From Generic Intelligence to Personalized AI: A Tutorial on Foundations of LLM Personalization

Ruijie Wang
Beihang University
School of Computer Science and Engineering
Beijing, China
ruijiew@buaa.edu.cn

Qingkai Zeng
Nankai University
College of Computer Science
Tianjin, China
qzengnkcs@gmail.com

Xuefei Wang
Beihang University
School of Computer Science and Engineering
Beijing, China
xuefeiw@buaa.edu.cn

Yutong Ye
Beihang University
School of Computer Science and Engineering
Beijing, China
yutongye@buaa.edu.cn

Li Sun
Beijing University of Posts and Telecommunications
School of Computer Science
Beijing, China
lsun@bupt.edu.cn

Jianxin Li
Beihang University
School of Computer Science and Engineering
Beijing, China
lijx@buaa.edu.cn

Phillip S. Yu
University of Illinois Chicago
Department of Computer Science
Chicago, Illinois, USA
psyu@uic.edu

Abstract

Large language models (LLMs) have achieved remarkable success across diverse applications, yet their generic training paradigm limits effectiveness in user-specific scenarios. LLM personalization aims to adapt large models to individual users or user groups by incorporating preferences, histories, and contextual signals while maintaining generalization ability. Despite growing interest, a unified understanding of personalization strategies for LLMs remains lacking, and key challenges—including data scarcity, efficiency, privacy, and fairness—are still open. This tutorial provides a systematic overview of LLM personalization, covering core paradigms such as user-aware representations, prompt-based adaptation, memory augmentation, parameter-efficient fine-tuning, and continual learning. We discuss evaluation protocols, theoretical foundations, and real-world applications in recommendation, dialogue systems, education, and healthcare. By synthesizing recent advances and identifying emerging research directions, this tutorial aims to equip the KDD community with a structured framework for advancing personalized large model research and deployment.

CCS Concepts

• **Computing methodologies → Artificial intelligence**.

Keywords

LLM Personalization, User Modeling, Personalized Evaluation

KDD '26, Jeju Island, Republic of Korea

ACM ISBN 979-8-4007-2259-2/2026/08
https://doi.org/10.1145/3770855.3816457

ACM Reference Format:
Ruijie Wang, Qingkai Zeng, Xuefei Wang, Yutong Ye, Li Sun, Jianxin Li, and Phillip S. Yu. 2026. From Generic Intelligence to Personalized AI: A Tutorial on Foundations of LLM Personalization. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '26), August 09–13, 2026, Jeju Island, Republic of Korea.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3770855.3816457

1 introduction

Large Language Models (LLMs) are increasingly deployed in user-centric applications that demand responses attuned to individual users [26, 31, 40, 45]. This demand emerges in personal assistants, educational platforms, and other interactive systems shaped by heterogeneous user contexts [44, 47, 49, 50, 58]. Accordingly, LLM personalization has emerged as a foundational paradigm for user-aware intelligence [38] By integrating user-specific signals such as profiles, interaction histories, and preference patterns into generation, personalized LLMs move beyond generic language competence toward outputs that are context-sensitive, preference-consistent, and more valuable in real-world applications [8].

However, deploying personalized LLMs in real-world settings remains challenging, as it requires effective adaptation, reliable user data, and principled evaluation [38, 45]. First, personalization under limited user supervision is difficult: although prior studies have explored retrieval augmentation, prompting, representation learning, and reinforcement learning from human feedback [17, 22, 28, 36], real-world user signals are often sparse, noisy, or privacy-sensitive, and cold-start scenarios remain underexplored [13, 34, 60]. Second, evaluating whether an LLM is genuinely personalized is nontrivial, as existing direct generation-based and indirect downstream-task evaluations [4, 20, 29, 38] are limited by scarce ground truth and the difficulty of jointly assessing style, relevance, and accuracy [60].

These issues make effectively leveraging imperfect user data and ensuring reliable evaluation central challenges for the field [45].

