

LLM-UP: SIGIR 2026 Workshop on LLM-powered User Profiling for Search and Recommendation

Hongzhi Yin
The University of Queensland
Brisbane, QLD, Australia
h.yin1@uq.edu.au

Joel Mackenzie
The University of Queensland
Brisbane, QLD, Australia
joel.mackenzie@uq.edu.au

Wei Yuan
The University of Queensland
Brisbane, QLD, Australia
w.yuan@uq.edu.au

Quoc Viet Hung Nguyen
Griffith University
Gold Coast, QLD, Australia
henry.nguyen@griffith.edu.au

Yi Zhang
Anhui University
Hefei, Anhui, China
zhangyi.ahu@gmail.com

Wayne Xin Zhao
Gaoling School of Artificial
Intelligence, Renmin University of
China
Beijing, China
batmanfly@gmail.com

Yong Li
Department of Electronic
Engineering, Tsinghua University
Beijing, China
liyong07@tsinghua.edu.cn

Lina Yao
University of New South Wales
Sydney, NSW
lina.yao@unsw.edu.au

Abstract

The rapid advancement of large language models (LLMs) has opened new possibilities for understanding users in search and recommendation. While traditional behavior-based or feature-driven user models rely primarily on explicit interactions or handcrafted representations, LLMs introduce a fundamentally different paradigm: LLM-powered user profiling, where user preferences, intents, and contextual attributes can be extracted, summarized, or reasoned about directly through natural language. This shift unlocks powerful new paths to achieve personalization but also raises pressing questions related to modeling fidelity, temporal dynamics, evaluation methodology, privacy, and responsible deployment. The LLM-UP workshop aims to bring together researchers and practitioners to systematize emerging progress in LLM-powered user profiling, identify open challenges, and explore opportunities for integrating such techniques into search and recommendation pipelines. The LLM-UP workshop adopts an interactive structure featuring lightning talks, panel discussions, and paper presentations to foster active engagement, cross-disciplinary dialogue, and community-driven agenda setting for this rapidly evolving field.

CCS Concepts

• Information systems → Recommender systems.

Keywords

Large Language Models, User Profiling, Personalization, Recommendation, Search

ACM Reference Format:

Hongzhi Yin, Wei Yuan, Yi Zhang, Joel Mackenzie, Quoc Viet Hung Nguyen, Wayne Xin Zhao, Yong Li, and Lina Yao. 2026. LLM-UP: SIGIR 2026 Workshop on LLM-powered User Profiling for Search and Recommendation. In *Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26)*, July 20–24, 2026, Melbourne, VIC, Australia. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3805712.3808652>

1 Motivation for the Workshop

Personalization has long been a core component of information retrieval (IR) and recommender systems (RecSys), where the primary objective is to construct reliable *user profiles* that enable systems to tailor their outputs to individual preferences. Traditional user profiling strategies, built on historical interactions, engineered features, or collaborative patterns, have supported a wide range of effective IR and RecSys applications [1, 21]. However, these approaches often struggle to capture underlying user intent, provide fine-grained semantic understanding and explanation, adapt to evolving user needs, and cope with data scarcity [10, 12]. With the rapid rise of large language models (LLMs), a new paradigm has emerged: *LLM-powered user profiling*, in which sparse behavioral traces, contextual cues, and latent preferences can be inferred, summarized, and reasoned about directly from natural language and heterogeneous signals [14, 18, 19].

The shift toward LLM-powered user profiling offers unprecedented opportunities for advancing personalized IR and recommendation [11, 17]. Yet this emerging paradigm simultaneously introduces a range of technical and conceptual challenges. A primary challenge lies in how to steer LLMs to construct trustworthy user profiles. For example, generative inference may produce unstable or inconsistent profiles [9]; model-driven reasoning can embed hidden biases [22]; reliance on raw interactions may amplify popularity-driven effect [7]; and long-term memory mechanisms



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGIR '26, Melbourne, VIC, Australia*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2599-9/2026/07
<https://doi.org/10.1145/3805712.3808652>

raise questions around temporal drift, preference forgetting, and representational validity [4]. Moreover, evaluating LLM-powered user profiles requires new methodological thinking, as conventional accuracy-based metrics do not fully capture subjective or contextual personalization quality [15]. Finally, the adaptive alignment and integration of user profiles with downstream models, along with persistent concerns around fairness and privacy, continue to pose significant open challenges [7, 19].

