

The 2nd Workshop on Recommendation with Generative Models

Website: <https://generative-rec.github.io/workshop/>

Wenjie Wang
wangwenjie@u.nus.edu
National University of Singapore

Yang Zhang
zy2015@mail.ustc.edu.cn
University of Science and Technology
of China

Xinyu Lin
xylin1028@gmail.com
National University of Singapore

Fuli Feng
fulifeng93@gmail.com
University of Science and Technology
of China

Weiwen Liu
liuweiwen8@huawei.com
Huawei Noah's Ark Lab

Yong Liu
liu.yong6@huawei.com
Huawei Noah's Ark Lab

Xiangyu Zhao
xianzhao@cityu.edu.hk
City University of Hong Kong

Wayne Xin Zhao
batmanfly@gmail.com
Renmin University of China

Yang Song
yangsong@kuaishou.com
Kuaishou Technology

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

ABSTRACT

The rise of generative models has driven significant advancements in recommender systems, leaving unique opportunities for enhancing users' personalized recommendations. This workshop serves as a platform for researchers to explore and exchange innovative concepts related to the integration of generative models into recommender systems. It primarily focuses on five key perspectives: (i) improving recommender algorithms, (ii) generating personalized content, (iii) evolving the user-system interaction paradigm, (iv) enhancing trustworthiness checks, and (v) refining evaluation methodologies for generative recommendations. With generative models advancing rapidly, an increasing body of research is emerging in these domains, underscoring the timeliness and critical importance of this workshop. The related research will introduce innovative technologies to recommender systems and contribute to fresh challenges in both academia and industry. In the long term, this research direction has the potential to revolutionize the traditional recommender paradigms and foster the development of next-generation recommender systems.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Generative Models for Recommendation, Large Language Models, Trustworthy Recommendation

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WWW '24 Companion, May 13–17, 2024, Singapore, Singapore

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0172-6/24/05.

<https://doi.org/10.1145/3589335.3641303>

ACM Reference Format:

Wenjie Wang, Yang Zhang, Xinyu Lin, Fuli Feng, Weiwen Liu, Yong Liu, Xiangyu Zhao, Wayne Xin Zhao, Yang Song, and Xiangnan He. 2024. The 2nd Workshop on Recommendation with Generative Models: Website: <https://generative-rec.github.io/workshop/>. In *Companion Proceedings of the ACM Web Conference 2024 (WWW '24 Companion)*, May 13–17, 2024, Singapore, Singapore. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3589335.3641303>

1 SCOPE AND TOPICS

The main objective of this workshop is to encourage pioneering research in the integration of generative models with recommender systems, with a specific focus on five key aspects. First, this workshop will motivate active researchers to utilize generative models for enhancing recommender algorithms and refining user modeling. Second, it promotes utilizing generative models to generate diverse content, *i.e.*, AI-generated content (AIGC), in certain situations, complementing human-generated content to satisfy a broader range of user preferences and information needs. Third, it embraces substantial innovations in user interactions with recommender systems, possibly driven by the boom of large language models (LLMs). Fourth, the workshop will highlight the significance of trust in employing generative models for recommendations, encompassing aspects like content trustworthiness, algorithmic biases, and adherence to evolving ethical and legal standards. Lastly, the workshop will prompt researchers to develop diverse methods for the evaluation, including novel metrics and human evaluation approaches.

The workshop provides an invaluable forum for researchers to present the latest advancements in the rapidly evolving field of recommender systems. We welcome original submissions focusing on generative models in recommender systems, including a range of relevant topics:

- Leveraging LLMs and other generative models such as diffusion models to improve user modeling and various recommendation

tasks, including sequential, cold-start, social, conversational, multimodal, and causal recommendation tasks.

