EARL: The 2nd Workshop on Evaluating and Applying Recommender Systems with Large Language Models

Irene Li
University of Tokyo
Tokyo, Japan
ireneli@ds.itc.u-tokyo.ac.jp

Ruihai Dong
University College Dublin
Dublin, Ireland
ruihai.dong@ucd.ie

Guillaume Salha-Galvan
Shanghai Jiao Tong University
Shanghai, China
gsalhagalvan@sjtu.edu.cn

Aonghus Lawlor
University College Dublin
Dublin, Ireland
aonghus.lawlor@ucd.ie

Dairui Liu
University College Dublin
Dublin, Ireland
dairui.liu@ucd.ie

Lei Li
Hong Kong Baptist University
Hong Kong, China
csleili@comp.hkbu.edu.hk

Abstract

This article presents our proposal to organize the 2nd Workshop on Evaluating and Applying Recommender Systems with Large Language Models (EARL), to be held in conjunction with the 19th ACM Conference on Recommender Systems (RecSys 2025) in Prague, Czech Republic, in September 2025. Building on the success of the first EARL edition at RecSys 2024, we will foster dynamic and interactive discussions and debates on the application and evaluation of large language models (LLMs) in recommender systems (RSs). The workshop will explore emerging techniques such as retrieval-augmented generation (RAG), multi-modal recommendation, reinforcement learning with human feedback (RLHF), scalable fine-tuning methods, dynamic prompting strategies, and personalized conversational agents. In addition to highlighting key innovations and showcasing applications across diverse sectors, we will emphasize critical challenges in LLM-driven personalization, including bias, fairness, and transparency, which are essential for ensuring trustworthy and responsible RSs. Ultimately, the workshop aims to inspire new research directions and foster collaboration in the rapidly evolving field of LLM-powered RSs.

CCS Concepts

• **Information systems** → *Recommender systems*; *Personalization*; *Evaluation of retrieval results*; *Language models*.

Keywords

Large Language Models, Recommender Systems, Evaluation.

ACM Reference Format:
Irene Li, Ruihai Dong, Guillaume Salha-Galvan, Aonghus Lawlor, Dairui Liu, and Lei Li. 2025. EARL: The 2nd Workshop on Evaluating and Applying Recommender Systems with Large Language Models. In *19th ACM Conference on Recommender Systems (RecSys '25), September 22–26, 2025, 2025, Prague, Czech Republic.* ACM, New York, NY, USA, 5 pages. https://doi.org/XX.XXXX/XXXXXXX.XXXXXXX

RecSys '25, Prague, Czech Republic

ACM ISBN XXX-X-XXXX-XXXX-XX/XX/XX
https://doi.org/XX.XXXX/XXXXXXX.XXXXXXX

1 Workshop Overview and Rationale

Context. In recent years, the emergence of large language models (LLMs), such as GPT [1], LLaMA [12], and QWen [2], has showcased remarkable capabilities across diverse domains, ushering in a new era for natural language processing, computer vision, and beyond [4]. Their integration into recommender systems (RSs), in particular, has gained significant traction [6, 9, 10, 14, 15, 17, 18], with promising applications ranging from LLM-based embeddings for personalization [8] to agent-based methods for improving RSs [14] and fine-tuning or prompt engineering for LLM-driven recommendations [15]. The emergence of LLMs in RS research is increasingly reflected in the RecSys conference series, as evidenced by the two dedicated paper sessions on LLMs at RecSys 2024 [5]. Looking ahead, we anticipate a continued surge in studies that will not only introduce innovative features to RSs but also drive the development of next-generation RS technologies, shaping both academic research and industrial applications.

Workshop Overview. In this context, we successfully held the 1st EARL workshop at RecSys 2024 [7], attracting significant interest from both academia and industry. This half-day event featured engaging invited talks, high-quality paper presentations, and fostered vibrant discussions at the intersection of LLMs and RSs. Building on this success and the insights gained, we propose the organization of the 2nd EARL edition at RecSys 2025, aiming to expand both the scope and impact of the workshop.