Despite the rapidly growing interest, research on LLM-powered user profiling in search and recommendation remains highly fragmented. Work on preference extraction, LLM memory management, personalization prompts, retrieval-generation integration, user simulation, privacy, and evaluation have made progress in the pure NLP and generative AI community [3, 5, 20], however, they still proceed independently with limited cross-fertilization in IR and RecSys communities. This fragmentation makes it difficult to establish common frameworks, shared resources, or systematic comparisons, ultimately slowing progress in an area that is becoming increasingly central to real-world search and RecSys applications.

Given these developments, there is an urgent need for a dedicated venue where researchers and practitioners can consolidate recent advances, clarify open questions, and develop a coherent research agenda for LLM-powered user profiling in the IR and RecSys areas. While SIGIR features strong tracks on retrieval, ranking, and recommendation, it currently lacks a unified forum focused specifically on the foundations, methodologies, challenges, and opportunities of LLM-powered profiling. The LLM-UP workshop seeks to fill this gap by fostering cross-disciplinary dialogue among IR, RecSys, and NLP communities, encouraging the exchange of insights across traditionally separate lines of work.

Ultimately, this workshop aims to catalyze a deeper understanding of how LLMs can transform user profiling and personalization in search and recommendation. By bringing together diverse perspectives, the workshop will help establish shared principles, identify actionable research directions, and foster a collaborative community committed to advancing robust, responsible, and user-centric LLM-powered profiling for next-generation IR and RecSys.

2 Theme and Purpose

The LLM-UP workshop is centered on understanding the challenges and risks and advancing the foundations, methodologies, and applications of integrating LLM-powered user profiling in IR and RecSys. The purpose includes: (1) To bring together researchers and practitioners working on improving search and recommendation via LLM-powered user profiling; (2) To develop a systematic and comprehensive understanding of emerging methods, underlying principles, and key technical challenges in LLM-powered user profiling; (3) To chart a research agenda that integrates technical advances with considerations of reliability, fairness, and responsible deployment for LLM-powered personalization search and recommendation. To support this purpose, the LLM-UP workshop is organized around four major themes outlined below.

Theme 1. Modeling and Representing User Profiles with LLMs. This theme explores the design and representation of LLM-powered user profiles, including:

- User profile extraction, analysis, and summarization from document content, interaction histories, contextual cues, multimodal data, and external knowledge.
- LLM-based reasoning about user needs, attributes, latent intents, and contextual constraints, enabling richer interpretations of user behavior and more adaptive personalization across diverse scenarios.
- Temporal and dynamic user profiling, addressing preference drift, sequential memory updates, evolving user needs, and long-term personalization.
- Designing efficient and cost-effective user profiling methods that can operate at scale and support large user populations in real-world IR and recommendation settings.

Theme 2. Integrating User Profiles into Search and Recommendation Pipelines. This theme investigates how LLM-powered user profiles can be incorporated into ranking pipelines, conversational agents, and hybrid retrieval-generation systems, including:

- Representation learning for LLM-powered user profiles, alignment with downstream models, and integration for feature augmentation.
- Incorporating user profiles into LLMs through fine-tuning or adaptation, enabling generative search, recommendation, and personalized response generation.
- Retrieval-augmented generation, using LLM-powered user profiles to guide personalized retrieval, relevance estimation, and generative response synthesis.
- Designing interactive and multi-agent-based system architectures that leverage user profiles to enable LLMs to act on behalf of users in refining queries and improving recommendations.

Theme 3. Evaluating the Fidelity, Quality, and Utility of LLM-powered Profiles. LLM-powered user profiles are generative, subjective, and highly context-dependent, posing unique challenges for evaluation. This theme focuses on developing appropriate methodologies for assessing their reliability and effectiveness, such as:

- Profile fidelity metrics, assessing the accuracy, stability, and consistency of inferred preferences, as well as the robustness of user profiles across different contexts and model conditions.
- Task-based evaluation, measuring how user profiles enhance retrieval, ranking, recommendation, and conversational outcomes in downstream IR and RecSys tasks.
- Human-centered evaluation, capturing user satisfaction, perceived relevance, and the quality of personalized experiences.
- Dataset and benchmark development, including synthetic user construction, controlled preference-tracking tasks, and multi-turn evaluation settings for assessing personalization.

Theme 4. Responsible, Fair, and Privacy-Preserving User Profiling. LLM-powered user profiling may amplify ethical concerns surrounding the interpretation and use of user data. This theme addresses critical issues related to:

- Debiasing and fairness, ensuring that user profiling methods mitigate the encoding or amplification of sensitive attributes, demographic stereotypes, and broader systemic inequities.
- Privacy protection and data governance, balancing the benefits of personalization with safeguards against unintended inference and unauthorized use of sensitive information.