- Improving generative recommender models (e.g., LLM-based recommenders) from different aspects, such as model architecture, and training and inference efficiency.
- Combining external knowledge from LLMs or other generative models to enhance user and item representation learning.
- Generative recommendation by harnessing generative AI to drive personalized item creation or editing, particularly in contexts such as advertisement, image, and micro-video.
- Innovation of user-system interaction paradigm for effective user feedback by leveraging strong conversational capability of LLMs.
- Real-world applications of generative recommender systems, ranging from finance to streaming platforms and social networks.
- Trustworthy recommendation with generative models, for example, developing the standards and technologies to improve or inspect the recommendations from the aspects of bias, fairness, privacy, safety, authenticity, legal compliance, and identifiability.
- Developing generative agents empowered by LLMs, motivating the recommendation agents from user simulation and data collection, to algorithm enhancement and evaluation.
- Evaluation of generative recommender systems, including new evaluation metrics, standards, and human evaluation approaches.

2 RATIONALE

2.1 Relevance

This workshop, centered on recommendation using generative models, aligns seamlessly with the Web conference, as it spotlights a pivotal research trend within recommender systems, which is an essential area at the Web conference. Moreover, incorporating generative models into recommender systems can improve information filtering services to cater to users' personalized information needs in a variety of Web applications.

2.2 Objectives and Expected Outcome

This workshop will encourage researchers to venture into new horizons within recommender systems by incorporating powerful generative models. These ongoing endeavors and emerging technologies will introduce fresh characteristics to industry products and motivate innovative research topics in academia. Extensive applications, spanning streaming media, social networks, and forums, are well-positioned to embrace these generative model techniques. In the long run, this research direction holds the potential to revolutionize the established recommender paradigm, giving rise to the evolution of next-generation recommender systems.

We expect to receive innovative contributions in this promising field with exciting ideas and novel methodologies. Some high-quality papers can also be encouraged to be submitted an extension version to our organized ACM TOIS special issue on using pre-trained models for recommendation.

2.3 Target Audience

This workshop's attractiveness stems from its dedication to an evolving area of recommender systems [2–4, 6, 9, 11, 15, 17, 18, 24]. It aims to draw the attention of a diverse audience, including researchers, industry experts, and academics. The workshop offers

a unique forum for these stakeholders to share innovative ideas, methods, and accomplishments, encouraging interdisciplinary collaboration and the exploration of novel applications.

2.4 Related Workshops

We have initiated the first workshop on recommendation with generative models, engaging many researchers and experts in the community to discuss research progress. The first workshop is held at the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23), in Birmingham UK, on October 22nd. It has attracted extensive attention with around 50 offline participants and over 60 online attendees via Zoom. The first workshop accepted 9 high-quality contributions out of 11 submissions, covering the research papers and surveys on LLM-based recommendation and diffusion recommender models. More information is at <https://rgm-cikm23.github.io/>. Besides, the Gen-IR workshop¹ at SIGIR'23 encourages the exploration of generative information retrieval. However, this workshop is more about general information retrieval, which is not specialized in leveraging generative models to drive the advancements in recommender systems.

The rapid progress of generative models for recommendation has spurred a wave of innovative research efforts in the past months [5, 9, 12, 18, 19, 21]. Notably, some recent studies explore new research directions, including improvements in data/model efficiency and architecture of generative recommendation [7, 10, 12], multimodal generative recommendation [5, 26], personalized outfit recommendation [20], LLM-based federated recommendation [25], and generative agents for recommendation [16, 21, 22]. As such, it is very necessary to host the second workshop to foster discussions on these evolving directions.