To address these challenges, substantial research has sought to improve how personalized LLMs acquire, utilize, and validate user-specific information [7, 9, 37]. Methodological studies investigate how user information can be incorporated into LLMs via retrieval-augmented generation, prompting, representation learning, or reinforcement learning from human feedback [10, 15, 36, 56]. Data-centric studies examine what forms of user information support personalization and under which settings it can be realized. Existing work distinguishes personalization by preference-modeling scope, including user-level, persona-level, and global-preference settings [25, 33, 63], and categorizes datasets by whether they provide ground-truth user-written text for direct personalized generation or only support downstream personalization tasks [20, 30, 58]. Evaluation studies assess whether personalized behavior is genuinely achieved, distinguishing intrinsic evaluation of personalized text from extrinsic evaluation in downstream applications, often with human judgments of user-specific satisfaction [5, 11, 23, 39, 46].

This tutorial systematically reviews LLM personalization for real-world user-centric applications, focusing on how user-specific information is acquired, incorporated, and evaluated. It is organized into four parts. Part I introduces the motivations, problem settings, and foundational concepts. Part II examines key challenges of incorporating user-specific information, and then introduce representative techniques, including retrieval-augmented generation, prompting, representation learning, and reinforcement learning from human feedback. Part III covers data foundations, personalization granularities, and dataset settings for direct and downstream personalization. Part IV discusses evaluation protocols, benchmarks, metrics, and open challenges toward effective, reliable, and responsible personalized LLMs.

2 Presentor Biography

In-person presenters of the tutorial include Ruijie Wang, Qingkai Zeng, Xuefei Wang, Li Sun, Jianxin Li, and Philip S. Yu, who will deliver the lecture-style presentations and interact with the audience during the session.

Ruijie Wang is a Professor at the School of Computer Science and Engineering, Beihang University. His research focuses on trustworthy and efficient AI, with an emphasis on large language models, graph learning, temporal data mining, and structured knowledge integration. He has published more than 40 papers in top-tier venues including NeurIPS, ACL, KDD, WWW, SIGIR, and UbiComp. His research has received multiple recognitions including the IEEE DCOSS Best Paper Award and UIUC PhD Fellowship. His work on graph-enhanced LLMs and structured knowledge integration directly aligns with the core technical focus of this tutorial.

Qingkai Zeng is an Associate Professor in the College of Computer Science at Nankai University. His research focuses on text mining, knowledge-enhanced LLMs, and knowledge graph reasoning, with an emphasis on improving reliability and interpretability in knowledge-intensive tasks. He has published more than 30 papers in leading venues such as KDD, WWW, ACL, EMNLP, and ICLR, and served as PC member or Area Chair for major conferences. His expertise in knowledge-enhanced LLMs and graph-based reasoning supports the methodological foundations of this tutorial.

Xuefei Wang is a Ph.D. student at Beihang University. His research focuses on LLM personalization and LLM system implementation, contributing implementation perspectives to this tutorial.

Yutong Ye is a Postdoctoral Researcher at Beihang University whose research focuses on graph data management, scalable graph analytics, and graph-enhanced LLMs. He has published over 30 papers in venues such as SIGMOD, VLDB, ICDE, AAAI, IJCAI, and CVPR, and his expertise supports scalable graph–LLM integration.

Li Sun is an Associate Professor at Beijing University of Posts and Telecommunications. His research focuses on large language models, graph representation learning, and knowledge mining. He has published over 70 papers in premier venues, including ICML, NeurIPS, ICLR, KDD, WWW, AAAI, TPAMI, TKDE, ACM TOIS, and TWEB, and received the ACM CIKM 2022 Best Paper Award. He has served as Area Chair or Program Committee member for major conferences such as KDD, ICML, NeurIPS, ICLR, WWW, IJCAI, and AAAI. His research on graph learning and LLM-based knowledge mining provides key algorithmic perspectives for the tutorial.

Jianxin Li is a Professor at the School of Computer Science and Engineering and the Beijing Advanced Innovation Center for Big Data and Brain Computing at Beihang University. His research interests include social network analysis, trustworthy machine learning, and large-scale data mining. He has published extensively in leading journals and conferences including TKDE, TDSC, TOIS, TKDD, KDD, AAAI, and WWW. He received the AAAI 2021 Best Paper Award and the DependSys 2017 Best Paper Award. He has also achieved first place in several international data science competitions. He currently serves as Associate Editor of IEEE Transactions on Computers and several other journals.

