EARL: The 3rd Workshop on Evaluating and Applying Recommender Systems with Large Language Models

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

Irene Li
University of Tokyo
Tokyo, Japan
irene.li@weblab.t.u-tokyo.ac.jp

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

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

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

Sairamvinay Vijayaraghavan
TikTok & University of California, Davis
Davis, United States of America
sairam16898@gmail.com

Lei Li
Shenzhen University
Shenzhen, China
lilei@szu.edu.cn

Abstract

This proposal outlines the organization of the 3rd Workshop on Evaluating and Applying Recommender Systems with Large Language Models (EARL), to be co-located with the 35th International ACM Conference on Knowledge and Information Management (CIKM 2026) in Rome, Italy. Following the strong engagement of the first two EARL editions, this third installment aims to solidify the workshop as a premier venue for advancing research at the intersection of large language models (LLMs) and recommender systems (RSs). EARL 2026 will foster discussions on the evaluation and practical deployment of LLM-driven RSs, emphasizing not only methodological soundness and reproducibility, but also the responsible integration of LLMs into real-world systems. By bringing together researchers and practitioners from academia and industry, the workshop will stimulate debate, encourage critical reflection, and shape future research directions in LLM-enhanced RSs.

CCS Concepts

• **Information systems** → *Recommender systems*; *Personalization*.

Keywords

Large Language Models, Recommender Systems, Evaluation.

ACM Reference Format:
Guillaume Salha-Galvan, Irene Li, Ruihai Dong, Aonghus Lawlor, Dairui Liu, Sairamvinay Vijayaraghavan, and Lei Li. 2026. EARL: The 3rd Workshop on Evaluating and Applying Recommender Systems with Large Language Models. In *Proceedings of the Thirty-Fifth International ACM Conference on Knowledge and Information Management (CIKM '26), November 7–11, 2026, Rome, Italy*. ACM, New York, NY, USA, 4 pages. https://doi.org/TBD

CIKM '26, Rome, Italy

ACM ISBN TBD
https://doi.org/TBD

1 Workshop Overview and Rationale

Context. In recent years, large language models (LLMs) such as GPT, Gemini, and DeepSeek have demonstrated remarkable capabilities across a wide range of tasks, marking a transformative shift in artificial intelligence (AI) [1, 5, 11, 17]. Originally advancing natural language processing (NLP), they have rapidly extended their influence to areas like computer vision and complex reasoning [3, 14].

Within recommender systems (RSs), the adoption of LLMs has also accelerated significantly. Researchers and practitioners are exploring diverse integration paradigms, including leveraging LLM-derived embeddings for personalization, using LLMs as content annotators, incorporating agent-based frameworks for interactive recommendation workflows, and applying prompt engineering techniques to generate recommendations [4, 6, 10, 12, 15, 18, 19]. These approaches suggest that LLMs are not merely auxiliary tools but are reshaping the foundations of modern RS research.

The growing prominence of LLM-enhanced RSs is increasingly visible in conferences such as CIKM, RecSys, and KDD, where dedicated sessions and rising submission numbers highlight strong interest [2, 16]. We expect continued expansion of this research at CIKM 2026, with LLM-driven approaches contributing to the evolution of RSs in both academic settings and industrial applications.

Workshop Overview. In this context, we successfully organized the first two EARL workshops at RecSys 2024 and 2025, attracting strong community interest [8, 9]. These events featured engaging invited talks, oral paper presentations, posters, and lively discussions at the intersection of LLMs and RSs. Building on this momentum, we propose organizing the third EARL workshop at CIKM 2026.

For this third edition, we will continue featuring invited talks and soliciting paper submissions on emerging directions and pressing challenges in LLM-enhanced RSs (see Section 5). Beyond oral talks and an interactive poster session, we will maintain a Best Paper Award to recognize outstanding contributions and encourage high-quality submissions. Overall, EARL 2026 seeks to foster continued community exchange and establish a rigorous, forward-looking forum dedicated to the evaluation and application of RSs with LLMs.

Complementarity with CIKM. We anticipate that LLM papers will appear in the CIKM 2026 main track. EARL is intentionally positioned as a complementary venue, focusing specifically on how LLMs are evaluated and applied within RSs, while connecting this topic to core CIKM interests such as information retrieval, knowledge management, user modeling, and data-driven intelligent systems. Moreover, while main-track contributions typically prioritize algorithmic novelty, EARL will provide a dedicated space for critical reflection, deployment considerations, and debates on the responsible integration of LLMs into RSs. EARL will also welcome more exploratory contributions that raise important questions.