3 WORKSHOP PROGRAM FORMAT

This workshop will be held for **half a day**. We will invite two researchers in this field to give 30-minute keynote talks. We will also invite several senior researchers and developers to organize a panel discussion on future directions. Besides, in this workshop, we have received 10 submissions and accepted 8 papers as follows:

- Diffusion Recommendation with Implicit Sequence Influence. Yong Niu, Xing Xing, Zhichun Jia, Ruidi Liu, Mindong Xin and Jianfu Cui.
- A Study of Implicit User Unfairness in Large Language Models for Recommendation. Chen Xu, Wenjie Wang, Yuxin Li, Liang Pang, Jun Xu and Tat-Seng Chua.
- Aligning GPTRec with Beyond-Accuracy Goals with Reinforcement Learning. Aleksandr Vladimirovich Petrov and Craig Macdonald.
- Controllable and Transparent Textual Latents for Recommender Systems. Emiliano Penalzoza, Haolun Wu, Olivier Gouvert and Laurent Charlin.
- How Reliable is Your Simulator? Analysis on the Limitations of Current LLM-based User Simulators for Conversational Recommendation. Lixi Zhu, Xiaowen Huang and Jitao Sang.
- Multimodal Conditioned Diffusion Model for Recommendation. Haokai Ma, Yimeng Yang, Lei Meng, Ruobing Xie and Xiangxu Meng.
- Bridging Items and Language: A Transition Paradigm for Large Language Model-Based Recommendation. Xinyu Lin, Wenjie Wang, Yongqi Li, Fuli Feng, See-Kiong Ng and Tat-Seng Chua.
- OutfitGPT: LLMs as Fashion Outfit Generator and Recommender. Yujuan Ding, Junrong Liao, Wenqi Fan, Yi Bin and Qing Li.

¹<https://coda.io/@sigir/gen-ir>.

Each paper would be given around 10 minutes for presentation and QA. We present the preliminary program schedule² as follows:

Table 1: Program schedule.

Event	Time
Opening Remarks from Co-Chairs	08:40–08:50
Keynote Talk #1 followed by QA	08:50–09:20
Paper Session #1 (4 papers)	09:20–10:00
Tea Break	10:00–10:30
Keynote Talk #2 followed by QA	10:30–11:00
Paper Session #2 (4 papers)	11:00–11:40
Panel Discussion & Closing Remarks	11:40–12:00

4 CALL FOR PAPERS

4.1 Introduction

The surge in generative models has catalyzed substantial progress within recommender systems. For example, pre-trained generative models have demonstrated their capability to effectively learn user preferences from historical interactions [1, 2]; generative models might help to produce item content to meet users' diverse information needs in some scenarios [17], generative models have shown promise in generating item content that caters to the diverse users' information needs in specific contexts; and the emergence of ChatGPT-like language models offers novel interaction modes to obtain users' feedback and intension [8, 14].

In this light, user experience can be potentially enhanced by advancing the traditional recommender paradigms via generative models. This workshop provides a platform to facilitate the integration of generative models into recommender systems, with a focus on user modeling, content generation, interaction patterns, trustworthiness evaluations [23], and evaluation methods [13].

4.2 Objectives and Scope

This workshop aims to encourage innovative research on integrating generative models with recommender systems, particularly on five key aspects: (i) enhancing algorithms for user modeling by generative models; (ii) generating personalized content to supplement human-generated content; (iii) evolving the user interaction modes with recommender systems; (iv) prioritizing trustworthiness in generative recommendation; (v) formulating evaluation techniques for generative model-based recommender systems. We summarize the detailed objectives and scope in Section 1.

4.3 Submission Information

Submission Guidelines: Submitted papers must be a single PDF file in the template of ACM WWW 2024. Submissions can be of varying length from 4 to 8 pages, plus unlimited pages for references. The authors may decide on the appropriate length of the paper as no distinction is made between long and short papers. All submitted papers will follow the "double-blind" review policy and undergo the same review process and duration. Expert peer reviewers in the

²Kindly note that the paper acceptance plan and the program schedule are tentative and may be subject to potential adjustments according to the requirements of the conference chairs.

field will assess all papers based on their relevance to the workshop, scientific novelty, and technical quality. The timeline is as follows:

- Submissions deadline: February 26, 2024
- Paper acceptance notification: March 4, 2024
- Workshop date: May 13, 2024

5 ORGANIZERS

• Wenjie Wang

Email: wangwenjie@u.nus.edu

Affiliation: National University of Singapore

Biography: Dr. Wenjie Wang is a research fellow at National University of Singapore (NUS). He received Ph.D. in Computer Science from NUS, supervised by Prof. Tat-Sent Chua. Dr. Wang was a winner of Google Ph.D. Fellowships. His research interests cover recommender systems, data mining, and causal inference. His first-author publications appear in top conferences and journals such as SIGIR, KDD, WWW, WSDM, and TOIS. His work has been selected into ACM MM 2019 Best Paper Final List.

