

# RMSE: Workshop on Recommendation in Multistakeholder Environments\*

Robin Burke  
Dept. of Information Science  
University of Colorado  
Boulder, Colorado

Himan Abdollahpouri  
Dept. of Information Science  
University of Colorado  
Boulder, Colorado

Edward C. Malthouse  
Spiegel Research Center  
Northwestern University  
Evanston, Illinois

KP Thai  
Squirrel AI Learning  
Shanghai, China

Yongfeng Zhang  
Dept. of Computer Science  
Rutgers University  
New Brunswick, New Jersey

## ABSTRACT

In research practice, recommender systems are typically evaluated on their ability to provide items that satisfy the needs and interests of the end user. However, in many recommendation domains, the user for whom recommendations are generated is not the only stakeholder in the recommendation outcome. For example, fairness and balance across stakeholders is important in some recommendation applications; achieving a goal such as promoting new sellers in a marketplace might be important in others. Such multistakeholder environments present unique challenges for recommender system design and evaluation, and these challenges were the focus of this workshop.

## CCS CONCEPTS

• **Information systems** → *Evaluation of retrieval results*; • **Social and professional topics** → *User characteristics*.

## KEYWORDS

multistakeholder recommendation; fairness; discrimination; bias; e-commerce

## 1 WORKSHOP TOPIC

One of the defining characteristics of recommender systems is personalization. Recommender systems are typically evaluated on their ability to provide items that satisfy the needs and interests of the end user. However, in many recommendation domains, the user for whom recommendations are generated is not the only stakeholder in the recommendation outcome. Other users, the providers of products, and even the system's own objectives may need to be considered when these perspectives differ from those of end users. Fairness and balance are important examples of system-level objectives, and these social-welfare-oriented goals may at times run counter to individual preferences. Sole focus on the end user hampers developers' ability to incorporate such objectives into recommendation algorithms and system designs.

In addition, in many e-commerce retail settings, recommendation is viewed as a form of marketing and, as such, the economic

considerations of the retailer will also enter into the recommendation function. A business may wish to highlight products that are more profitable or that are currently on sale, for example. Commercial recommender systems often use separate "business rules" functionality to integrate such items into the personalized recommendations generated through conventional means. Adding the retailer as a stakeholder allows such considerations to be integrated throughout the recommendation process.

The workshop encouraged submissions addressing the challenges of producing recommendations in multistakeholder settings, including but not limited to the following topics:

- The requirements of different multistakeholder applications such as:
  - Recommendation in multisided platforms
  - Fairness-aware recommendation
  - Multi-objective optimization in Recommendation
  - Value-aware recommendation in commercial settings
  - Reciprocal recommendation
- Algorithms for multistakeholder recommendation including multi-objective optimization, re-ranking and others.
- The evaluation of multistakeholder recommendation systems.
- User experience considerations in multistakeholder recommendation including ethics, transparency, and interfaces for different stakeholders.

The RMSE 2019 workshop is a continuation of the discussion of these topics in prior RecSys workshops including Value-Aware and Multistakeholder Recommendation (VAMS 2017) [3] and Responsible Recommendation (FATRec 2017 and 2018) [6, 7].

## 2 WORKSHOP ORGANIZATION

The program committee for the workshop included (in addition to the organizers listed above):

- Gediminas Adomavicius, University of Minnesota
- James Caverlee, Texas A&M University
- Dietmar Jannach, University of Klagenfurt
- Toshihiro Kamishima, National Institute of Advanced Industrial Science and Technology (AIST)
- Rishabh Mehrotra, Spotify Research
- Nasim Sonboli, University of Colorado Boulder
- Yong Zheng: Illinois Institute of Technology

\*Copyright 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Presented at the RMSE workshop held in conjunction with the 13th ACM Conference on Recommender Systems (RecSys), 2019, in Copenhagen, Denmark.

- Morteza Zihayat, Ryerson University

### 3 WORKSHOP PROGRAM

The workshop schedule included a keynote address by Ed Chi of Google, and three paper sessions featuring 10 papers (6 long, 4 short) in the following areas:

- **Fairness:** defining, evaluating, and implementing different aspects of fairness for recommender systems. In this session, three papers were presented and discussed. In [4], Deldjoo et al. introduced a fairness measure based on the generalized cross entropy to ensure the outcome of the recommendations matches a pre-defined utility for each group. Abdollahpouri and Burke [1] developed a taxonomy of different types of multi-stakeholder recommender systems based on the architecture of the system. They categorized such systems into 1) multi-receiver recommenders, 2) multiprovider recommenders and 3) recommenders with side stakeholders. The authors also showed the close connection between multi-stakeholder recommendation and multi-sided fairness. In [12], Strucks et al. investigated BlurMe, a gender obfuscation technique that has been shown to block classifiers from inferring binary gender from users' profiles and proposed an extension to BlurMe, called BlurM(or)e, that addresses the privacy issues associated with BlurMe.
- **Calibration:** matching recommendation output to user preferences and examining disparities between types of users and types of items. There were four papers presented in this session. Tsintzou et al. [13] proposed a metric called bias disparity to measure the difference between the bias towards different movie genres in user profiles and in recommendations. Some of the other papers in this session were inspired by this work. Mansoury et al. [11] explored how different recommendation algorithms reflect the trade-off between ranking quality and bias disparity. Their experiments included neighborhood-based, model-based, and trust-aware recommendation algorithms. Another paper in this session was a work by Lin et al. [8] where authors examined bias disparity over a range of different algorithms and for different item categories and demonstrate significant differences between model-based and memory-based algorithms. Finally, Abdollahpouri et al. in [2] investigated the unfairness of popularity bias in recommender systems and how it affects different user groups differently. They show that in many recommendation algorithms the recommendations the users get are extremely concentrated on popular items even if a user is interested in long-tail and non-popular items showing an extreme bias disparity.
- **Multistakeholder Recommendation:** implementing recommender systems that balance interests of users and those of other stakeholders. There were three papers presented in this session. Louca et al. [9] presented an approach to jointly optimize for relevance and profit. Another paper regarding the optimization of revenue while keeping an acceptable level of accuracy for allocating sponsored content in recommendation was presented by Malthouse et al. [10] where they used a multi-objective binary integer programming model

