SURE 2024: Workshop on Strategic and Utility-aware REcommendations

Himan Abdollahpouri[1], Tonia Danylenko[2], Masoud Mansoury[3], Babak Loni[4], Daniel Russso[5] and Mihajlo Grbovic[6]

[1]Spotify, USA
[2]Spotify, Sweden
[1]Delft University of Technology
[1]Meta, Netherlands
[1]Columbia University, USA
[1]Airbnb, USA

1. Workshop Description

Nowadays, recommender systems are employed across a diverse set of application domains, not only supporting us in our decision making and choices but also helping us to discover and find new items, products, and services much more efficiently. The commonly used approach in recommender systems is receiver-centric (or user-centric) where the focus is on satisfying the receiver of the recommendations without any considerations for business & strategic objectives or the objectives of the item providers.

A recommender system that does not include strategic or business-related objectives is often referred to as "organic" recommendations, emphasizing personalized recommendations with exclusive consideration for user relevance. Conversely, strategic, sponsored or utility-aware recommendations adopt a different perspective, with the objective to optimize for both user relevancy and some kind of utility associated with those recommendations. In contrast to an organic recommender system, the focus of a strategic (or so called "non-organic") recommender system is to identify the most "relevant" users for a given item to maximize utility.

Utility can be defined in many ways depending on the problem we are solving and the domain on which the recommender systems in operating. For example, for a job recommender system on Linkedin, the utility could be to ensure the person who receive a certain job recommendation actually is qualified for it and fits the criteria of the recruiter who listed the job. It could also be monetary where a one recommendation might have a higher profit margin than another when they may have similar relevance from the user's perspective.

In real-world applications of recommender systems, aligning user-centric recommendation with overarching strategies supporting creators in their growth has become imperative in many multi-sided platforms. With advanced development of technology and research in the domain, the pressing need for a closer integration of recommender systems with both long and short term strategic goals becomes more clear. This introduces various challenges that are worth further investigation by the research community and industry practitioners, including but not limited to :

- Involvement of multiple stakeholders: The objectives in non-organic and strategic recommendations are often driven via a diverse set of stakeholders. For example, promotional content is frequently sponsored by different companies and creators, requiring consideration of the financial aspect in the recommendation process. Another example may include the strategic necessity to direct a user's attention toward content that is exclusively owned by the company, and a company

18th ACM Conference on Recommender Systems (RecSys '24), October 14–18, 2024, Bari, Italy
himana@spotify.com (H. Abdollahpouri); antoninad@spotify.com (T. Danylenko); m.mansoury@tudelft.nl (M. Mansoury); bloni@meta.com (B. Loni); djr2174@columbia.edu (D. Russso); mihajlo.grbovic@airbnb.com (M. Grbovic)

becomes a stakeholder of this recommender system. To achieve this, the recommender system would need a strategic adjustment as long as it still fits user preferences.
- Existence of multiple objectives: Non-organic recommendations inherently involve multiple objectives. For instance, there may be a trade-off between user relevance and the profit generated by the recommended item, necessitating a delicate balance. Similarly, when the goal is to diversify the recommendation experience, the relevance of the recommendations, the content popularity and its strategic potential are three different objectives that need to be balanced.
- Balancing short-term and long-term user engagement: A singular focus on optimizing for strategic recommendations that may yield immediate success and boost short-term gains. However, this approach could potentially compromise user trust over time as less quality content is recommended for a quick meaningless engagement, impacting the system's credibility in the long run.
- The need for explanation and transparency: While explanations could also enhance organic recommendations to assist users understand why such items are recommended to them, they become even more crucial in the context of non-organic and strategic recommendations due to involving other objectives in addition to item relevance. Providing rationales for recommendations is vital to maintaining user trust and encouraging consumption of strategic and utility-focused content.
- The frequency of strategic recommendations: If a non-organic recommendation is perceived as less relevant by the end-user, it becomes essential to establish reasonable limits on the frequency of such recommendations within a specified time frame. This limit can also be personalized for each user, taking their preferences and tolerance into account.
- Balance between reach and relevance: For a given item that needs to be recommended, it is important to balance the number of users to reach and the relevance of the item to those users. Reaching too many users may risk harming users' satisfaction and their trust in the system. Reaching too few users may not satisfy the creator's need who wants their item to be promoted to enough users.
- User's mental model of non-organic recommendations: How do users perceive a recommendation that combines several aspects? How does having an explicit label indicating that a recommendation is sponsored or promoted impact users' acceptance of the recommendations? Does that vary depending on what type of item is recommended (for example, a song, a tweet, an image, a product etc)?
- Competition between recommendation items: In the situations when the inventory is limited and there's a wide variety of items with diverse utilities, it's essential to establish fair competition among these items. This should be transparent and beneficial for various stakeholders.
- Forecasting and analysis: Multiple stakeholders with multiple objectives and with a limited inventory might require an alignment before a recommender system starts incorporating all the utilities. This means that we need a robust approach to forecasting the potential performance of the recommender system and the impact on its strategic objectives.