For this second edition, we will again issue a call for long and short papers, aiming to broaden participation by encouraging submissions that explore emerging topics in LLM-powered RSs, key innovations and applications, and critical challenges in evaluation, fairness, trustworthiness, and responsibility. Beyond invited talks and paper presentations, we will introduce new interactive elements, including a best paper award and an interactive panel discussion. Overall, we aim to reinforce components that actively engage participants, fostering a dynamic and interactive forum for discussions on evaluating and applying RSs with LLMs.

Complementary with RecSys 2025. We anticipate that several LLM-related papers will be presented at RecSys 2025. However, this workshop is designed to complement rather than duplicate the main conference by fostering deeper discussions and debates on the challenges and opportunities of LLM-driven RSs. While main-track

papers will likely focus on algorithmic and performance advancements, we will welcome submissions that critically assess the application and evaluation of LLM-based RSs. We will also encourage submissions that discuss practical challenges in deploying existing LLMs. We will provide a platform for researchers whose work may be of interest but not yet fully mature for the main conference, offering them an opportunity to engage with the RecSys community. Overall, EARL is designed to be highly participatory. We will place a stronger emphasis on interactivity, debate, and inclusive discussions, aiming to foster critical reflections. We aim to propose an engaging event that promotes collaboration and sparks innovation.

2 Organizers

List of Organizers. This workshop will be organized by an experienced team of six researchers with expertise in RSs and LLMs, combining a strong track record in event coordination with active involvement in the RecSys community through previous publications, workshops, tutorials, and program committee roles:

- **Dr. Irene Li**[1] (*ireneli@ds.itc.u-tokyo.ac.jp*) is a Project Assistant Professor at the University of Tokyo. Her primary research focuses on natural language processing and artificial intelligence, with a particular emphasis on medical and clinical texts, alongside educational applications. She actively contributes to leading conferences such as ACL, NAACL, EMNLP, and AAAI through publications and program committee roles. Moreover, she has been the lead organizer of various online and in-person seminars for over three years, with one event drawing around 300 participants. She was among the co-authors of the **Best Student Paper at RecSys 2023** [16]. She was also the principal organizer of the first EARL workshop at RecSys 2024.
- **Dr. Ruihai Dong** (*ruihai.dong@ucd.ie*) is an Assistant Professor at University College Dublin. His research interests lie in machine learning, deep learning, and their applications to recommender systems and finance. He has published in top peer-reviewed journals and leading conferences such as WWW, RecSys, IUI, ACL, and IJCAI, and has served on the program committee for various conferences, including ACL, AAAI, ECML, and EMNLP. He was also a co-author of the **Best Student Paper at RecSys 2023** [16]. He was a founder and organizer of the Deep Learning meetup in Dublin and organized a series of events sponsored by multiple industry supporters, including Zalando, Accenture, and Deloitte.
- **Dr. Guillaume Salha-Galvan** (*gsalhagalvan@sjtu.edu.cn*) is an Associate Professor at Shanghai Jiao Tong University. He previously served as the Director of Machine Learning Research at Kibo Ryoku, and as a Research Scientist and Coordinator for Music Recommendation at Deezer. His research interests include graph learning, recommender systems, LLMs, and their music-related applications. He regularly publishes and serves as a reviewer in peer-reviewed journals and leading conferences such as ICML, KDD, WWW, IJCAI, and RecSys. In particular, he has co-authored eight articles in the past five RecSys editions, including works nominated for **Best Short Paper at RecSys 2020** [3], **Best Student Paper at RecSys 2021** [11], and **Best Full Paper at RecSys 2024** [13].
- **Dr. Aonghus Lawlor** (*aonghus.lawlor@ucd.ie*) is an Assistant Professor in the School of Computer at UCD and Funded Investigator in the Insight Centre for Data Analytics. The Insight Centre for Data Analytics is one of Europe's largest data analytics research organisations, with 400+researchers, more than 80 industry partners and over €100m in funding. Dr. Lawlor's current research interests are in recommender systems and machine/deep learning, and the application of computer vision for medical imaging. He has worked on several large scale industry projects and EU projects, developing novel recommender systems approaches to many application domains. This work has led to publications in top conferences and multiple patents.
- **Dr. Dairui Liu** (*dairui.liu@ucd.ie*) is a Post-doctoral Research Fellow at University College Dublin. His research focuses on natural language processing and recommender systems, specializing in explainable news recommendations. Recently, he has spearheaded the development of privacy-by-design and governed recommendation frameworks at Huawei Ireland Research Centre. He has published extensively in top-tier venues, including ACL, RecSys, CIKM, TORS, and ICCBR. He actively contributes to the academic community as a reviewer for prestigious conferences (e.g., ACL, RecSys) and journals (e.g., TOIS, TORS). He was also a co-author of the **Best Student Paper at RecSys 2023** [16].
- **Dr. Lei Li** (*csleili@comp.hkbu.edu.hk*) is a Postdoctoral Research Fellow at Hong Kong Baptist University, specializing in recommender systems and natural language processing. His recent research focuses on leveraging LLMs for recommendation systems, supported by the Hong Kong Research Grants Council (RGC) and Huawei. His work in this area includes LLM-based explainable recommendation, efficient LLMs for recommendation, and a survey on LLM-driven recommendation techniques. He has served as a Program Committee member for RecSys and WWW, a reviewer for leading journals such as TKDE, TOIS, TORS, and TBD, and a guest editor for a TORS special issue on "Large Language Models for Recommender Systems." Additionally, he co-presented a tutorial on LLM-based recommendation at RecSys 2023,titled "**Large Language Models for Recommendation**" [6], which attracted hundreds of attendees.