- Transparency and explainability, enabling users and practitioners to understand how LLM-powered user profiles are constructed, interpreted, and applied in downstream tasks.
- Safety and misuse prevention, mitigating risks associated with hallucinated user traits, overspecification, harmful personalization, and other unintended model behaviors.

3 Workshop Sketch

3.1 Workshop Format and Planned Interaction

LLM-UP will be an interactive hybrid half-day workshop, which concludes with, two keynotes delivered by leading researchers and industrial practitioners, a research paper presentation with three to five high-quality research papers selected via rigorous review, an interactive panel discussion, fostering an exchange of ideas and visions for the future LLM-powered user profiling in search and recommendation. The specific time schedule is as follows:

- 08:30-09:00 Welcome and Opening Remarks
- 09:00-09:45 Invited Keynote 1
- 09:45-10:30 Invited Keynote 2
- 10:30-10:45 Coffee Break
- 10:45-11:30 Panel Discussions
- 11:30-12:30 Paper Encore

3.2 Distinction from Main Conference Topics

While SIGIR 2026 will feature work on IR and recommendation foundations, and possibly the application of LLMs within these areas, our workshop focuses specifically on how LLM-powered user profiling enables the generative inference, representation, and evolution of user preferences and intents, which offers a new pathway toward personalized search and recommendation. Moreover, the workshop highlights the unique challenges introduced by LLM-driven profiling, such as profile fidelity, temporal dynamics, privacy, and responsible personalization, that are not commonly addressed in the main conference tracks. By foregrounding these emerging opportunities and risks, the workshop complements SIGIR's core topics and draws attention to a critical yet underexplored dimension of next-generation personalized IR systems.

3.3 Target Audiences and Selection Process

The workshop targets a broad and interdisciplinary audience interested in advancing LLM-powered personalization for search and recommendation, including IR/RecSys/NLP researchers working on user modeling and generative inference, industry practitioners and engineers building personalized search/recommendation pipelines, and graduate students or newcomers seeking exposure to emerging methods and open problems.

Rather than call for extra submissions, we will invite some accepted papers' authors from the SIGIR 2026 main conference, whose work follows our workshop's topic, to share their thoughts, interesting experiences, and best practices during their experiments. This sharing is much more different from the main conference's presentation as it contains more about practical experience rather than theoretical introduction.

4 Organizers Information

Prof. Hongzhi Yin (attend onsite) is a Full Professor and ARC Future Fellow at the University of Queensland. He received his Ph.D. degree in computer science from Peking University. His research focuses on recommender systems, data mining, and large-scale personalization, with recent emphasis on LLM-based recommendation, federated learning, and privacy-preserving user modeling. He has published over 300 papers in top-tier venues such as SIGIR, KDD, WWW, ICDE, and TOIS, and so on. He has won eight Best Paper Awards at conferences like CIKM, WSDM, and ICDE. He has served as a main organizer for workshops and tutorials at WWW2025/2024/2022, CIKM2021, ICDE2019, and KDD2017. He is area chair for SIGIR2026, ACL2025, ICDM2025, and KDD2026/2025.

Dr. Wei Yuan (attend onsite) received his Ph.D. from the University of Queensland, where he is currently a Postdoctoral Research Fellow. His research focuses on language models, recommender systems, conversational information retrieval, as well as many other data-driven applications. He has published extensively in top-tier venues such as SIGIR, KDD, WWW, ICDE, AAAI, ACL, ICDM, WSDM, TOIS, and TKDE. He has also served as a reviewer for leading conferences such as SIGIR, KDD, WWW, WSDM, and CIKM.

Yi Zhang (attend onsite) is a Ph.D. student at Anhui University, specializing in recommendation systems, and LLM-enhanced personalization. His work explores user modeling, generative inference, and representation learning for recommendation. He has published in major IR and ML venues such as SIGIR, KDD, AAAI, WWW, TOIS, and TKDE.