EARL also distinguishes itself through continuity and community-building. Having been successfully organized for two consecutive years, the workshop has established a stable organizing team and Program Committee with domain expertise spanning both RSs and LLMs. This continuity enables consistent quality standards, institutional memory, and sustained thematic focus, positioning EARL as a mature and evolving forum rather than a one-off event.

2 Organizers

List of Organizers. The workshop will be organized by seven researchers with established expertise in RSs, LLMs, and academic event organization. Most of the organizers were involved in previous editions, ensuring continuity and accumulated experience:

- **Dr. Guillaume Salha-Galvan**[1] (*gsalhagalvan@sjtu.edu.cn*) is an Associate Professor at Shanghai Jiao Tong University. He has a decade of industry experience, including as Director of Research at Kibo Ryoku and Research Coordinator for Music Recommendation at Deezer. His research interests include RSs, LLMs, and their real-world applications. He regularly publishes and reviews for conferences such as ICML, KDD, WWW, RecSys, and CIKM. He was the main on-site host for the second EARL edition at RecSys 2025.
- **Dr. Irene Li** (*irene.li@weblab.t.u-tokyo.ac.jp*) is a Project Lecturer at the University of Tokyo. Her research centers on NLP and AI, with a focus on medical and educational applications. She actively contributes to top venues such as ACL, NAACL, EMNLP, AAAI, and CIKM through publications, Program Committee (PC), and Area Chair (AC) service. She organized multiple academic seminars with up to 300 participants, and was the principal organizer of the first two EARL editions.
- **Dr. Ruihai Dong** (*ruihai.dong@ucd.ie*) is an Associate Professor at University College Dublin. His research focuses on machine/deep learning and applications to RSs and finance. He publishes and serves on PCs for venues such as WWW, AAAI, RecSys, ACL, and CIKM. He founded the Dublin Deep Learning meetup and organized events sponsored by industry supporters including Zalando, Accenture, and Deloitte. He also co-organized both previous EARL editions.
- **Dr. Aonghus Lawlor** (*aonghus.lawlor@ucd.ie*) is an Assistant Professor at University College Dublin and a Funded Investigator at the Insight Centre for Data Analytics. His research interests include machine/deep learning and RSs. Through various large-scale industry and EU projects, he has developed novel RS approaches across multiple domains, leading to top conference publications and several patents. He co-hosted the second EARL edition on-site at RecSys 2025.
- **Dr. Dairui Liu** (*dairui.liu@ucd.ie*) is a Postdoctoral Research Fellow at University College Dublin. His research focuses on NLP and RSs, specializing in explainable news recommendations. Recently, he developed privacy-by-design and governed RS frameworks at the Huawei Ireland Research Centre. He publishes and reviews for top-tier journal and conferences including ACL, RecSys, CIKM, and TORS, and co-organized both previous EARL editions.
- **Dr. Sairamvinay Vijayaraghavan** (*sairam16898@gmail.com*) is a Machine Learning Engineer at TikTok. Bringing an additional industry perspective to the team, he also holds a Ph.D. from the University of California, Davis, and publishes in top venues like RecSys. His research centers on LLMs and RSs, specifically studying the reliability of explanations within RSs and how LLMs can be leveraged for this purpose.
- **Dr. Lei Li** (*lilei@szu.edu.cn*) is an Assistant Professor at Shenzhen University, specializing in RSs and NLP. His research focuses on leveraging LLMs for RSs, including explainable and efficient recommendation, etc. He has served as a PC member for KDD, RecSys, WWW, and CIKM, a reviewer for journals such as TKDE, TOIS, and TORS, and a guest editor for a TORS special issue on LLMs for RSs. He co-presented a tutorial on LLM-based recommendation at RecSys 2023 which attracted hundreds of attendees [7], and co-organized both previous EARL editions.

Organizer Roles. We intentionally formed a large organizing team to ensure operational robustness and program quality. As EARL has grown over past editions, the workload has expanded. Our collective, complementary expertise across RSs and LLMs also strengthens the reviewing process. Finally, distributing responsibilities across seven organizers improves resilience to unforeseen constraints, reflecting EARL's maturity as an established workshop.