Philip S. Yu is a Distinguished Professor of Computer Science and the Wexler Chair in Information Technology at the University of Illinois Chicago (UIC). Before joining UIC, he managed the Software Tools and Techniques Department at IBM. He is a Fellow of ACM, IEEE, and AAAS. His research spans data mining, graph learning, user behavior modeling, and multimodal large language models. He has published over 1,700 papers with more than 252,000 citations and an h-index of 214. He received the ACM SIGKDD Innovation Award, IEEE Computer Society Technical Achievement Award, and multiple Test-of-Time awards. His pioneering work in graph mining and large-scale data analytics provides foundational insights for the tutorial's discussion of graph-enhanced AI and LLM systems.

3 Tutorial Details

3.1 Outline of the Tutorial

The tutorial will last for 3 hours, and the outline is as follows:

- Part I: Introduction
  - Background and Motivations
  - Problem Definitions and Settings
- Part II: Personalization Techniques for LLMs
  - Personalization via Retrieval-Augmented Generation
  - Personalization via Prompting
  - Personalization via Representation Learning
  - Personalization via Reinforcement Learning from Human Feedback
- Part III: Data Foundations
  - Personalization Granularity

  - Dataset Settings for Direct and Downstream Personalization
- Part IV: Evaluation and Open Questions
  - Evaluation Protocols and Benchmarks
  - Open Challenges and Future Directions

The details for Parts II-IV are introduced as follows.

3.2 Part II: Personalization Techniques for LLMs

Personalization techniques for LLMs aim to incorporate user-specific information in a principled and effective manner [6, 55]. Existing studies can be broadly categorized into retrieval-augmented generation, prompting, representation learning, and reinforcement learning from human feedback [35, 61, 62], which differ in how user signals are represented, injected, and optimized [32, 53, 54]. Retrieval-augmented generation treats user information as external knowledge, typically encoded as vectors and retrieved at inference time through similarity-based search [36]. Prompting incorporates user information into the input context to guide personalization tasks [18, 21, 43]. Representation learning maps user information into neural representations for fine-tuning, parameter-efficient adaptation, or user-specific embeddings [27]. Reinforcement learning from human feedback leverages user information as alignment signals [3, 33, 35, 54]. This part reviews how these paradigms enable LLMs to exploit user information in real-world applications.

3.3 Part III: Data Foundations

Data foundations of personalized LLMs determine what forms of user information can support personalization and how finely such personalization can be realized [48, 51]. The heterogeneous nature of user data, together with sparsity, dynamic preferences, and privacy concerns, makes personalization substantially more challenging than generic language modeling [41, 42]. In this part, we discuss the challenges of building personalized LLMs under realistic data settings, including limited user-specific data, skewed user coverage, and cold-start scenarios [19]. We categorize related studies based on the granularity of personalization objectives and introduce three major settings: user-level personalization, persona-level personalization, and global preference personalization [14, 17, 52]. Moreover, we discuss the implications of data availability and structure for building personalized systems that are effective, robust, and adaptable across different application scenarios [1, 2, 57].

3.4 Part IV: Evaluation and Open Questions

To assess whether personalized LLMs truly satisfy user-specific requirements, comprehensive evaluation has become a central problem in this area [16, 39]. In this setting, intrinsic evaluation compares personalized outputs with ground-truth user-written text or preference signals, while extrinsic evaluation measures their utility in downstream applications [24, 59]. Reliable evaluation further requires appropriate protocols, benchmarks, and metrics, since existing benchmarks often fail to capture realistic multi-turn interactions, evolving preferences, and diverse contexts [12, 20]. This part discusses evaluation foundations and broader open questions for the field, including benchmark design, cold-start settings, bias and fairness, privacy, and multimodal personalization.