We believe that delving deeper into the challenges outlined above would significantly enhance research in recommender systems and their practical applications in various industries. SURE-2024 seeks to foster collaboration between researchers from academia and industry, providing a platform for in-depth exploration and discussion of this crucial problem. The objective is to explore and exchange the ideas of innovative approaches, methodologies, and case studies that can effectively tackle the aforementioned challenges.

In particularly, the SURE 2024 workshop encourages submissions addressing the following topics of interest:

- Recommendation with multiple stakeholders
- Applications of personalization in advertising and promotions

- Recommender systems with multiple objectives
- Studying different domains where strategic and utility-aware recommendations can be important. For instance, in job recommendation, providers (recruiter) may have preference for who should receive their recommendations but that may not be the case in a typical movie recommendation.
- Estimation and optimization methods for the long-term value of recommender systems
- Long-term community or audience growth for recommended items
- Explainable recommendations, especially for strategic and utility-aware recommendations.
- Methods for estimating trade-offs between user retention and satisfaction and a utility value of recommendations
- The impact of different objectives on the short-term and long-term success of the recommender systems.
- Simulation for experimentation in multi-objective, multi-stakeholder recommenders, and long-term user satisfaction.
- Users' trust in recommendations for non-organic recommendations.

2. Program committee

The following is the confirmed list of program committee members:

- Lucas Maystre (Spotify, London)
- Claudia Hauff (Spotify, Netherlands)
- Yu Zhao (Spotify, Sweden)
- Ludovico Boratto (University of Cagliari, Italy)
- Dietmar Jannach (Alpen-Adria-Universität Klagenfurt, Austria)
- Toshihiro Kamishima (National Institute of Advanced Industrial Science and Technology, Japan)
- Yashar Deldjoo (Polytechnic University of Bari, Italy)
- Olivier Jeunen (ShareChat, UK)
- Mirko Marras (University of Cagliari, Italy)
- Hossein A. Rahmani (University College London, UK)
- Milad Sabouri (DePaul University)
- Manel Slokom (Delft University of Technology)

3. Workshop organizers

This year's workshop is organized by the following researchers:
Himan Abdollahpouri (Spotify, USA)
Himan Abdollahpouri is a Research Scientist at Spotify. He was one of the co-chairs of MORS 2022 [1], MORS 2021 [2], RMSE 2019 (Recommendation in Multi-Stakeholder Environments) [3], and VAMS 2017 (Value-Aware and Multi-Stakeholder recommendation) [4] workshops at the ACM Conference on Recommender Systems (RecSys). He received his Ph.D. in Computer & Information Science at the University of Colorado Boulder in 2020. His research interests include popularity bias, multi-stakeholder and multi-objective recommendation, and long-term optimization in recommender systems.

Tonia Danylenko (Spotify, Sweden)
Tonia Danylenko is a Senior Machine Learning Engineering Manager at Spotify, and a co-organizer of WiDS AI and ML Sweden. At Spotify Tonia is focusing on strategic and utility-aware recommendations as an essential part of core personalization. Before joining Spotify, Tonia led an applied machine learning team at Viaplay and strategic data science initiatives at IKEA. Tonia holds a PhD in Computer Science from Linnaeus University, Sweden, and has an interest in machine learning and Generative AI

in personalization and advertising.

