

MORS 2022:

The Second Workshop on Multi-Objective Recommender Systems*

HIMAN ABDOLLAHPOURI, Spotify, United States

SHAGHAYEGH SAHEBI, State University of New York, United States

MEHDI ELAHI, University of Bergen, Norway

MASOUD MANSOURY, University of Amsterdam, Discovery Lab Elsevier, Netherlands

BABAK LONI, Meta, Netherlands

ZAHRA NAZARI, Spotify, United States

MARIA DIMAKOPOULOU, Spotify, United States

Recommender Systems are becoming an inherent part of today's Internet. They can be found anywhere from e-commerce platforms (eBay, Amazon) to music or movie streaming (Spotify, Netflix), social media (Facebook, Instagram, TikTok), travel platforms (Booking.com, Expedia), and much more. Whether a recommendation is successful or not can rely on multiple objectives such as user satisfaction, business value, and societal issues. In addition, the long-term happiness (along with short-term excitements and delight) of the users is critical for a recommender system to be considered successful. MORS workshop brings together researchers and practitioners to discuss the importance of these aspects of recommender systems and find ways to develop algorithms to build multi-objective recommenders and also evaluation metrics to assess their success.

Additional Key Words and Phrases: multi-objective recommendation, Value-aware recommendation

1 WORKSHOP DESCRIPTION

Recommender systems are software tools that are used in a variety of application domains supporting users to find relevant items, products, and services easier. Historically, the main criterion for a successful recommender system was the relevance of the recommended items to the user. In other words, the only objective for the recommendation algorithm was to learn user preferences for different items and generate recommendations accordingly. However, real-world recommender systems are well beyond a simple objective and often take into account multiple objectives. Indeed, different objectives can be important and should be considered for generating recommendations. These objectives can be either from the users' perspective or they could come from other stakeholders such as item providers and the ones that could be impacted by the recommendations.

From the users' perspective, often multiple objectives need to be considered for generating the recommendations. For example, in restaurant recommendations, several factors, such as users' taste, diet restrictions, the proximity of the restaurant, and even the price could be taken into account. Each of these aspects can be important for the users, some of which, more than the others. Therefore, it is crucial for the recommender system to incorporate all these different objectives and aspects into account when recommending some restaurants to the user. Similarly, in the education domain, a student may prefer working on simpler problems to achieve higher scores. However, some struggle is inevitable for the students to learn the new concepts properly. As a result, a problem recommender algorithm should balance the simplicity of the recommended problems and their utility to help students learn more.

*Copyright 2022 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0). Presented at the MORS workshop held in conjunction with the 16th ACM Conference on Recommender Systems (RecSys), 2022, in Seattle, WA, USA.

The objectives may also come from other stakeholders such as the item providers (those who provide the items to the platform to be recommended) or other stakeholders such as the platform owner or even side stakeholders such as the society. For example, on a music streaming service, the platform may want to balance multiple objectives, some of which are related to the users and some related to the providers (artists) and even the society as a whole. For instance, the platform may want to ensure a certain degree of exposure for different artists, so they reach their desired audience and to avoid monopoly by some superstars. The platform may also want to make sure it does not negatively affect the music culture of some smaller countries by over-exposing the users in those countries to some popular western music. These types of objectives and considerations exist in many other domains including social media, transportation, news recommendation, and food recommendation, to name a few.

MORS 2022 Main Theme. In MORS 2022, we focused on group discussions in the research community around the topics of multi-objective recommender systems. In addition, we had a main "theme" of *Long-term Optimization in Recommender Systems*. We believe this aspect of recommendations has been largely overlooked in the research community, and it is crucial to pay more attention to this aspect. Long-term optimization refers to the fact that the recommendations given to the users should help achieve long-term satisfaction for the users rather than focusing only on the short-term metrics (e.g., precision, recall, etc.). The MORS 2022 workshop was a continuation of the discussion of these topics in prior RecSys workshops including MORS 2021 [1], Workshop on Recommendation in Multi-stakeholder Environments (RMSE 2019 [5]), and Value-Aware and Multistakeholder Recommendation (VAMS 2017 [6]).

