

# Large Language Models for Recommendation: Progresses and Future Directions

Keqin Bao\*  
baokq@mail.ustc.edu.cn  
University of Science and Technology  
of China

Wenjie Wang  
wenjiewang96@gmail.com  
National University of Singapore

Jizhi Zhang  
cdzhangjizhi@mail.ustc.edu.cn  
University of Science and Technology  
of China

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

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

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

## ABSTRACT

The powerful large language models (LLMs) have played a pivotal role in advancing recommender systems. Recently, in both academia and industry, there has been a surge of interest in developing LLMs for recommendation, referred to as LLM4Rec. This includes endeavors like leveraging LLMs for generative item retrieval and ranking, as well as the exciting possibility of building universal LLMs for diverse open-ended recommendation tasks. These developments hold the potential to reshape the traditional recommender paradigm, paving the way for the next-generation recommender systems.

In this tutorial, we aim to retrospect the evolution of LLM4Rec and conduct a comprehensive review of existing research. In particular, we will clarify how recommender systems benefit from LLMs through a variety of perspectives, including the model architecture, learning paradigm, and the strong abilities of LLMs such as chatting, generalization, planning, and generation. Furthermore, we will discuss the critical challenges and open problems in this emerging field, for instance, the trustworthiness, efficiency, and model retraining issues. Lastly, we will summarize the implications of previous work and outline future research directions. We believe that this tutorial will assist the audience in better understanding the progress and prospects of LLM4Rec, inspiring them for future exploration. This, in turn, will drive the prosperity of LLM4Rec, possibly fostering a paradigm shift in recommendation systems.

## CCS CONCEPTS

• Information systems → Recommender systems.

## KEYWORDS

Large Language Models, Recommender Systems, Generative Recommendation, Generative Models

\*Main contact author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*SIGIR-AP '23, November 26–28, 2023, Beijing, China*

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-0408-6/23/11...\$15.00  
<https://doi.org/10.1145/3624918.3629550>

## ACM Reference Format:

Keqin Bao\*, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Large Language Models for Recommendation: Progresses and Future Directions. In *Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region (SIGIR-AP '23), November 26–28, 2023, Beijing, China*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3624918.3629550>

## 1 FORMAT AND INTENDED AUDIENCE

• **Format.** This is a half-day (3 hours plus breaks) lecture-style tutorial, conducted on-site. All tutorial presenters plan to attend the conference in person.

• **Intended audience.** This tutorial is tailored to data scientists, machine learning engineers, and researchers interested in harnessing the potential of large language models (LLMs) for recommender systems. It presupposes that the audience possesses an introductory understanding of recommendation techniques. The tutorial will give an induction of the fundamental concepts and capabilities of LLMs, facilitating the audience's comprehension of the progress and future trajectories of LLM-enhanced recommender systems, and offering insights into their practical applications or research possibilities. Furthermore, this tutorial extends its relevance to researchers and practitioners in various information retrieval domains beyond recommendation, who aspire to leverage the power of LLMs for their specific tasks.

## 2 PRESENTERS

**Keqin Bao**<sup>1</sup> is a Ph.D. student at University of Science and Technology of China (USTC), supervised by Prof. *Fengli Feng and Xiangnan He*. His research interest lies in the recommender system and large language models. He has several publications in top conferences such as RecSys.

**Jizhi Zhang**<sup>2</sup> is a Ph.D. student at University of Science and Technology of China (USTC), supervised by Prof. Fuli Feng and Prof. Xiangnan He. His research interest lies in the recommender system and large language models. He has several publications in top conferences such as SIGIR, RecSys, WWW, and ACL.

**Yang Zhang**<sup>3</sup> is a Ph.D. student at the University of Science and Technology of China (USTC), supervised by Prof. Xiangnan He. His

<sup>1</sup>[http://data-science.ustc.edu.cn/\\_upload/tpl/12/b5/4789/template4789/author/keqin-bao.html](http://data-science.ustc.edu.cn/_upload/tpl/12/b5/4789/template4789/author/keqin-bao.html).

<sup>2</sup>[https://data-science.ustc.edu.cn/\\_upload/tpl/12/b5/4789/template4789/author/jizhi-zhang.html](https://data-science.ustc.edu.cn/_upload/tpl/12/b5/4789/template4789/author/jizhi-zhang.html).

<sup>3</sup><https://scholar.google.com/citations?user=M9NcazMAAAAJ>.

research interest lies in the recommender system. He has over ten publications that appeared in top conferences/journals. His work on recommendation has received the Best Paper Honorable Mention in SIGIR 2021. He has served as the PC member and reviewer for the top conferences and journals including TKDE, TOIS, RecSys, WSDM, WWW, SIGIR, and KDD. He has presented the tutorial on recommendation in WWW 2022 and SIGIR 2023.

**Wenjie Wang**<sup>4</sup> is a research fellow at the School of Computing, National University of Singapore. He received Ph.D. in Computer Science from National University of Singapore in 2023. His research interests cover causal recommendation, data mining, and multimedia. His publications appear in several top conferences and journals such as SIGIR, KDD, WWW, ACM MM, and TIP. Moreover, he has served as the PC member and reviewer for the top conferences and journals including TPAMI, TOIS, SIGIR, WWW, and KDD. He has presented tutorials in WWW 2022 and SIGIR 2023.

