Special Issue on Responsible Recommender Systems Part 2

LINA YAO, CSIRO's Data61 and University of New South Wales, Australia
JULIAN MCAULEY, University of California, USA
XIANZHI WANG, University of Technology Sydney, Australia
DIETMAR JANNACH, University of Klagenfurt, Austria

CCS Concepts: • **Information systems** → **Recommender systems**.

Additional Key Words and Phrases: Recommender Systems, Fairness, Explainability, Privacy

1 INTRODUCTION

Recommender systems are information filtering systems that suggest items tailored to individual users or user groups. They represent a powerful machine learning tool to support various human decision-making activities in e-commerce, social networks, entertainment, transportation, healthcare, and cybersecurity. Existing recommender systems typically focus on accuracy and personalization but increasingly call for an effective means of ensuring the systems work responsibly. Without appropriate responsible techniques, recommender systems could have undesired effects on users, communities, and society. For example, a recommendation algorithm trained on imbalanced data might be biased toward catering to the preferences of a majority group of users while overlooking minority groups; a system without countermeasures to misinformation may amplify the spread of misinformation, and a recommendation that lacks appropriate disclosure of the decision-making process or intuitive explanations may not be easily trusted by users.

Responsible recommender systems require innovative ways of assessing recommendation contexts and processing, communicating and presenting the recommendations. While responsible recommender systems have attracted attention in recent years, the corresponding challenges and threats in technical, social/societal, and ethical contexts are yet to be fully addressed. Furthermore, while existing approaches to the issue mostly focus on a single aspect, such as fairness, explainability, or social trust, a systematic/holistic approach to responsible recommender systems is still in demand. To address the challenges and seize the opportunities posed by responsible recommender systems, existing methodologies, models, techniques, and applications must be reassessed, adapted or transformed significantly. Besides, emerging techniques for analyzing all types of side information (e.g., multimodal attributes, social information, external knowledge) should be fully explored to advance or expand the scope of state-of-the-art recommendation research toward delivering more responsible systems.

This second part of the special issue comprises fourteen papers that present recent advances and novel contributions in the emerging yet promising field of responsible recommender systems. They represent some of the most recent progress in advancing responsible recommender systems research in four directions below.

Authors' addresses: Lina Yao, lina.yao@unsw.edu.au, CSIRO's Data61 and University of New South Wales, Sydney, Australia; Julian McAuley, jmcauley@eng.ucsd.edu, University of California, San Diego, USA; Xianzhi Wang, xianzhi.wang@uts.edu.au, University of Technology Sydney, Sydney, Australia; Dietmar Jannach, dietmar.jannach@aau.at, University of Klagenfurt, Klagenfurt, Austria.

ACM 2157-6912/2024/10-ART
https://doi.org/10.1145/3689367

Fairness and Biases. Zhao *et al.* [13] survey fairness and diversity issues in recommender systems. A notable feature of this work is that it considers fairness and diversity to be two strongly connected domains, extending the notion of diversity to encompass both the item and user levels. This enables a re-interpretation of fairness studies from the diversity viewpoint, promoting an enhanced understanding of fairness in recommendation contexts. Aiming at reducing hidden bias in data and ensuring fairness in algorithmic data analysis, Anagnostopoulos *et al.* [1] formulate a theoretic problem of deriving a fair densest subgraph or a clustering with fairness constraints. The authors prove and experimentally validate that a suitable spectral embedding allows the recovery of an almost optimal, fair, dense subgraph hidden in the input data (if one exists). Further, efficient algorithms are presented for the densest subgraph and clustering problems. Yan *et al.* [10] focus on a particular type of bias, popularity bias, in a mashup creation context, where the task is to recommend a set of APIs for reuse in building a new mashup application. The work analyses popularity bias introduced by correlation graph-based API recommendation approaches. Subsequently, we empirically validate the presence of popularity bias in API recommendations, followed by presenting a recommendation approach that can mitigate popularity bias. Similarly, Zhang *et al.* [12] look at a particular source of fairness—inaccurate user feedback in dynamic learning-to-rank; they first demonstrate the adverse impact of the overlooked interesting items on ranking results; then, by treating those overlooked interesting items as noise, they propose a Co-teaching Rank (CoTeR) framework incorporated with a customized loss function and a communicated data sampling strategy to improve the utilities while preserving ranking fairness. adaptability of the ranking system in dynamic environments. The strong adaptability of the framework makes it particularly useful for scenarios where new items are introduced without pre-existing historical data. Instead of mitigating biases or unfairness, Medda *et al.* [6] aim to leverage counterfactuality methods for explaining unfairness in CNN-based recommender systems. The authors identify sources of unfairness by removing suitable edges from the original bipartite graph that models user-item interactions. Although the approach does not mitigate unfairness directly, it offers insights into where and how unfairness could be reduced.