Remark: Why Six Organizers? We have chosen to involve six organizers to ensure a rigorous, well-rounded, and inclusive workshop. This larger team is essential not only to manage the anticipated increase in submissions and broader interest—driven by the growing prominence of LLMs in recommender systems—but also to handle the expanded responsibilities of event management. In addition to reviewing a diverse array of submissions from various RS subdomains, our team is tasked with overseeing the event website and actively seeking sponsorships, tasks that benefit from having more hands on deck. Our organizers bring a shared focus on LLMs and RSs, complemented by diverse expertise across both foundational research and applied domains such as finance, news, and music. This diversity strengthens the quality and breadth of

[1] **Dr. Irene Li** is the main point of contact for communications.

the reviewing process, ensuring a high-quality program elaboration. Moreover, we have deliberately prioritized broad geographic representation—organizers span three different countries—and a mix of academic and industry experience, further enhancing the workshop's relevance. Lastly, with only one organizer able to attend in person last year, our expanded team is strategically structured to ensure greater on-site presence at RecSys 2025, thereby improving logistical coordination, session moderation, and participant engagement.

3 Workshop Duration and Participant Numbers

Duration. We propose a **half-day** workshop, as for the first edition at RecSys 2024. The half-day format proved effective last year, allowing for a balanced program featuring presentations and discussions while maintaining engagement throughout the session.

Estimated Attendance. Based on our experience and the first EARL workshop, we anticipate receiving 15–25 submissions and expect an audience of **50–80 participants**. The growing research interest in LLMs for RSs, along with our expanded promotional efforts (see Section 8), is expected to drive strong participation.

4 Workshop Format

Format. The workshop will primarily take place **in person** at the RecSys 2025 venue in Prague, Czech Republic, in September 2025. However, we will support **remote participation** through livestreamed and recorded sessions with participant consent, ensuring broader accessibility. Social events will be held primarily in person, fostering networking opportunities among attendees.

Acknowledgment of Organizer Attendance. We acknowledge that at least one organizer will attend in person. In practice, we aim to have **the majority of the organizing team present** to facilitate smooth operations and enhance on-site engagement.

5 Workshop Activities

Activities from the First EARL Edition. We plan to follow the outline of the first EARL workshop, which consisted of a well-balanced mix of the following elements:

- **Invited Talks:** This year, the workshop will feature 2–3 invited talks delivered by distinguished experts from both academia and industry. They will provide cutting-edge insights into the integration of LLMs with RSs, covering recent advancements, emerging challenges, and future directions.
- **Accepted Contribution Talks:** The workshop will also include highlight talks from accepted papers (see Section 6), where authors will present their research.