for this optimization problem. The last paper was presented by [5] where authors investigated the simplification of the objective function in Groupon's search and recommendation system which is two-sided marketplace to capture the essence of short, mid and long term benefits while preserving fairness and moving users forward in the customer lifecycle.

Additional information about the workshop can be found at the following URL: <https://sites.google.com/view/rmse>

### REFERENCES

- [1] Himan Abdollahpouri and Robin Burke. 2019. Multi-stakeholder Recommendation and its Connection to Multi-sided Fairness. In *Workshop on Recommendation in Multi-stakeholder Environments (RMSE'19), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys'19, 2019*.
- [2] Himan Abdollahpouri, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. 2019. The Unfairness of Popularity Bias in Recommendation. In *Workshop on Recommendation in Multi-stakeholder Environments (RMSE'19), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys'19, 2019*.
- [3] Robin Burke, Gediminas Adomavicius, Ido Guy, Jan Krasnodebski, Luiz Pizzato, Yi Zhang, and Himan Abdollahpouri. 2017. VAMS 2017: Workshop on Value-Aware and Multistakeholder Recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. ACM, 378–379.
- [4] Yahshar Deldjoo, Vito Walter Anelli, Hamed Zamani, Alejandro Bellogin Kouki, and Tommaso Di Noia. 2019. Recommender Systems Fairness Evaluation via Generalized Cross Entropy. In *Workshop on Recommendation in Multi-stakeholder Environments (RMSE'19), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys'19, 2019*.
- [5] Joaquin Delgado, Samuel Lind, Carl Radecke, and Satish Konijeti. 2019. Simple Objectives Work Better. In *Workshop on Recommendation in Multi-stakeholder Environments (RMSE'19), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys'19, 2019*.
- [6] Michael D Ekstrand and Amit Sharma. 2017. FATREC Workshop on Responsible Recommendation. In *Proc. ACM RecSys '18*. ACM, 382–383. <https://doi.org/10.1145/3109859.3109960>
- [7] Toshihiro Kamishima, Pierre-Nicolas Schwab, and Michael D Ekstrand. 2018. 2nd FATREC workshop: responsible recommendation. In *Proc. ACM RecSys '18*. ACM, 516–516. <https://doi.org/10.1145/3240323.3240335>
- [8] Kun Lin, Nasim Sonboli, Bamshad Mobasher, and Robin Burke. 2019. Crank up the volume: preference bias amplification in collaborative recommendation. In *Workshop on Recommendation in Multi-stakeholder Environments (RMSE'19), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys'19, 2019*.
- [9] Raphael Louca, Moumita Bhattacharya, Diane Hu, and Liangjie Hong. 2019. Joint Optimization of Profit and Relevance for Recommendation Systems in E-commerce. In *Workshop on Recommendation in Multi-stakeholder Environments (RMSE'19), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys'19, 2019*.
- [10] Edward Malthouse, Khadija Ali Vakeel, Yasaman Hessary Kamyab, Robin Burke, and Morana Fuduric. 2019. A Multistakeholder Recommender Systems Algorithm for Allocating Sponsored Recommendations. In *Workshop on Recommendation in Multi-stakeholder Environments (RMSE'19), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys'19, 2019*.
- [11] Masoud Mansoury, Bamshad Mobasher, Robin Burke, and Mykola Pechenizkiy. 2019. Bias Disparity in Collaborative Recommendation: Algorithmic Evaluation and Comparison. In *Workshop on Recommendation in Multi-stakeholder Environments (RMSE'19), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys'19, 2019*.
- [12] Christopher Strucks, Manel Slokom, and Martha Larson. 2019. BlurM(or): Revisiting Gender Obfuscation in the User-Item Matrix. In *Workshop on Recommendation in Multi-stakeholder Environments (RMSE'19), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys'19, 2019*.
- [13] Virginia Tsintzou, Evangelia Pitoura, and Panayiotis Tsaparas. 2019. Bias Disparity in Recommendation Systems. In *Workshop on Recommendation in Multi-stakeholder Environments (RMSE'19), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys'19, 2019*.