Explainability, Trustworthiness, and Robustness. Moller *et al.* [7] aim to understand the extent to which the recommendations of news recommender systems are based on (or can be explained by) content-related evidence. To this end, they present a neural news recommender that can attribute individual recommendations to news items and words in the input reading histories, disclosing several interesting findings, e.g., many recommendations are not based on content-related evidence, many users' clicks on the news are not explainable from reading histories, and recommendations stem from a spurious bias in user representations. Liao *et al.* [5] aim for explainable recommendations and propose a novel multi-task model that combines multimodal contrastive learning and personalized aspect selection for recommendation and explanation tasks. The work features integrating implicit feedback derived from user-item interactions and user-written text reviews, demonstrating superior recommendation accuracy and explanation quality to a number of baselines. Zhuang *et al.* [14] also target explainable recommendations but focus on generating faithful and factually consistent reviews as natural language explanations for recommendations, where contrastive learning is leveraged to optimize the model to distinguish faithful explanations from unfaithful explanations with factual errors. Carragher *et al.* [3] also aim for non-factual data (or misinformation) resilience but set in an information retrieval context. The authors introduce and evaluate interventions to reduce traffic to unreliable news domains from search engines while maintaining traffic to reliable domains. This contributes to the development of targeted strategies to enhance the trustworthiness and quality of search engine results, Rather than explainability or trustworthiness, Ye *et al.* [11] focus on enhancing the robustness of recommender systems under rating flip noise, which widely exists in the interaction history of recommender systems. Through a matrix modeling of flip noise and applying loss correction methods, they demonstrate the advantages of the proposed robust approach in coping with noisy data. Uslu *et al.* [9] follow a control-theoretical approach to decision-making for trustworthy and responsible natural resource management. Specifically, they propose a hybrid framework that builds bottom-up trust to enhance cooperation among decision-makers in the food, energy, and water sectors. The work focuses on how community decisions can be affected by individuals' trust and acceptance of proposed solutions, including those generated by AI-based recommender systems.

Privacy. Sun *et al.* [8] address the data heterogeneous and data sparsity challenges for graph neural network (GNN) based social recommendation in federated learning (FL) setting, a privacy-preserving alternative to traditional centralized training methods. By leveraging the hypergraph structure to address the heterogeneity of data, the approach shows improvements over several matrix factorization-based, GNN-based, and Federated learning methods. Ariyarathna *et al.* [2] consider a specific adversary that recovers user trajectories used to train the federated recommendation models with high proximity accuracy in FL-based personalized travel route recommendation, where FL protects users' sensitive location data. The proposed solution uses shared gradients from global model training in FL to reconstruct private user trajectories. It shows better recovery results than state-of-the-art FL data leakage attacks. It is also more robust than existing attacks when noises are added to defend against the recovery.

Position Paper. Golbeck *et al.* [4] recognize that recommender systems could negatively impact people in recovery from eating disorders (ED), i.e., recommending content that features severely underweight bodies or encourages disordered eating could cause ED to relapse. Specifically, they analyze the content of tweets about recommended ED content and show ED content is recommended through a small-scale study on Pinterest, revealing implications for responsible recommendation and harm prevention. Although the work focuses on a specific domain of ED, it raises concerns over potentially harmful content for vulnerable populations beyond ED. This work can be considered a case example of how responsible recommender systems are paramount for broad applications.

The guest editors would like to thank all the authors for contributing to the special issue and the reviewers for offering their constructive reviews to uphold the quality of publications in the journal. Again, we thank the journal's editorial team for their professional support and collaboration.

REFERENCES

[1] Aris Anagnostopoulos, Luca Becchetti, Matteo Bohm, Adriano Fazzone, Stefano Leonardi, Cristina Menghini, and Chris Schwiegelshohn. 2024. Fair Projections as a Means Towards Balanced Recommendations. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[2] Thirasara Ariyarathna, Meisam Mohommady, Hye-Young Paik, and Salil S.; Kanhere. 2024. DeepSneak: User GPS Trajectory Reconstruction from Federated Route Recommendation Models. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[3] Peter Carragher, Evan Williams, and Kathleen Carley. 2024. Misinformation Resilient Search Rankings with Webgraph-based Interventions. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[4] Jennifer Golbeck. 2024. Recommender System-Induced Eating Disorder Relapse: Harmful Content and the Challenges of Responsible Recommendation. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[5] Hao Liao, Shuo Wang, Hao Cheng, Wei Zhang, Jiwei Zhang, Mingyang Zhou, Kezhong Lu, Rui Mao, and Xing Xie. 2024. Aspect-enhanced Explainable Recommendation with Multi-modal Contrastive Learning. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[6] Giacomo Medda, Francesco Fabbri, Mirko Marras, Ludovico Boratto, and Gianni Fenu. 2024. GNNUERS: Fairness Explanation in GNNs for Recommendation via Counterfactual Reasoning. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[7] Lucas Möller and Sebastian Pado. 2024. Explaining Neural News Recommendation with Attributions onto Reading Histories. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[8] Hongliang Sun, Zhiying Tu, Dianbo Sui, Bolin Zhang, and Xiaofei Xu. 2024. A Federated Social Recommendation Approach with Enhanced Hypergraph Neural Network. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[9] Suleyman Uslu, Davinder Kaur, Samuel Rivera, Arjan Durresi, Meghna Babbar-Sebens, and Jenna Tilt. 2024. A Trustworthy and Responsible Decision-Making Framework for Resource Management in Food-Energy-Water Nexus: A Control-Theoretical Approach. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[10] Chao Yan, Arif Ali Zhong, Weiyinaand Khan, Wenwen Gong, Yanwei Xu, and Baogui Xin. 2024. Popularity Bias in Correlation Graph based API Recommendation for Mashup Creation. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[11] Shanshan Ye and Jie Lu. 2024. Robust Recommender Systems with Rating Flip Noise. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[12] Chenhao Zhang, Weitong Chen, Wei Zhang, and Miao Xu. 2024. Mitigating the Impact of Inaccurate Feedback in Dynamic Learning-to-Rank: A Study of Overlooked Interesting Item. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[13] Yuying Zhao, Yu Wang, Yunchao Liu, Xueqi Cheng, Charu Aggarwal, and Tyler Derr. 2024. Fairness and Diversity in Recommender Systems: A Survey. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.

[14] Haojie Zhuang, Wei Zhang, Weitong Chen, Jian Yang, and Quan Z. Sheng. 2024. Improving Faithfulness and Factuality with Contrastive Learning in Explainable Recommendation. *ACM Transactions on Intelligent Systems and Technology* 15 (2024). Issue 6.