In practice, **Dr. Salha-Galvan** will be the main contact for the organization of EARL 2026. He will attend CIKM 2026 and handle on-site logistics alongside **Dr. Dong** and **Dr. Lawlor**. **Dr. Liu** will be responsible for the website and online presence of the workshop. **Dr. Irene Li**, **Dr. Lei Li**, and **Dr. Vijayaraghavan** will primarily supervise the talk and paper selection process, although all seven organizers will actively participate in reviews.

3 Workshop Organization

Format. EARL 2026 will be conducted **in person** at CIKM 2026 in Rome, Italy. To enhance accessibility, we will support **remote participation** via Zoom (subject to approval by the CIKM chairs). The poster session will be held in person.

Duration. We propose a **half-day** workshop, consistent with the previous editions. This format allows for a focused yet balanced program combining oral talks, a poster session, and interactive discussions, while maintaining strong participant engagement.

Estimated Attendance. Based on past editions, we anticipate around 30 submissions and **60–80 participants**. Given the growing interest in LLMs for RSs, we expect sustained and potentially increasing engagement from both academia and industry.

[1] **Dr. Guillaume Salha-Galvan** is the main point of contact for communications.

4 Workshop Program

We will maintain the format established in previous years, which has proven effective in balancing technical depth with interactive exchange. The workshop will consist of the following components:

- **Invited Talks:** The workshop will feature 1–2 invited talks delivered by leading experts, providing high-level perspectives on recent advances and emerging challenges.
- **Highlight Oral Presentations:** A few accepted papers (see the call for papers in Section 5) will be selected for oral presentation based on reviews and subsequent deliberation.
- **Interactive Poster Session:** The remaining accepted papers will be presented during a poster session. Last year, this session proved to be a key component for fostering interactivity and exchange among attendees.
- **Best Paper Award and Runner-Up Recognition:** We will continue the Best Paper Award introduced in 2025, along with two Best Paper Runner-Up recognitions. The selection process will be based on reviews, followed by deliberations among the organizers to ensure fairness and rigor. These awards aim to recognize originality, technical excellence, and impactful contributions to LLM-based RSs.

5 Submission and Participation Process

Call for Papers. We will invite submissions in 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 references): Suitable for promising work not yet mature enough for a long paper.

Submissions will adhere to the ACM Proceedings template using the double-column format. All accepted papers will be non-archival. We will particularly welcome contributions from early-career researchers and students. The workshop will provide an opportunity for them to gain visibility and receive constructive feedback.

Topics of Interest. Consistent with the rationale presented in Section 1, EARL 2026 will particularly welcome submissions on the evaluation and application of LLMs within recommender systems, especially work that connects recommendation with core CIKM themes such as information retrieval, knowledge representation, knowledge management, user modeling, personalization, and data-driven intelligent systems. Topics of interest include retrieval-augmented and knowledge-grounded recommendation, conversational and multimodal recommendation, and efficient deployment of LLM-enhanced RSs. We also encourage work on evaluation, reproducibility, fairness, safety, and real-world deployment.

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

Program Committee. EARL 2025 was supported by a high-quality PC comprising **twenty** external reviewers in addition to all organizers. This included a diverse mix of experts from both academia (e.g., University of Tokyo, UC Dublin) and industry (e.g., Meta, Xiaomi). The complete PC from last year is available online[2].

[2] Website of EARL 2025: https://earl-workshop.github.io

This year, we will continue to benefit from a similar high-quality team. A **majority** of last year's PC members have already agreed to participate again, and we expect to attract additional reviewers. The final PC composition will be finalized in the coming weeks.

6 Promotion and Dissemination

Workshop Website. As in previous editions, we will maintain an official public website[2] for the EARL workshop. This website will serve as the central hub for all workshop-related information, including the call for papers, submission instructions, program details, accepted contributions, and other news.

Promotion on Social Media and Professional Networks. We will continue to promote the workshop through academic and professional networks, including LinkedIn, X, and Bluesky, leveraging the combined reach of the organizers, speakers, and affiliated institutions. Coordinated announcements will help sustain broad visibility and engagement from both academia and industry.

Sponsorship and Community Support. Building on the success of EARL 2025, which secured three sponsors, we will seek continued and new support for this upcoming edition. Sponsors will receive visibility through logo placement on the website, acknowledgments during the event, and inclusion in promotional materials. Additionally, we will continue to print best paper certificates and prepare workshop souvenirs (e.g., tote bags) to foster a welcoming environment and engance community engagement.