In the first EARL edition, invited talks and accepted paper talks lasted 50 minutes and 20 minutes, respectively, including Q&A sessions. This year, the duration of each talk will be adjusted based on the total number of presentations, ensuring a well-balanced program that optimizes both content depth and audience engagement.

New Activities. To enhance engagement and interactivity, we will introduce the following additions to this year's program:

- **Best Paper Award:** To encourage high-quality submissions and recognize outstanding contributions, the workshop will present a best paper award. The selection process will involve an initial shortlisting based on reviewer scores, followed by a rigorous discussion among the organizers to ensure a fair and comprehensive evaluation. The award will highlight research that demonstrates originality, technical excellence, and significant contributions to LLM-based RSs.
- **Interactive Panel Discussion:** We will also introduce an interactive panel discussion featuring a diverse group of researchers from academia and industry. This session will explore key challenges in LLM-based RSs. To ensure audience involvement, we will collect questions from participants during the workshop registration process, allowing attendees to shape the discussion topics. Additionally, the panel will be highly interactive, with participants encouraged to contribute insights, ask follow-up questions, and engage in real-time discussions. This format will foster active participation, meaningful dialogue, and the exchange of diverse perspectives across the research and industry landscape.

6 Submission and Review Process

Call for Papers. As in last year's edition, we will invite submissions of original work in the following two formats:

- **Long Papers** (up to 8 pages, plus additional pages for references): intended for significant and enduring contributions.
- **Short Papers** (up to 4 pages, plus additional pages for references): suitable for new and promising work that has not yet reached the maturity required for a long paper.

Submissions will adhere to the ACM Proceedings template[2] using the double-column format. All submissions will be non-archival. We will especially encourage submissions from early-career researchers, whose work may be of interest but not mature enough for the RecSys 2025 main tracks. This workshop will allow them to engage with the RecSys community and receive valuable feedback.

Topics of Interest. Consistent with the rationale presented in Section 1, we will particularly welcome submissions on (1) emerging techniques for LLM-based RSs (including retrieval-augmented generation, multi-modal recommendation, reinforcement learning with human feedback, scalable fine-tuning methods, dynamic prompting strategies, and personalized conversational agents), (2) real-world applications of existing LLMs, and (3) critical challenges in ensuring the trustworthiness and responsibility of LLM-driven RSs. In detail, topics of interest include, but are not limited to, the following:

- Integrating LLMs to enhance RSs.
- Leveraging LLM-generated data to improve traditional RSs.
- LLM fine-tuning or prompt engineering techniques for RSs.
- Developing interactive and conversational RSs with LLMs.
- Integration of reinforcement learning with LLMs to adapt recommendations based on user feedback.
- Leveraging retrieval-augmented generation (RAG) to improve relevance and diversity in recommendations.
- Multi-modal RSs powered by LLMs, e.g., techniques integrating text, images, and audio data.
- Few-shot and zero-shot learning for LLM-based RSs.

[2] https://www.acm.org/publications/proceedings-template

- Personalization strategies for LLM-powered RSs, including dynamic user modeling and real-time adaptation.
- Cross-domain and cross-lingual RSs utilizing multilingual and generalist LLMs.
- Applications of LLM-enhanced RSs in domains such as finance, streaming platforms, and social networks.
- Scalability challenges in deploying LLM-powered RSs.
- Efficiency challenges in deploying LLM-powered RSs.
- Evaluation of LLM-powered RSs using novel metrics/standards.
- Evaluation of LLM-powered RSs using human feedback.
- Enhancing transparency in LLM-powered RSs.
- Enhancing fairness in LLM-powered RSs.
- Enhancing explainability in LLM-powered RSs.
- Trustworthy recommendation with LLMs, addressing bias, safety, privacy, and authenticity issues.
- Responsible AI practices in LLM-powered RSs, emphasizing ethical considerations and sustainable AI.

Peer-Review Process. Submitted papers will undergo a **double-blind** peer-review process and will be evaluated by at least three reviewers from our Program Committee. Submissions will be evaluated based on academic soundness, relevance to the workshop themes, and overall interest to the RecSys community.