Particularly, the MORS workshop encouraged submissions addressing the following topics of interest:

- Recommender systems with multiple objectives
- Balancing the long-term impacts of the recommendations and the users' short-term preferences
- Reinforcement Learning for long-term optimization in recommender systems
- Evaluation metrics and methodologies for long-term optimization in recommendation
- Feedback loops and the impact of recommendations in long term
- Value-aware recommendation (profit, value, purpose, etc.)
- Recommendation with multiple stakeholders
- Conflict handling in multi-stakeholder recommendation
- Fairness-aware recommender systems
- News recommendation with editorial values
- Educational recommender systems with multiple, potentially conflicting, objectives

2 WORKSHOP ORGANIZERS

Himan Abdollahpouri (Spotify, USA)

Himan Abdollahpouri is a Research Scientist at Spotify. He was one of the co-chairs of MORS 2021 [1], RMSE 2019 (Recommendation in Multi-Stakeholder Environments) [5], and VAMS 2017 (Value-Aware and Multi-Stakeholder recommendation) [6] workshops at RecSys 2021, 2019 and 2017, respectively. He received his Ph.D. in Information Science at the University of Colorado Boulder. His research interests include popularity bias, multi-stakeholder recommendation, and long-term optimization in recommender systems.

Shaghayegh (Sherry) Sahebi (University at Albany – SUNY, United States)

Sherry Sahebi is an assistant professor of Computer Science at the University At Albany – SUNY and the founder of

Personalized AI (PersAI) Lab. Sherry has completed her Ph.D. in Intelligent Systems from University of Pittsburgh in 2016. Since then, she has received several internal (SUNY) and external (NSF) fundings for her projects on recommender systems and educational data mining. She has been actively involved in the research community, e.g., being a program chair of the 14th Educational Data Mining conference (EDM'21) and a proceedings chair of the 14th ACM RecSys (RecSys'20). She has prior experience in successfully co-organizing the Workshop on Intelligent Recommender Systems by Knowledge Transfer & Learning (RecSysKTL) in RecSys 2017 [18] and 2018 [14].

Mehdi Elahi (University of Bergen, Norway)

Mehdi Elahi is an Associate Professor at University of Bergen (UiB), Department of the Information Science & Media Studies (InfoMedia) where he is the co-founder of the Data Analytics and Recommender Systems lab (DARS) in Norway. He has been involved in the authorship of several grant proposals such as a large-scale grant, recently funded SFI MediaFutures with a budget of nearly NOK 300 Million, where he will serve as WP Leader for 8 years. He has published more than 80 peer-reviewed publications in top-tier journals and conferences. His current #citation is 2300+ and his H-index is 22. Together with big IT companies such as Spotify, XING, and Trivago, he has organized international challenges in 2017-2019 [2].

Masoud Mansoury (University of Amsterdam, Netherlands)

Masoud Mansoury is a postdoctoral researcher at Amsterdam Machine Learning Lab (AMLab) at University of Amsterdam, Netherlands. He is also a member of Discovery Lab collaborating with Elsevier Company in the area of recommender systems. Masoud received his PhD in Computer and Information Science from Eindhoven University of Technology, Netherlands, in 2021. He was one of the co-chair of MORS 2021 workshop [1]. His research interests include recommender systems, algorithmic bias, and contextual bandits.

Babak Loni (Meta, Netherlands)

Babak Loni is a senior Machine Learning Engineer at Meta. Babak has a Ph.D. in Machine Learning and Recommender Systems and a MS.c. in Computer Science, both from Delft University of Technology. He has been organizing RecSysNL meetups and a few RecSys workshops in the past, including the last year's MORS workshop. Babak has worked in ING, Padora Media, and DPG Media in the past where he built different solutions for personalization and recommendations.

Zahra Nazari (Spotify, United States)

Zahra Nazari is a senior research scientist at Spotify working on the design and evaluation of recommender systems with a focus on cold-start problem, multi-objective recommendations and long-term optimization. Zahra did her Ph.D. at the University of Southern California with a focus on preference elicitation and human behavior modeling in complex situations such as negotiations. Her proposed agent won the top prize at the International Automated Negotiating Agent Competitions in 2017 [11]. She has published in top conferences such as SIGIR, EC, The Web Conference, IJCAI and AAMAS and served as PC member for several of them.

Maria Dimakopoulou (Spotify, United States)

Maria Dimakopoulou is Senior Manager of ML Engineering in the Home Personalization product area at Spotify. She leads the Home Ranking team and the Home Assembly team, which focus on combining cutting-edge advances in ML, causal inference and reinforcement learning with scalable engineering to deliver Spotify Home recommendations

that optimize for the long-term satisfaction of the 400+ million Spotify users. Before joining Spotify, Maria was at Netflix, building out the causal recommenders research investment and subsequently starting and leading the Adaptive Experimentation team. Prior to that, she did her PhD on reinforcement learning at Stanford and worked at Google.