• Yang Zhang

Email: zy2015@mail.ustc.edu.cn

Affiliation: University of Science and Technology of China

Biography: Yang Zhang is a Ph.D. candidate at the University of Science and Technology of China (USTC), under the supervision of Prof. Xiangnan He. His research interests lie in recommender systems and causal inference, and he has published several first-author papers in top conferences. In particular, he received the Best Paper Honorable Mention in SIGIR 2021 for his work on causal recommendation.

• Xinyu Lin

Email: xylin1028@gmail.com

Affiliation: National University of Singapore

Biography: Xinyu Lin is a Ph.D. candidate at the University of Singapore, under the supervision of Prof. Tat-seng Chua. Her research interests lie in recommender systems, and her work has been published in top conferences and journals such as SIGIR, WWW, CIKM, and TOIS. Moreover, she has also served as the reviewer and PC member for the top conferences and journals, including SIGIR, WSDM, and TOIS.

• Fuli Feng

Email: fulifeng93@gmail.com

Affiliation: University of Science and Technology of China

Biography: Dr. Fuli Feng is a professor in University of Science and Technology of China. He received Ph.D. in Computer Science from National University of Singapore in 2019. His research interests include information retrieval, data mining, causal inference, and multi-media processing. He has over 60 publications appeared in several top conferences such as SIGIR, WWW, and SIGKDD, and journals including TKDE and TOIS. He has received the Best Paper Honourable Mention of SIGIR 2021 and Best Poster Award of WWW 2018. Moreover, he organized the 1st workshop on Information Retrieval in Finance at SIGIR'20.

• Weiwen Liu

Email: liuweiwen8@huawei.com

Affiliation: Huawei Noah's Ark Lab

Biography: Dr. Weiwen Liu is currently a senior researcher at

Huawei Noah's Ark Lab. She received her Ph.D. in Computer Science and Engineering from the Chinese University of Hong Kong in 2020. Her research is broadly concerned with recommender systems, information retrieval, and user preference learning. She has published over 40 papers on top conferences including KDD, SIGIR, and WWW. She gave a tutorial on neural re-ranking at RecSys'22. She will co-organize the DLP workshop at RecSys'23.

• Yong Liu

Email: liu.yong6@huawei.com

Affiliation: Huawei Noah's Ark Lab

Biography: Dr. Yong Liu is a Senior Principal Researcher at Huawei Noah's Ark Lab, Singapore. Prior to joining Huawei, he was a Senior Research Scientist at Nanyang Technological University (NTU), a Data Scientist at NTUC Enterprise, and a Research Scientist at Institute for Infocomm Research (I2R), A*STAR, Singapore. Moreover, he has served as the Challenge Co-chair for RecSys 2023, PC Co-chair for ICCSE 2021.

• Xiangyu Zhao

Email: xianzhao@cityu.edu.hk

Affiliation: City University of Hong Kong

Biography: Prof. Xiangyu Zhao is a tenure-track assistant professor of Data Science at City University of Hong Kong (CityU). His research has been awarded ICDM'22 and ICDM'21 Best-ranked Papers, Global Top 100 Chinese New Stars in AI, CCF-Tencent Open Fund (twice), CCF-Ant Research Fund, CCF-BaiChuan-Ebtech Foundation Model Fund, Ant Group Research Fund, Tencent Focused Research Fund, Criteo Faculty Research Award, Bytedance Research Collaboration Program, and nomination for Joint AAAI/ACM SIGAI Doctoral Dissertation Award. He also co-organizes DRL4KDD and DRL4IR workshops at KDD'19, WWW'21, SIGIR'20/21/22 and CIKM'23.

