shi, Yan Feng, and Chun Chen. 2019. SamWalker: Social Recommendation with Informative Sampling Strategy. In *WWW*. ACM, 228–239.
[10] Zhihong Chen, Rong Xiao, Chenliang Li, Gangfeng Ye, Haochuan Sun, and Hongbo Deng. 2020. ESAM: Discriminative Domain Adaptation with Non-Displayed Items to Improve Long-Tail Performance. In *SIGIR*. 579–588.
[11] Andrew Collins, Dominika Tkaczyk, Akiko Aizawa, and Joeran Beel. 2018. A Study of Position Bias in Digital Library Recommender Systems. *arXiv preprint arXiv:1802.06565* (2018).
[12] Georges E Dupret and Benjamin Piwowarski. 2008. A user browsing model to predict search engine click data from past observations.. In *SIGIR*. 331–338.
[13] Ruoyuan Gao and Chirag Shah. 2020. Counteracting Bias and Increasing Fairness in Search and Recommender Systems. In *RecSys*. 745–747.
[14] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search. In *KDD*. 2221–2231.
[15] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *SIGIR*. 639–648.
[16] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *WWW*. ACM, 173–182.
[17] Ray Jiang, Silvia Chiappa, Tor Lattimore, András György, and Pushmeet Kohli. 2019. Degenerate feedback loops in recommender systems. In *AIES*. 383–390.
[18] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2012. Enhancement of the Neutrality in Recommendation. In *RecSys*, Vol. 893. 8–14.
[19] Dawen Liang, Laurent Charlin, James McInerney, and David M Blei. 2016. Modeling user exposure in recommendation. In *WWW*. ACM, 951–961.
[20] Dugang Liu, Pengxiang Cheng, Zhenhua Dong, Xiuqiang He, Weike Pan, and Zhong Ming. 2020. A General Knowledge Distillation Framework for Counterfactual Recommendation via Uniform Data. In *SIGIR*. 831–840.
[21] Yiming Liu, Xuezhi Cao, and Yong Yu. 2016. Are You Influenced by Others When Rating?: Improve Rating Prediction by Conformity Modeling. In *RecSys*. ACM, 269–272.
[22] Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. Feedback Loop and Bias Amplification in Recommender Systems. *arXiv preprint arXiv:2007.13019* (2020).
[23] Benjamin M Marlin, Richard S Zemel, Sam Roweis, and Malcolm Slaney. 2007. Collaborative filtering and the missing at random assumption. In *UAI*. 267–275.
[24] Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. 2020. Controlling Fairness and Bias in Dynamic Learning-to-Rank. In *SIGIR*. ACM, 429–438.
[25] Razieh Nabi and Ilya Shpitser. 2018. Fair Inference on Outcomes. In *AAAI*. 1931–1940.
[26] Rong Pan, Yunhong Zhou, Bin Cao, Nathan N. Liu, Rajan Lukose, Martin Scholz, and Qiang Yang. 2008. One-class collaborative filtering. *ICDM* (2008), 502–511.
[27] Dino Pedreschi, Salvatore Ruggieri, and Franco Turini. 2008. Discrimination-aware data mining. In *KDD*. 560–568.
[28] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. 2016. Recommendations as treatments: Debiasing learning and evaluation. In *ICML (JMLR Workshop and Conference Proceedings, Vol. 48)*. 1670–1679.
[29] Yixin Wang, Dawen Liang, Laurent Charlin, and David M Blei. 2018. The deconfounded recommender: A causal inference approach to recommendation. *arXiv preprint arXiv:1808.06581* (2018).
[30] Tianxin Wei, Fuli Feng, Jiawei Chen, Chufeng Shi, Ziwei Wu, Jinfeng Yi, and Xiangnan He. 2020. Model-Agnostic Counterfactual Reasoning for Eliminating Popularity Bias in Recommender System. *arXiv preprint arXiv:2010.15363* (2020).
[31] Richard S. Zemel, Yu Wu, Kevin Swersky, Toniann Pitassi, and Cynthia Dwork. 2013. Learning Fair Representations. In *ICML*, Vol. 28. 325–333.
[32] Junzhe Zhang and Elias Bareinboim. 2018. Fairness in Decision-Making - The Causal Explanation Formula. In *AAAI*. 2037–2045.
[33] Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021. Causal Intervention for Leveraging Popularity Bias in Recommendation. In *SIGIR*.
[34] Xiangyu Zhao, Long Xia, Liang Zhang, Zhuoye Ding, Dawei Yin, and Jiliang Tang. 2018. Deep reinforcement learning for page-wise recommendations. In *RecSys*. 95–103.
[35] Yu Zheng, Chen Gao, Xiang Li, Xiangnan He, Yong Li, and Depeng Jin. 2021. Disentangling User Interest and Conformity for Recommendation with Causal Embedding. In *WWW*. 2980–2991.