Program Committee. The workshop is supported by a high-quality Program Committee with expertise in RSs and LLMs, which will be finalized in the coming weeks and already includes:

- Aonghus Lawlor (UCD, Ireland)
- Neil Hurley (UCD, Ireland)
- Elias Tragos (Insight Centre for Data Analytics, Ireland)
- Hang Jiang (MIT, USA)
- Sixun Ouyang (Manulife Singapore Pte. Ltd., Singapore)
- Dairui Liu (UCD, Ireland)
- Honghui Du (UCD, Ireland)
- Zheng Ju (UCD, Ireland)
- Qin Ruan (UCD, Ireland)
- Aayush Singha Roy (UCD, Ireland)
- Boming Yang (UTokyo, Japan)
- Daxiang Dong (Baidu Inc, China)
- Yuang Jiang (Meta Inc, USA)
- Rui Yang (NUS, Singapore)
- Tianwei She (Smartor, Japan)
- Jinghui Lu (Smartor, Japan)
- Bichen Shi (Huawei Ireland Research Center, Ireland)
- Xingsheng Guo (Huawei Ireland Research Center, Ireland)
- Michael O'Mahony (UCD, Ireland)
- Xingyu Song (UTokyo, Japan)
- Jiaying Xu (Kibo Ryoku, France)
- Guillaume Salha-Galvan (SJTU, China)

7 Participant Selection

Drawing on the experience and established networks of our organizing team within the RecSys and AI communities, we will gather experts in the relevant domains of this workshop for both the **Interactive Panel Discussion** and **Invited Talks**. Invitations will be extended based on expertise and relevance to the workshop themes. Importantly, we will strive to ensure diversity and representation across multiple dimensions, including geographic distribution, gender balance, and a mix of academic and industry perspectives. A list of potential invitees is currently being compiled, and invitations will be sent upon the acceptance of this workshop proposal.

8 Promotion and Dissemination

Workshop Website. As we did last year[3], we will maintain a public website for this second EARL edition. This website will serve as the central hub for all workshop-related information, including the call for papers, submission guidelines, speaker announcements, program details, and accepted papers. By providing a structured and accessible online presence, we aim to maximize visibility and facilitate engagement from prospective participants.

Promotion on Social Media and Professional Networks. To effectively promote the workshop, we will leverage social media platforms, including LinkedIn, Bluesky, and X, to engage professionals and academics. We will enhance visibility through targeted hashtags and interactive engagement initiatives. We will extend our reach through targeted advertising and cross-promotion within the networks of organizers, speakers, and affiliated institutions.

Sponsorships. Last year, we successfully secured over 500 euros in sponsorship from industry partners. Building on this success, we will actively seek new sponsorships from companies and organizations with interests in LLMs. Currently, the Insight Center [4] has agreed to sponsor us if the proposal can be approved. Sponsors will receive brand visibility through logo placement, acknowledgments during the event, and promotion in workshop materials and across social media channels.

Dissemination. Accepted papers and presentation materials will be made available on the workshop website. We are also considering making the workshop recordings publicly available, enabling post-event access and broader diffusion of the content. This initiative will allow participants, including those unable to attend in person, to engage asynchronously. Invited speakers and authors will benefit from extended visibility, as they will have the opportunity to share their presentations and insights with a wider audience. Recordings will be shared only upon consent from each participant.

9 Workshop History

EARL 2024. The first edition of the EARL workshop, held at RecSys 2024 in Bari, Italy, attracted strong interest from both academia and industry. This successful half-day event featured two invited talks by Prof. Scott Sanner (University of Toronto) and Prof. Yongfeng Zhang (Rutgers University), as well as four accepted paper presentations, fostering high-quality and dynamic discussions on the integration of LLMs in RSs. The detailed program is available online[3].

Insights and Improvements. We aim to refine this edition by boosting promotion, aligning with the main conference, and ensuring on-site coordination. New features include a best paper award to recognize top research and an interactive panel uniting experts to discuss challenges and future directions of LLM-powered RSs, echoing RecSys 2025's focus on engaging workshops.