Dr. Joel Mackenzie (attend onsite) is a senior lecturer and DECRA Fellow at the University of Queensland, where he conducts research in the field of information retrieval. He focuses on efficient and effective representations for large-scale search engines, including indexing, compression, and retrieval, as well as system-oriented effectiveness measurements and user behaviour analysis. He has been (co)chair for many conferences, including the resources track at SIGIR2026, the short paper track at SIGIR2025, the demo track at SIGIR 2024, and proceedings chair at CHIIR 2025/2021, SIGIR-AP2023 and WSDM2019. He has organized the workshop on efficient neural IR at SIGIR2023 and 2024,¹ and a workshop on Agentic AI at AAAI2026.²

Prof. Quoc Viet Hung Nguyen (attend onsite) received the PhD degree from EPFL, Switzerland. He is currently an associate professor with Griffith University, Australia. His research focuses on Data Augmentation and Integration, Data Quality, Recommender Systems, and Big Data Visualization, with special emphasis on web data and social data. He has published 200+ articles in top-tier venues, such as SIGIR, SIGMOD, VLDB, KDD, AAAI, ICDE, IJCAI, JVLDB, TKDE, TOIS, and TIST.

Prof. Wayne Xin Zhao (attend onsite/online) is currently a professor at Renmin University of China. He obtained the doctoral degree from Peking University in July 2014. He has broad research interests in the fields of information retrieval and natural language processing, with 400+ published papers at top-tier conferences/journals and 40000+ academic citations from Google Scholar. He has led the development of several popular open-source tools

¹<https://reueir.org/>

²<https://ai-4-finance.pages.dev/>

such as RecBole (recommendation system library) and TextBox (text generation library). 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.

Prof. Yong Li (attend onsite/online) is a Professor at Tsinghua University. His research interests are in the area of artificial intelligence, data science, and mobile computing. Dr. Li has served as General Chair, TPC Chair, TPC Member for several international workshops and conferences, and he is on the editorial board of four international journals. His papers have total citations more than 36000+ citations.

Prof. Lina Yao (attend onsite) is Professor and ARC Future Fellow at UNSW. Her research focuses on machine learning, data mining, and natural language processing, with applications in LLM-based agents, recommender systems, brain-computer interfaces, and the Internet of Things. She has served as an Associate Editor for several international journals and as a member of organizing committees for international conferences.

5 Related Workshops

Recent workshops indicate strong community interest in LLMs for IR/RecSys. GENNEXT@SIGIR2025 [2] and AgentIR@SIGIR2025 [6] emphasize language agents, generative models, and agent-based IR/RecSys paradigms, while GenAIR@KDD2025 [16] and IR-LLM@WWW2024 [8] broadly explore LLMs and generative AI for IR and RecSys. LLM4Eval@SIGIR2024 [13] solely investigates IR evaluation. Our workshop differs by explicitly framing LLM-powered user profiling as the primary research object across both search and recommendation, with emphasis on the full lifecycle of profiles and profile-specific evaluation and governance. This dedicated focus motivates a new workshop to consolidate emerging work, clarify best practices, and identify open research directions for next-generation personalized search and recommender systems.

Acknowledgments

The Australian Research Council supports this work under the streams of Future Fellowship (Grant No. FT210100624), the Discovery Project (Grant No. DP240101108), the Linkage Project (Grant No. LP230200892), and the Linkage Project (Grant No. LP240200546).