4 Related Tutorials

- **Tutorial on Large Language Models for Recommendation**
  - **Presenters**: Wenyue Hua, Lei Li, Shuyuan Xu, Li Chen, and Yongfeng Zhang
  - **Conference**: RecSys '23, September 18–22, 2023, Singapore
  - **Connection**: This tutorial introduces large language models for recommendation, covering datasets, models, evaluation, toolkits, and real-world systems.
  - **Differences**: This tutorial focuses on recommendation as a downstream application of LLMs. Our tutorial instead studies LLM personalization more broadly, covering techniques, data foundations, evaluation, and open questions beyond recommendation.
- **Large Language Models for Recommendation: Progresses and Future Directions**
  - **Presenters**: Jizhi Zhang, Keqin Bao, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He
  - **Conference**: WWW '24, May 13–17, 2024, Singapore
  - **Connection**: This tutorial presents recent advances in using large language models for recommendation, which is closely related to personalized generation and user-aware modeling.
  - **Differences**: This tutorial centers on LLMs for recommendation. Our tutorial instead focuses on the broader problem of LLM personalization, covering core techniques, personalization data, and evaluatio.

5 Tutorial Participation

Target Audience and Expected Reach. The intended audience of this tutorial includes students, researchers, and practitioners interested in LLMs, personalization, and user-centric AI systems. Basic familiarity with machine learning and natural language processing is helpful but not required. Given the growing use of LLMs in assistants, recommendation, customer interaction, and education, the tutorial is relevant to both academic and industrial audiences. We will promote the tutorial through academic networks, personal webpages, social media, and related research communities. The estimated number of participants is around 50–100.

Engagement Plan and Expected Impact. As LLM personalization has become a rapidly developing research topic, this tutorial is expected to attract broad interest. We will prepare clear tutorial materials and recorded videos with intuitive examples to support accessibility and participation. To improve the interactive quality of the tutorial, we will encourage questions both during and after the session. This tutorial is expected to serve as a useful reference and source of research directions for researchers and practitioners.

6 Social Impact Statement

This tutorial is closely related to the KDD community, as personalization is a longstanding topic in recommendation, search, user modeling, and interactive systems. With the rapid deployment of LLMs, personalization is becoming an important direction for data-driven AI research and applications. From a societal perspective, personalized LLMs may support more equitable and inclusive AI services by adapting to users with diverse needs, preferences, and contexts. These capabilities are relevant to education, recommendation, and digital assistants, while also raising important concerns about fairness, bias, and privacy that this tutorial will discuss.

References

[1] Mahyar Abbasian, Iman Azimi, Amir M. Rahmani, and Ramesh Jain. Conversational health agents: A personalized llm-powered agent framework, 2023. arXiv preprint.

[2] Mahyar Abbasian, Zhongqi Yang, Elahe Khatibi, Pengfei Zhang, Nitish Nagesh, Iman Azimi, Ramesh Jain, and Amir M. Rahmani. Knowledge-infused llm-powered conversational health agent: A case study for diabetes patients, 2024. arXiv preprint.

[3] Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce style optimization for learning from human feedback in llms, 2024. arXiv preprint.

[4] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In *Proceedings of the 17th ACM conference on recommender systems*, pages 1007–1014, 2023.

[5] Anna Bavaresco, Raffaella Bernardi, Leonardo Bertolazzi, Desmond Elliott, Raquel Fernández, Albert Gatt, Esam Ghaleb, Mario Giulianelli, Michael Hanna, and Alexander Koller. Llms instead of human judges? a large scale empirical study across 20 nlp evaluation tasks, 2024. arXiv preprint.

[6] Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, et al. From persona to personalization: A survey on role-playing language agents, 2024. arXiv preprint.

[7] Jin Chen, Zheng Liu, Xu Huang, Chenwang Wu, Qi Liu, Gangwei Jiang, Yuanhao Pu, Yuxuan Lei, Xiaolong Chen, and Xingmei Wang. When large language models meet personalization: Perspectives of challenges and opportunities. *World Wide Web*, 27(4):42, 2024.

[8] Jin Chen, Zheng Liu, Xu Huang, Chenwang Wu, Qi Liu, Gangwei Jiang, Yuanhao Pu, Yuxuan Lei, Xiaolong Chen, Xingmei Wang, et al. When large language models meet personalization: Perspectives of challenges and opportunities. *World wide web*, 27(4):42, 2024.

[9] Junyi Chen. A survey on large language models for personalized and explainable recommendations, 2023. arXiv preprint.

[10] Ruizhe Chen, Xiaotian Zhang, Meng Luo, Wenhao Chai, and Zuozhu Liu. Pad: Personalized alignment of llms at decoding-time, 2024. arXiv preprint.