• Wayne Xin Zhao

Email: batmanfly@gmail.com

Affiliation: Renmin University of China

Biography: Dr. (Wayne) Xin Zhao is currently a professor at Renmin University of China. He obtained the doctoral degree from Peking University in July 2014. He has broad research interest in the fields of information retrieval and natural language processing, with 100+ published papers at top-tier conferences/journals and 10000+ academic citations from Google Scholar. He received ECIR 2021 Test of Time Award, RecSys 2022 Best Student Paper Runner-up, CIKM 2022 Best Resource Paper Runner-up, and other awards.

• Yang Song

Email: yangsong@kuaishou.com

Affiliation: Kuaishou Technology

Biography: Dr. Yang Song is currently the Head of Recommendation at Kwai, overseeing both core modeling and data mining teams. He has published over 70 papers in conferences and journals. He has served as PC&Area Chairs in Recsys, WSDM, TheWebConf(WWW), IEEE Big Data etc.

• Xiangnan He

Email: xiangnanhe@gmail.com

Affiliation: University of Science and Technology of China

Biography: Dr. Xiangnan He is a professor at the University of Science and Technology of China (USTC). He has over 100

publications appeared in top conferences such as SIGIR, WWW, and KDD, and journals including TKDE, TOIS, and TNNLS. His work on recommender system has received the Best Paper Award Honourable Mention in SIGIR (2021, 2016) and WWW (2018). He has rich experience in organizing workshops and tutorials at SIGIR'18, WSDM'19&20, WWW'21&22, and RecSys'21.