References

- [1] Tong Chen, Hongzhi Yin, Guanhua Ye, Zi Huang, Yang Wang, and Meng Wang. 2020. Try this instead: Personalized and interpretable substitute recommendation. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*. 891–900.
- [2] Yashar Deldjoo, Scott Sanner, Enrico Palumbo, Hugues Bouchard, Shuai Zhang, Pablo Castells, and Julian McAuley. 2025. GENNEXT: The Next Generation of IR and Recommender Systems with Language Agents, Generative Models, and Conversational AI. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 4195–4198.
- [3] Ge Gao, Alexey Taymanov, Eduardo Salinas, Paul Mineiro, and Dipendra Misra. 2024. Aligning llm agents by learning latent preference from user edits. *Advances in Neural Information Processing Systems* 37 (2024), 136873–136896.
- [4] Kostas Hatalas, Despina Christou, Joshua Myers, Steven Jones, Keith Lambert, Adam Amos-Binks, Zohreh Dannenhauer, and Dustin Dannenhauer. 2023. Memory matters: The need to improve long-term memory in llm-agents. In *Proceedings of the AAAI Symposium Series*, Vol. 2. 277–280.
- [5] Feng He, Tianqing Zhu, Dayong Ye, Bo Liu, Wanlei Zhou, and Philip Yu. 2024. The emerged security and privacy of llm agent: A survey with case studies. *Comput. Surveys* (2024).
- [6] Pengyue Jia, Qingpeng Cai, Xiangyu Zhao, Ling Pan, Xin Xin, Jin Huang, Weinan Zhang, Li Zhao, Dawei Yin, and Grace Hui Yang. 2025. Agentir: 2nd workshop on agent-based information retrieval. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 4180–4183.
- [7] Chumeng Jiang, Jiayin Wang, Weizhi Ma, Charles LA Clarke, Shuai Wang, Chuhan Wu, and Min Zhang. 2025. Beyond Utility: Evaluating LLM as Recommender. In *Proceedings of the ACM on Web Conference 2025*. 3850–3862.
- [8] Zheng Liu, Yujia Zhou, Yutao Zhu, Jianxun Lian, Chaozhao Li, Zhicheng Dou, Defu Lian, and Jian-Yun Nie. 2024. Information retrieval meets large language models. In *Companion Proceedings of the ACM Web Conference 2024*. 1586–1589.
- [9] Matéo Mahaut, Laura Aina, Paula Czarnowska, Momchil Hardalov, Thomas Mueller, and Lluís Màrquez. 2024. Factual Confidence of LLMs: on Reliability and Robustness of Current Estimators. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 4554–4570.
- [10] Fengran Mo, Kelong Mao, Ziliang Zhao, Hongjin Qian, Haonan Chen, Yiruo Cheng, Xiaoxi Li, Yutao Zhu, Zhicheng Dou, and Jian-Yun Nie. 2025. A survey of conversational search. *ACM Transactions on Information Systems* 43, 6 (2025), 1–50.
- [11] Fengran Mo, Chuan Meng, Mohammad Aliannejadi, and Jian-Yun Nie. 2025. Conversational search: From fundamentals to frontiers in the LLM era. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 4094–4097.
- [12] Shutong Qiao, Chen Gao, Wei Yuan, Yong Li, and Hongzhi Yin. 2024. LLM-based Bi-level Multi-interest Learning Framework for Sequential Recommendation. *arXiv preprint arXiv:2411.09410* (2024).
- [13] Hossein A Rahmani, Clemencia Siro, Mohammad Aliannejadi, Nick Craswell, Charles LA Clarke, Guglielmo Faggioli, Bhaskar Mitra, Paul Thomas, and Emine Yilmaz. 2024. Llm4eval: Large language model for evaluation in ir. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 3040–3043.
- [14] Lei Sang, Weichen Fei, Yi Zhang, Yuce Huang, and Yiwen Zhang. 2025. Heterogeneous adaptive preference learning for recommendation. *ACM Transactions on Recommender Systems* 4, 1 (2025), 1–25.
- [15] Shreya Shankar, JD Zamfirescu-Pereira, Björn Hartmann, Aditya Parameswaran, and Ian Arawjo. 2024. Who validates the validators? aligning llm-assisted evaluation of llm outputs with human preferences. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–14.
- [16] Narges Tabari, Aniket Deshmukh, Wang-Cheng Kang, Julian McAuley, James Caverlee, Neil Shah, and George Karypis. 2025. Second Workshop on Generative AI for Recommender Systems and Personalization. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*. 6294–6295.
- [17] Zongwei Wang, Min Gao, Junliang Yu, Xinyi Gao, Quoc Viet Hung Nguyen, Shazia Sadiq, and Hongzhi Yin. 2025. Id-free not risk-free: Llm-powered agents unveil risks in id-free recommender systems. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1902–1911.
- [18] Jiaqi Zhang, Junliang Yu, Zongwei Wang, Wei Yuan, Tong Chen, Quoc Viet Hung Nguyen, Bin Cui, and Hongzhi Yin. 2025. Towards Reasoning-Aware Recommender Systems: A Survey in the LLM Era. *Authorea Preprints* (2025).
- [19] Yi Zhang, Yiwen Zhang, Yu Wang, Tong Chen, and Hongzhi Yin. 2025. ProEx: A Unified Framework Leveraging Large Language Model with Profile Extrapolation for Recommendation. *arXiv preprint arXiv:2512.00679* (2025).
- [20] Zeyu Zhang, Quanyu Dai, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. 2025. A survey on the memory mechanism of large language model-based agents. *ACM Transactions on Information Systems* 43, 6 (2025), 1–47.
- [21] Jianling Zhong, Weiwei Guo, Huiji Gao, and Bo Long. 2020. Personalized query suggestions. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1645–1648.
- [22] Hanzhang Zhou, Zijian Feng, Zixiao Zhu, Junliang Qian, and Kezhi Mao. 2024. Unibias: Unveiling and mitigating llm bias through internal attention and fn manipulation. *Advances in Neural Information Processing Systems* 37 (2024).