[11] Sourish Dasgupta, Ankush Chander, Parth Borad, Isha Motiyani, and Tanmoy Chakraborty. Perseval: Assessing personalization in text summarizers, 2024. arXiv preprint.

[12] Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. Mind2web: Towards a generalist agent for the web. In *Advances in Neural Information Processing Systems*, volume 36, pages 28599–28623. Curran Associates, Inc., 2023.

[13] Horatiu Dumitru, Marek Gibiec, Negar Hariri, Jane Cleland-Huang, Bamshad Mobasher, Carlos Castro-Herrera, and Mehdi Mirakhorli. On-demand feature recommendations derived from mining public product descriptions. In *Proceedings of the 33rd international conference on software engineering*, pages 181–190, 2011.

[14] Iason Gabriel. Artificial intelligence, values, and alignment. *Minds and Machines*, 30(3):411–437, 2020.

[15] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. Retrieval-augmented generation for large language models: A survey, 2023. arXiv preprint.

[16] Zhengyu Hu, Linxin Song, Jieyu Zhang, Zheyuan Xiao, Jingang Wang, Zhenyu Chen, Jieyu Zhao, and Hui Xiong. Rethinking llm-based preference evaluation, 2024. arXiv preprint.

[17] Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. Personalized soups: Personalized large language model alignment via post-hoc parameter merging. *arXiv preprint arXiv:2310.11564*, 2023.

[18] Jaehyung Kim and Yiming Yang. Few-shot personalization of llms with misaligned responses, 2024. arXiv preprint.

[19] Hannah Rose Kirk, Alexander Whitefield, Paul Rottger, Andrew M Bean, Katerina Margatina, Rafael Mosquera-Gomez, Juan Ciro, Max Bartolo, Adina Williams, He He, et al. The prism alignment dataset: What participatory, representative and individualised human feedback reveals about the subjective and multicultural alignment of large language models. *Advances in Neural Information Processing Systems*, 37:105236–105344, 2024.

[20] Ishita Kumar, Snigdha Viswanathan, Sushrita Yerra, Alireza Salemi, Ryan A Rossi, Franck Dernoncourt, Hanieh Deilamsalehy, Xiang Chen, Ruiyi Zhang, Shubham Agarwal, et al. Longlamp: A benchmark for personalized long-form text generation. *arXiv preprint arXiv:2407.11016*, 2024.

[21] Cheng Li, Mingyang Zhang, Qiaozhu Mei, Weize Kong, and Michael Bendersky. Learning to rewrite prompts for personalized text generation. In *Proceedings of the ACM Web Conference 2024*, WWW '24, pages 3345–3356, Singapore, may 2024. Association for Computing Machinery.

[22] Cheng Li, Mingyang Zhang, Qiaozhu Mei, Yaqing Wang, Spurthi Amba Hombaiah, Yi Liang, and Michael Bendersky. Teach llms to personalize–an approach inspired by writing education. *arXiv preprint arXiv:2308.07968*, 2023.

[23] Minzhi Li, Taiwei Shi, Caleb Ziems, Min-Yen Kan, Nancy Chen, Zhengyuan Liu, and Diyi Yang. Coannotating: Uncertainty-guided work allocation between human and large language models for data annotation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1487–1505, Singapore, December 2023. Association for Computational Linguistics.

[24] Pan Li and Alexander Tuzhilin. Towards controllable and personalized review generation, 2019. arXiv preprint.

[25] Xinyu Li, Ruiyang Zhou, Zachary C Lipton, and Liu Leqi. Personalized language modeling from personalized human feedback. *arXiv preprint arXiv:2402.05133*, 2024.

[26] Yuanchun Li, Hao Wen, Weijun Wang, Xiangyu Li, Yizhen Yuan, Guohong Liu, Jiacheng Liu, Wenxing Xu, Xiang Wang, Yi Sun, et al. Personal llm agents: Insights and survey about the capability, efficiency and security. *arXiv preprint arXiv:2401.05459*, 2024.

[27] Jiongnan Liu, Yutao Zhu, Shuting Wang, Xiaochi Wei, Erxue Min, Yu Lu, Shuaiqiang Wang, Dawei Yin, and Zhicheng Dou. Llms + persona-plug = personalized llms, 2024. arXiv preprint.