REFERENCES

- [1] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. TALLRec: An Effective and Efficient Tuning Framework to Align Large Language Model with Recommendation. In *RecSys*. ACM.
- [2] Zeyu Cui, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. M6-Rec: Generative Pretrained Language Models are Open-Ended Recommender Systems. *arXiv:2205.08084*.
- [3] Romain Deffayet, Thibaut Thonet, Jean-Michel Renders, and Maarten de Rijke. 2023. Generative Slate Recommendation with Reinforcement Learning. In *The Sixteenth International Conference on Web Search and Data Mining*. ACM, 580–588.
- [4] Yunfan Gao, Tao Sheng, Youlin Xiang, Yun Xiong, Haofen Wang, and Jiawei Zhang. 2023. Chat-REC: Towards Interactive and Explainable LLMs-Augmented Recommender System. *arXiv:2303.14524*.
- [5] Shijie Geng, Juntao Tan, Shuchang Liu, Zuohui Fu, and Yongfeng Zhang. 2023. VIP5: Towards Multimodal Foundation Models for Recommendation. In *EMNLP*.
- [6] Jiming Li, Wentao Zhang, Tian Wang, Guanglei Xiong, Alan Lu, and Gerard Medioni. 2023. GPT4Rec: A Generative Framework for Personalized Recommendation and User Interests Interpretation.
- [7] Lei Li, Yongfeng Zhang, and Li Chen. 2023. Prompt distillation for efficient llm-based recommendation. In *CIKM*. 1348–1357.
- [8] Lizi Liao, Grace Hui Yang, and Chirag Shah. 2023. Proactive Conversational Agents. In *The Sixteenth International Conference on Web Search and Data Mining*.
- [9] Xinyu Lin, Wenjie Wang, Yongqi Li, Fuli Feng, See-Kiong Ng, and Tat-Seng Chua. 2023. A Multi-facet Paradigm to Bridge Large Language Model and Recommendation. *arXiv:2310.06491*.
- [10] Xinyu Lin, Wenjie Wang, Yongqi Li, Shuo Yang, Fuli Feng, Yinwei Wei, and Tat-Seng Chua. 2024. Data-efficient Fine-tuning for LLM-based Recommendation. *arXiv:2401.17197*.
- [11] Junling Liu, Chao Liu, Renjie Lv, Kang Zhou, and Yan Zhang. 2023. Is ChatGPT a Good Recommender? A Preliminary Study.
- [12] Kai Mei and Yongfeng Zhang. 2023. LightLM: A Lightweight Deep and Narrow Language Model for Generative Recommendation. *arXiv:2310.17488*.
- [13] Aixun Sun. 2023. Take a Fresh Look at Recommender Systems from an Evaluation Standpoint. In *SIGIR*. ACM.
- [14] Zhiwen Tang, Hrishikesh Kulkarni, and Grace Hui Yang. 2021. High-Quality Diversification for Task-Oriented Dialogue Systems. In *ACL-IJCNLP 2021*.
- [15] Lei Wang and Ee-Peng Lim. 2023. Zero-Shot Next-Item Recommendation using Large Pretrained Language Models. *arXiv preprint arXiv:2304.03153* (2023).
- [16] Lei Wang, Jingsen Zhang, Xu Chen, Yankai Lin, Ruihua Song, Wayne Xin Zhao, and Ji-Rong Wen. 2023. When Large Language Model based Agent Meets User Behavior Analysis: A Novel User Simulation Paradigm. *arXiv:2306.02552*.
- [17] Wenjie Wang, Xinyu Lin, Fuli Feng, Xiangnan He, and Tat-Seng Chua. 2023. Generative Recommendation: Towards Next-generation Recommender Paradigm. *arXiv:2304.03516*.
- [18] Wenjie Wang, Yiyang Xu, Fuli Feng, Xinyu Lin, Xiangnan He, and Tat-Seng Chua. 2023. Diffusion Recommender Model. In *SIGIR*. ACM, 832–841.
- [19] Yunjia Xi, Weiwu Liu, Jianghao Lin, Jieming Zhu, Bo Chen, Ruiming Tang, Weinan Zhang, Rui Zhang, and Yong Yu. 2023. Towards Open-World Recommendation with Knowledge Augmentation from Large Language Models. *arXiv:2306.10933*.
- [20] Yiyang Xu, Wenjie Wang, Fuli Feng, Yunshan Ma, Jizhi Zhang, and Xiangnan He. 2024. DiFashion: Towards Personalized Outfit Generation. *arXiv:2402.17279*.
- [21] An Zhang, Leheng Sheng, Yuxin Chen, Hao Li, Yang Deng, Xiang Wang, and Tat-Seng Chua. 2023. On Generative Agents in Recommendation. *arXiv:2310.10108*.
- [22] Jizhi Zhang, Keqin Bao, Wenjie Wang, Yang Zhang, Wentao Shi, Wanhong Xu, Fuli Feng, and Tat-Seng Chua. 2024. Prospect Personalized Recommendation on Large Language Model-based Agent Platform. *arXiv:2402.18240* (2024).
- [23] Jizhi Zhang, Keqin Bao, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Is ChatGPT Fair for Recommendation? Evaluating Fairness in Large Language Model Recommendation. In *RecSys*. ACM.
- [24] Yuhui Zhang, Hao Ding, Zeren Shui, Yifei Ma, James Zou, Anoop Deoras, and Hao Wang. 2021. Language Models as Recommender Systems: Evaluations and Limitations. In *NeurIPS 2021 Workshop on I (Still) Can't Believe It's Not Better*.
- [25] Jujia Zhao, Wenjie Wang, Chen Xu, Zhaochun Ren, See-Kiong Ng, and Tat-Seng Chua. 2024. LLM-based Federated Recommendation. *arXiv:2402.09959*.
- [26] Peilin Zhou, Cao Meng, You-Liang Huang, Qichen Ye, Peiyang Zhang, Liu Junling, Yueqi Xie, Yining Hua, and Jaeboum Kim. 2023. Exploring Recommendation Capabilities of GPT-4V(ision): A Preliminary Case Study. *arXiv:2311.04199*.