3 PROGRAM COMMITTEE

The following is a list of program committee members.

- Konstantin Bauman (Temple University, USA)
- Ludovico Boratto (University of Cagliari, Italy)
- Robin Burke (University of Colorado, Boulder, USA)
- Dietmar Jannach (Alpen-Adria-Universität Klagenfurt, Austria)
- Toshihiro Kamishima (National Institute of Advanced Industrial Science and Technology, Japan)
- Mesut Kaya (Aalborg University, Denmark)
- Peter Knees (Vienna University of Technology, Austria)
- Dominik Kowald (Know-Center, Austria)
- Kun Lin (DePaul University, USA)
- Yue Shi (Meta, CA, USA)
- Annelien Smets (Brussels Centre for Urban Studies, Belgium)
- Marko Tkalcić (University of Primorska, Slovenia)

4 TIMELINE

The following is the timeline of MORS 2022:

- First call for participation: April 15th, 2022
- Paper submission deadline: August 5th, 2022
- Notification of paper acceptance: August 27th, 2022
- Camera-ready version deadline: September 10th, 2022
- Workshop (at RecSys 2022): September 23rd, 2022

5 PROGRAM

The workshop starts with a keynote by Prof. Dietmar Jannach titled “Multi-Objective Recommendation: Overview and Challenges”[8]. The workshop then follows by six paper presentations, consisting of two long and four short contributions. The workshop is then finalized with a poster session. The following is the list of accepted papers and posters:

- [15] Sinan Seymen, Anna-Lena Sachs and Edward Malthouse. *Making smart recommendations for perishable and stockout products* (long)
- [13] Yiding Ran, Hengchang Hu and Min-Yen Kan. *PM K-LightGCN: Optimizing for Accuracy and Popularity Match in Course Recommendation* (short)
- [16] Chunpai Wang, Shaghayegh Sahebi and Peter Brusilovsky. *Proximity-Based Educational Recommendations: A Multi-Objective Framework* (long)

- [10] Oleg Lesota, Stefan Brandl, Matthias Wenzel, Alessandro Benedetto Melchiorre, Elisabeth Lex, Navid Rekabsaz and Markus Schedl. *Exploring Cross-group Discrepancies in Calibrated Popularity for Accuracy/Fairness Trade-off Optimization* (short)
- [9] Peter Knees, Andres Ferraro and Moritz Hübler. *Bias and Feedback Loops in Music Recommendation: Studies on Record Label Impact* (short)
- [3] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, Antonio Ferrara, Daniele Malitesta and Claudio Pomo. *How Neighborhood Exploration influences Novelty and Diversity in Graph Collaborative Filtering* (short)
- [4] Tahereh Arabghalizi and Alexandros Labrinidis. *A Ranked Bandit Approach for Multi-stakeholder Recommender Systems* (poster)
- [7] Mounir Hafsa, Pamela Wattebled, Julie Jacques and Laetitia Jourdan. *A Multi-Objective E-learning Recommender System at Mandarine Academy* (poster)
- [12] Renata Pelissari, Paulo Alencar, Sarah Ben Amor and Leonardo Tomazeli Duarte. *A systematic review of the use of multiple criteria decision aiding methods in recommender systems: preliminary results* (poster)
- [17] Yan Zhao, Mitchell Goodman, Sameer Kanase, Shenghe Xu, Yannick Kimmel, Brent Payne, Saad Khan and Patricia Grao. *Mitigating Targeting Bias in Content Recommendation with Causal Bandits* (poster)

6 ACKNOWLEDGEMENT

This work was supported by industry partners and the Research Council of Norway with funding to MediaFutures: Research Centre for Responsible Media Technology and Innovation, through The Centres for Research-based Innovation scheme, project number 309339. It was also partially supported by the United States National Science Foundation under Grant No. 2047500. This project was also funded by Elsevier's Discovery Lab.