[28] Qijiong Liu, Nuo Chen, Tetsuya Sakai, and Xiao-Ming Wu. Once: Boosting content-based recommendation with both open-and closed-source large language models. In *Proceedings of the 17th ACM international conference on web search and data mining*, pages 452–461, 2024.

[29] Hanjia Lyu, Song Jiang, Hanqing Zeng, Yinglong Xia, Qifan Wang, Si Zhang, Ren Chen, Chris Leung, Jiajie Tang, and Jiebo Luo. Llm-rec: Personalized recommendation via prompting large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 583–612, 2024.

[30] Pierre-Emmanuel Mazaré, Samuel Humeau, Martin Raison, and Antoine Bordes. Training millions of personalized dialogue agents. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2775–2779, Brussels, Belgium, October-November 2018. Association for Computational Linguistics.

[31] Microsoft. Meet microsoft copilot. https://www.microsoft.com/en-us/microsoft-copilot/meet-copilot, 2024. Accessed: 2024-09-17.

[32] Lin Ning, Luyang Liu, Jiaxing Wu, Neo Wu, Devora Berlowitz, Sushant Prakash, Bradley Green, Shawn O'Banion, and Jun Xie. User-llm: Efficient llm contextualization with user embeddings, 2024. arXiv preprint.

[33] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.

[34] Luis G Perez, Manuel Barranco, and Luis Martinez. Building user profiles for recommender systems from incomplete preference relations. In *2007 IEEE International Fuzzy Systems Conference*, pages 1–6. IEEE, 2007.

[35] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems*, volume 36, pages 53728–53741. Curran Associates, Inc., 2023.

[36] Chris Richardson, Yao Zhang, Kellen Gillespie, Sudipta Kar, Arshdeep Singh, Zeynab Raeesy, Omar Zia Khan, and Abhinav Sethy. Integrating summarization and retrieval for enhanced personalization via large language models. *arXiv preprint arXiv:2310.20081*, 2023.

[37] Alireza Salemi, Surya Kallumadi, and Hamed Zamani. Optimization methods for personalizing large language models through retrieval augmentation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 752–762, 2024.

[38] Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. Lamp: When large language models meet personalization. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7370–7392, 2024.

[39] Shreya Shankar, JD Zamfirescu-Pereira, Björn Hartmann, Aditya G. Parameswaran, and Ian Arawjo. Who validates the validators? aligning llm-assisted evaluation of llm outputs with human preferences, 2024. arXiv preprint.

[40] Significant Gravitas. Autogpt. https://github.com/Significant-Gravitas/AutoGPT, 2024. Accessed: 2024-09-18.

[41] Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Mireshghallah, Christopher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, Tim Althoff, and Yejin Choi. A roadmap to pluralistic alignment, 2024. arXiv preprint.

[42] Robin Staab, Mark Vero, Mislav Balunović, and Martin Vechev. Beyond memorization: Violating privacy via inference with large language models, 2023. arXiv preprint.

[43] Chenkai Sun, Ke Yang, Revanth Gangi Reddy, Yi R. Fung, Hou Pong Chan, ChengXiang Zhai, and Heng Ji. Persona-db: Efficient large language model personalization for response prediction with collaborative data refinement, 2024. arXiv preprint.

[44] Ruixiang Tang, Xiaotian Han, Xiaoqian Jiang, and Xia Hu. Does synthetic data generation of llms help clinical text mining? *arXiv preprint arXiv:2303.04360*, 2023.

[45] Yu-Min Tseng, Yu-Chao Huang, Teng-Yun Hsiao, Wei-Lin Chen, Chao-Wei Huang, Yu Meng, and Yun-Nung Chen. Two tales of persona in llms: A survey of role-playing and personalization. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16612–16631, 2024.

[46] Rahul Vansh, Darsh Rank, Sourish Dasgupta, and Tanmoy Chakraborty. Accuracy is not enough: Evaluating personalization in summarizers. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2582–2595, Singapore, December 2023. Association for Computational Linguistics.