Masoud Mansoury (Delft University of Technology, Netherlands)
Masoud Mansoury is an Assistant Professor at Delft University of Technology in the Netherlands. He earned his PhD in Computer and Information Science from Eindhoven University of Technology. Masoud has twice co-organized the MORS workshop at the ACM Conference on Recommender Systems (RecSys) in both 2021 [2] and 2022 [1]. His research primarily focuses on the development of trustworthy and explainable recommender systems, with a particular interest in 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 an 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 first and the second Workshops of Multi-Objective Recommender Systems (MORS). Babak has worked in ING, Pandora Media, and DPG Media in the past where he built different solutions for personalization and recommendations.

Daniel Russo (Colombia University, USA)
Daniel Russo is a Philip H. Geier Jr. Associate Professor in the Decision, Risk, and Operations division of Columbia Business School. His research lies at the intersection of statistical machine learning and online decision making, mostly falling under the broad umbrella of reinforcement learning. His research has been recognized by the Frederick W. Lanchester Prize, a Junior Faculty Interest Group Best Paper Award, and first place in the George Nicholson Student Paper Competition. He serves as an associate editor of the journals Operations Research, Management Science, and Stochastic Systems. Outside academia, Daniel works with Spotify to leverage reinforcement learning techniques and AI foundation models in audio recommendations.

Mihajlo Grbovic (Airbnb, USA)
Mihajlo Grbovic is a Machine Learning Scientist at Airbnb. He holds a PhD in Machine Learning from Temple University in Philadelphia. He has more than 15 years of technical experience in applied Machine Learning, acting as a Science Lead in a portfolio of projects at Yahoo and now at Airbnb. During his time at Yahoo, from 2012 to 2016, he worked on integrating Machine Learning in various Yahoo Products, such as Yahoo Mail, Search, Tumblr & Ads. Some of his biggest accomplishments include building Machine Learning-powered Ad Targeting for Tumblr, being one of the key developers of Email Classification for Yahoo Mail and introducing the next generation of query-ad matching algorithms to Yahoo Search Ads. Dr. Grbovic joined Airbnb in 2016 as a Machine Learning Scientist, specializing in Machine Learning. He works mostly on Search & Recommendation problems for Airbnb Homes and Experiences. Some of his key accomplishments include building the first Airbnb Search Autocomplete algorithm, building out Machine Learning-powered Search for Airbnb Experiences, building algorithms that power Airbnb Categories that are currently showcased on Airbnb Homepage. Currently, he is working on building an AI Travel Concierge at Airbnb. Dr. Grbovic published more than 60 peer-reviewed publications at top Machine Learning and Web Science Conferences, and co-authored more than 10 patents (h-index: 25; citations: 3073; i10-index: 37). He was awarded the Best Paper Award at KDD 2018 Conference. His work was featured in Wall Street Journal, Scientific American, MIT Technology Review, Popular Science and Market Watch.

References

[1] H. Abdollahpouri, S. Sahebi, M. Elahi, M. Mansoury, B. Loni, Z. Nazari, M. Dimakopoulou, Mors 2022: The second workshop on multi-objective recommender systems, in: Proceedings of the 16th ACM Conference on Recommender Systems, 2022, pp. 658–660.

[2] H. Abdollahpouri, M. Elahi, M. Mansoury, S. Sahebi, Z. Nazari, A. Chaney, B. Loni, Mors 2021: 1st workshop on multi-objective recommender systems, in: Fifteenth ACM Conference on Recommender Systems, 2021, pp. 787–788.
[3] R. Burke, H. Abdollahpouri, E. C. Malthouse, K. Thai, Y. Zhang, Recommendation in multistakeholder environments, in: Proceedings of the 13th ACM Conference on Recommender Systems, 2019, pp. 566–567.
[4] R. Burke, G. Adomavicius, I. Guy, J. Krasnodebski, L. Pizzato, Y. Zhang, H. Abdollahpouri, Vams 2017: Workshop on value-aware and multistakeholder recommendation, in: Proceedings of the Eleventh ACM Conference on Recommender Systems, 2017, pp. 378–379.