[3] Website of the first EARL workshop at RecSys 2024: https://earl-workshop.github.io
[4] https://www.insight-centre.org/

References

[1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774* (2023).

[2] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen Technical Report. *arXiv preprint arXiv:2309.16609* (2023).

[3] Walid Bendada, Guillaume Salha, and Théo Bontempelli. 2020. Carousel Personalization in Music Streaming Apps with Contextual Bandits. In *Proceedings of the 14th ACM Conference on Recommender Systems*. 420–425.

[4] Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2024. A Survey on Evaluation of Large Language Models. *ACM Transactions on Intelligent Systems and Technology* 15, 3 (2024), 1–45.

[5] Tommaso Di Noia, Pasquale Lops, Thorsten Joachims, Katrien Verbert, Pablo Castells, Zhenhua Dong, and Ben London. 2024. RecSys '24: Proceedings of the 18th ACM Conference on Recommender Systems.

[6] Wenyue Hua, Lei Li, Shuyuan Xu, Li Chen, and Yongfeng Zhang. 2023. Tutorial on Large Language Models for Recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 1281–1283.

[7] Irene Li, Ruihai Dong, Lei Li, and Li Chen. 2024. EARL: Workshop on Evaluating and Applying Recommendation Systems with Large Language Models. In *Proceedings of the 18th ACM Conference on Recommender Systems*. 1262–1264.

[8] Zheng Li and Kai Zhange. 2024. Personalized News Recommendation System via LLM Embedding and Co-Occurrence Patterns. *arXiv preprint arXiv:2411.06046* (2024).

[9] Jianghao Lin, Xinyi Dai, Yunjia Xi, Weiwen Liu, Bo Chen, Hao Zhang, Yong Liu, Chuhan Wu, Xiangyang Li, Chenxu Zhu, et al. 2025. How Can Recommender Systems Benefit from Large Language Models: A Survey. *ACM Transactions on Information Systems* 43, 2 (2025), 1–47.

[10] Dairui Liu, Boming Yang, Honghui Du, Derek Greene, Aonghus Lawlor, Ruihai Dong, and Irene Li. 2023. Recprompt: A Prompt Tuning Framework for News Recommendation Using Large Language Models. *arXiv preprint arXiv:2312.10463* (2023).

[11] Guillaume Salha-Galvan, Romain Hennequin, Benjamin Chapus, Viet-Anh Tran, and Michalis Vazirgiannis. 2021. Cold Start Similar Artists Ranking with Gravity-Inspired Graph Autoencoders. In *Proceedings of the 15th ACM Conference on Recommender Systems*. 443–452.

[12] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971* (2023).

[13] Viet-Anh Tran, Guillaume Salha-Galvan, Bruno Sguerra, and Romain Hennequin. 2024. Transformers Meet ACT-R: Repeat-Aware and Sequential Listening Session Recommendation. In *Proceedings of the 18th ACM Conference on Recommender Systems*. 486–496.

[14] Lei Wang, Jingsen Zhang, Hao Yang, Zhi-Yuan Chen, Jiakai Tang, Zeyu Zhang, Xu Chen, Yankai Lin, Hao Sun, Ruihua Song, Wayne Xin Zhao, Jun Xu, Zhicheng Dou, Jun Wang, and Ji-Rong Wen. 2024. User Behavior Simulation with Large Language Model-based Agents for Recommender Systems. *ACM Transactions on Information Systems* (2024).

[15] Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, et al. 2024. A Survey on Large Language Models for Recommendation. *World Wide Web* 27, 5 (2024), 60.

[16] Boming Yang, Dairui Liu, Toyotaro Suzumura, Ruihai Dong, and Irene Li. 2023. Going Beyond Local: Global Graph-Enhanced Personalized News Recommendations. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 24–34.

[17] 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 *Proceedings of the 17th ACM Conference on Recommender Systems*. 993–999.

[18] Zihuai Zhao, Wenqi Fan, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Zhen Wen, Fei Wang, Xiangyu Zhao, Jiliang Tang, et al. 2024. Recommender Systems in the Era of Large Language Models (LLMs). *IEEE Transactions on Knowledge and Data Engineering* (2024).