[47] Shen Wang, Tianlong Xu, Hang Li, Chaoli Zhang, Joleen Liang, Jiliang Tang, Philip S Yu, and Qingsong Wen. Large language models for education: A survey and outlook. *IEEE Signal Processing Magazine*, 42(6):51–63, 2026.

[48] Zekun Moore Wang, Zhongyuan Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, and Man Zhang. Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models, 2023. arXiv preprint.

[49] Zichao Wang, Jakob Valdez, Debshila Basu Mallick, and Richard G. Baraniuk. Towards human-like educational question generation with large language models. In *Artificial Intelligence in Education*, pages 153–166, Cham, 2022. Springer International Publishing.

[50] Stanisław Woźniak, Bartłomiej Koptyra, Arkadiusz Janz, Przemysław Kazienko, and Jan Kocoń. Personalized large language models. In *2024 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 511–520. IEEE, 2024.

[51] Junda Wu, Hanjia Lyu, Yu Xia, Zhehao Zhang, Joe Barrow, Ishita Kumar, Mehrnoosh Mirtaheri, Hongjie Chen, Ryan A. Rossi, and Franck Dernoncourt. Personalized multimodal large language models: A survey, 2024. arXiv preprint.

[52] Likang Wu, Zhaopeng Qiu, Zhi Zheng, Hengshu Zhu, and Enhong Chen. Exploring large language model for graph data understanding in online job recommendations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 9178–9186. AAAI Press, 2024.

[53] Yunjia Xi, Weiwen Liu, Jianghao Lin, Jieming Zhu, Bo Chen, Ruiming Tang, Weinan Zhang, Rui Zhang, and Yong Yu. Towards open-world recommendation with knowledge augmentation from large language models, 2023. arXiv preprint.

[54] Zhangchen Xu, Fengqing Jiang, Luyao Niu, Yuntian Deng, Radha Poovendran, Yejin Choi, and Bill Yuchen Lin. Magpie: Alignment data synthesis from scratch by prompting aligned llms with nothing, 2024. arXiv preprint.

[55] Diyi Yang and Lucie Flek. Towards user-centric text-to-text generation: A survey. In *Text, Speech, and Dialogue: 24th International Conference, TSD 2021, Olomouc, Czech Republic, September 6–9, 2021, Proceedings*, pages 3–22, Berlin, Heidelberg, 2021. Springer-Verlag.

[56] Bin Yin, Junjie Xie, Yu Qin, Zixiang Ding, Zhichao Feng, Xiang Li, and Wei Lin. Heterogeneous knowledge fusion: A novel approach for personalized recommendation via llm. In *Proceedings of the 17th ACM Conference on Recommender Systems*, pages 599–601, New York, NY, USA, 2023. Association for Computing Machinery.

[57] Xiaohan Yu, Li Zhang, Xin Zhao, Yue Wang, and Zhongrui Ma. Ra-rec: An efficient id representation alignment framework for llm-based recommendation, 2024. arXiv preprint.

[58] Jiayi Yuan, Ruixiang Tang, Xiaoqian Jiang, and Xia Hu. Large language models for healthcare data augmentation: An example on patient-trial matching. In *AMIA Annual Symposium Proceedings*, volume 2023, page 1324, 2024.

[59] Xinfeng Yuan, Siyu Yuan, Yuhan Cui, Tianhe Lin, Xintao Wang, Rui Xu, Jiangjie Chen, and Deqing Yang. Evaluating character understanding of large language models via character profiling from fictional works, 2024. arXiv preprint.

[60] Zhehao Zhang, Ryan A Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, et al. Personalization of large language models: A survey. *arXiv preprint arXiv:2411.00027*, 2024.

[61] Zhi Zheng, Zhaopeng Qiu, Xiao Hu, Likang Wu, Hengshu Zhu, and Hui Xiong. Generative job recommendations with large language model, 2023. arXiv preprint.

[62] Yuchen Zhuang, Haotian Sun, Yue Yu, Qifan Wang, Chao Zhang, and Bo Dai. Hydra: Model factorization framework for black-box llm personalization, 2024. arXiv preprint.

[63] Thomas P. Zollo, Andrew Wei Tung Siah, Naimeng Ye, Ang Li, and Hongseok Namkoong. Personalllm: Tailoring llms to individual preferences, 2024. arXiv preprint.