Dissemination. Accepted papers will be made publicly available on the website. Subject to speaker consent, we will also provide access to workshop recordings to facilitate broader dissemination and asynchronous engagement. This approach ensures extended visibility for authors and invited speakers while allowing those unable to attend in person to benefit from the workshop content.

7 Workshop History

Previous Editions. EARL 2024, held at RecSys 2024 in Bari, Italy, featured invited talks by Scott Sanner from the University of Toronto and Yongfeng Zhang from Rutgers University, along with four accepted paper presentations from eight submissions. EARL 2025, held at RecSys 2025 in Prague, Czech Republic, featured invited talks by Wei-Wei Du and Tommaso Carraro from Sony, and Wenyue Hua from Microsoft Research. This edition saw a significant increase in submissions, with **27 submissions**. Of these, **18 papers** were accepted, including **6 for highlight talks** and **12 for posters**. The submissions came from both industry (e.g., Google, Pinterest, Spotify) and academia (e.g., University of Amsterdam, Sorbonne University). The Best Paper Award was introduced for the first time and presented to Palumbo et al. [13] from Spotify.

Insights. A key insight from the second edition was that submissions more than tripled, leading us to organize a poster session, while all papers were presented as oral talks at EARL 2024. This year, we anticipate similar submission numbers and will keep the poster session in the program. Additionally, we will reinforce our PC to ensure a smooth reviewing process.

References

[1] 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).

[2] Maria Bielikova, Pavel Kordik, Markus Schedl, et al. 2025. RecSys '25: Proceedings of the 19th ACM Conference on Recommender Systems. *Association for Computing Machinery* (2025).

[3] 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.

[4] Enzo Charolois-Pasqua, Eléa Vellard, Youssra Rebboud, Pasquale Lisena, and Raphaël Troncy. 2025. A Language Model-Based Playlist Generation Recommender System. In *Proceedings of the 19th ACM Conference on Recommender Systems*. 1–11.

[5] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. 2025. Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities. *arXiv preprint arXiv:2507.06261* (2025).

[6] Mathieu Delcluze, Léa Briand, Benjamin Chapus, Deniz Mekik, and Guillaume Salha-Galvan. 2026. Music Playlist Captioning at Scale with Large Language Models. In *Proceedings of the 2026 European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*.

[7] 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.

[8] 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.

[9] 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 *Proceedings of the 19th ACM Conference on Recommender Systems*. 1366–1370.

[10] 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.

[11] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. DeepSeek-V3 Technical Report. *arXiv preprint arXiv:2412.19437* (2024).

[12] Dairui Liu, Boming Yang, Honghui Du, Derek Greene, Neil Hurley, Aonghus Lawlor, Ruihai Dong, and Irene Li. 2024. RecPrompt: A Self-tuning Prompting Framework for News Recommendation Using Large Language Models. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*. 3902–3906.

[13] Enrico Palumbo, Gustavo Penha, Alva Liu, Marcus Eltscheminov, Jefferson Carvalho dos Santos, Alice Wang, Hugues Bouchard, Humberto Jesús Corona Pampin, and Michelle Tran Luu. 2025. AudioBoost: Increasing Audiobook Retrievability in Spotify Search with Synthetic Query Generation. In *2nd EARL Workshop on Evaluating and Applying Recommender Systems with Large Language Model*.

[14] Dhruv Shah, Blazej Osinski, Brian Ichter, and Sergey Levine. 2022. LM-Nav: Robotic Navigation with Large Pre-Trained Models of Language, Vision, and Action. In *Proceedings of the 6th Annual Conference on Robot Learning*. 492–504.

[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] Carl Yang, Senjuti Basu Roy, Zhenhui Jessie Li, Jaap Kamps, Meeyoung Cha, Chanyoung Park, and Noseong Park. 2025. Report on the 34th ACM Conference on Information and Knowledge Management (CIKM 2025). In *ACM SIGIR Forum*, Vol. 59. ACM New York, NY, USA, 1–14.

[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] Yubin Zhang, Haiming Xu, Guillaume Salha-Galvan, Ruiyan Han, Feiyang Xiao, Yanhua Huang, Li Lin, Yang Luo, and Yao Hu. 2026. IDProxy: Cold-Start CTR Prediction for Ads and Recommendation at Xiaohongshu with Multimodal LLMs. *arXiv preprint arXiv:2603.01590* (2026).

[19] 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).