REFERENCES

- [1] Himan Abdollahpouri, Mehdi Elahi, Masoud Mansoury, Shaghayegh Sahebi, Zahra Nazari, Allison Chaney, and Babak Loni. 2021. MORS 2021: 1st Workshop on Multi-Objective Recommender Systems. In *Fifteenth ACM Conference on Recommender Systems*. 787–788.
- [2] Fabian Abel, Yashar Deldjoo, Mehdi Elahi, and Daniel Kohlsdorf. 2017. Recsys challenge 2017: Offline and online evaluation. In *Proceedings of the eleventh acm conference on recommender systems*. 372–373.
- [3] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, Antonio Ferrara, Daniele Malitesta, and Claudio Pomo. 2022. How Neighborhood Exploration influences Novelty and Diversity in Graph Collaborative Filtering. In *Workshop of Multi-Objective Recommender Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022*.
- [4] Tahereh Arabghalizi and Alexandros Labrinidis. 2022. A Ranked Bandit Approach for Multi-stakeholder Recommender Systems. In *Workshop of Multi-Objective Recommender Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022*.
- [5] Robin Burke, Himan Abdollahpouri, Edward C Malthouse, KP Thai, and Yongfeng Zhang. 2019. Recommendation in multistakeholder environments. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 566–567.
- [6] Robin Burke, Gediminas Adomavicius, Ido Guy, Jan Krasnobebski, 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*. 378–379.
- [7] Mounir Hafsa, Pamela Wattebled, Julie Jacques, and Laetitia Jourdan. 2022. A Multi-Objective E-learning Recommender System at Mandarine Academy. In *Workshop of Multi-Objective Recommender Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022*.
- [8] Dietmar Jannach. 2022. Multi-Objective Recommendation: Overview and Challenges. In *Workshop of Multi-Objective Recommender Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022*.
- [9] Peter Knees, Andres Ferraro, and Moritz Hübler. 2022. Bias and Feedback Loops in Music Recommendation: Studies on Record Label Impact. In *Workshop of Multi-Objective Recommender Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022*.
- [10] Oleg Lesota, Stefan Brandl, Matthias Wenzel, Alessandro Benedetto Melchiorre, Elisabeth Lex, Navid Rekabsaz, and Markus Schedl. 2022. Exploring Cross-group Discrepancies in Calibrated Popularity for Accuracy/Fairness Trade-off Optimization. In *Workshop of Multi-Objective Recommender*

- Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022.*
- [11] Johnathan Mell, Jonathan Gratch, Tim Baarslag, Reyhan Aydođan, and Catholijn M. Jonker. 2018. Results of the First Annual Human-Agent League of the Automated Negotiating Agents Competition. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents (Sydney, NSW, Australia) (IVA '18)*. Association for Computing Machinery, New York, NY, USA, 23–28. <https://doi.org/10.1145/3267851.3267907>
 - [12] Renata Pelissari, Paulo Alencar, Sarah Ben Amor, and Leonardo Tomazeli Duarte. 2022. A systematic review of the use of multiple criteria decision aiding methods in recommender systems: preliminary results. In *Workshop of Multi-Objective Recommender Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022.*
 - [13] Yiding Ran, Hengchang Hu, and Min-Yen Kan. 2022. PM K-LightGCN: Optimizing for Accuracy and Popularity Match in Course Recommendation. In *Workshop of Multi-Objective Recommender Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022.*
 - [14] Shaghayegh Sahebi, Yong Zheng, Weike Pan, and Ignacio Fernández. 2018. The 2nd workshop on intelligent recommender systems by knowledge transfer & learning (recsysKTL). In *Proceedings of the 12th ACM Conference on Recommender Systems*. 523–524.
 - [15] Sinan Seymen, Anna-Lena Sachs, and Edward Malthouse. 2022. Making smart recommendations for perishable and stockout products. In *Workshop of Multi-Objective Recommender Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022.*
 - [16] Chunpai Wang, Shaghayegh Sahebi, and Peter Brusilovsky. 2022. Proximity-Based Educational Recommendations: A Multi-Objective Framework. In *Workshop of Multi-Objective Recommender Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022.*
 - [17] Yan Zhao, Mitchell Goodman, Sameer Kanase, Shenghe Xu, Yannick Kimmel, Brent Payne, Saad Khan, and Patricia Grao. 2022. Mitigating Targeting Bias in Content Recommendation with Causal Bandits. In *Workshop of Multi-Objective Recommender Systems (MORS'22), in conjunction with the 16th ACM Conference on Recommender Systems, RecSys'22, 2022.*
 - [18] Yong Zheng, Weike Pan, Shaghayegh Sahebi, and Ignacio Fernández. 2017. The 1st Workshop on Intelligent Recommender Systems by Knowledge Transfer & Learning: (RecSysKTL). In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. 366–367.