**Fuli Feng**<sup>5</sup> is a professor at the University of Science and Technology of China. His research interests include information retrieval, data mining, causal inference, and multi-media processing. He has over 60 publications appeared in several top conferences such as SIGIR, WWW, and SIGKDD, and journals including TKDE and TOIS. He has received the Best Paper Honorable Mention of SIGIR 2021 and Best Poster Award of WWW 2018. Moreover, he has served as the reviewer for several top conferences and journals, including SIGIR, WWW, SIGKDD, NeurIPS, ICML, ICLR, ACL, TOIS, TKDE, TNNLS, TPAMI. He has rich teaching experience and has organized tutorials at SIGIR'23, WWW'21&23, and RecSys'21.

**Xiangnan He**<sup>6</sup> is a professor at the University of Science and Technology of China (USTC). His research interests span information retrieval, data mining, and applied machine learning. He has over 100 publications appeared in top conferences such as SIGIR, WWW, and KDD, and journals including TKDE, TOIS, and TNNLS. His work on recommender system has received the Best Paper Award Honorable Mention in SIGIR (2023, 2021, 2016) and WWW (2018). He has rich teaching experience and has organized tutorials at SIGIR'18&20&23, WSDM'19&20&22, WWW'18&21&22, CIKM'19, and RecSys'21.

## 3 TOPIC AND RELEVANCE

### 3.1 Motivation

Recommender systems have become an integral part of our digital lives, providing personalized recommendations regarding products, services, and content [1]. Looking back at the history of recommender systems, language models (LMs) are essential in shaping the landscape of recommender algorithms.

In terms of model architecture, sequential recommendation has drawn inspiration from the Transformer architecture, giving rise to innovative approaches like BERT4Rec [45] and SASRec [31]. Regarding learning paradigms, recommender models have also embraced the two-stage paradigm of pre-training and fine-tuning [35, 41], as well as prompt learning for multiple tasks [14, 23]. Notably, earlier models such as P5 [23] and M6-Rec [14] have explored the potential of a unified LM to seamlessly handle various recommendation

tasks, such as click-through rate prediction, sequential recommendation, explanation generation, and conversational recommendation. However, it is worth noting that these earlier studies depend on medium-size LMs with limited generalization capabilities, often necessitating extensive recommendation data for training [23, 28].

Recently, the advent of powerful LLMs has significantly sparked technical advances in the field of recommender systems, opening up new horizons of leveraging LLMs for recommendation (shorted as LLM4Rec). These LLMs exhibit remarkable language understanding and generation capabilities, generalization prowess, and planning abilities [39, 67]. In light of these, in-context learning-based [15, 22] and instruction tuning-based recommendations [6, 62] with few-shot data, have garnered substantial attention within the research community. As such, it is really opportune to organize this tutorial to examine the progress of LLM4Rec.

In addition to the progress of LLM4Rec, this tutorial will demonstrate and analyze its critical challenges. The trustworthiness issues in recommendations generated by LLMs remain a primary concern, including problems like bias & fairness, privacy, and safety. Besides, efficiency is another essential aspect. LLMs are computationally demanding, thus deploying LLMs into real-time recommendation systems requires careful optimization and resource management. Additionally, the dynamic nature of user preferences necessitates timely model retraining to keep recommendations up to date. It is also challenging to achieve timely retraining in LLM4Rec. Lastly, recommendation data might encompass collaborative filtering signals, long interaction sequences from users, and high-dimensional features, which are also hard for LLMs to capture or utilize.

Last but not least, this tutorial will summarize the implications of previous work and present future research directions. In the future, it is essential to construct the evaluation benchmarks for LLM4Rec. Moreover, personalized AI-generated content might significantly facilitate the boom of generative recommendations [49]. As we delve into this tutorial, we aim to empower researchers and practitioners with knowledge and inspiration about LLM4Rec, inspiring further exploration and innovation in this emerging field.

**Necessity and timely of this tutorial.** LLM4Rec is a new topic, but developing vigorously. In the past years, there emerge extensive relevant papers [5, 6, 19, 56, 63]. Besides, some relevant workshops are hosted at CIKM'23<sup>7</sup>. Furthermore, researchers also actively organize special issues on TOIS<sup>8</sup> and TORS<sup>9</sup> (coming soon). It is undoubted that more research endeavors will surface in this field, leading to more reliable and powerful recommender systems. Given the potential and rapid development of LLM4Rec, it is a suitable time to conduct this tutorial, benefiting researchers and industrial developers to learn the progress and future directions of LLM4Rec.

### 3.2 Objective

The aim of this tutorial is to offer a comprehensive grasp of LLM development for recommendation, encompassing the motivation, significance, current advancements, hurdles, and future directions. Additionally, this tutorial strives to empower participants with the expertise and capabilities needed to harness LLMs for their own

<sup>4</sup><https://scholar.google.com/citations?user=Ma5DtmoAAAAJ&hl=en>.

<sup>5</sup><https://scholar.google.com.sg/citations?user=QePM4u8AAAAJ&hl=en>.

<sup>6</sup><https://scholar.google.com.sg/citations?user=X45Go24AAAAJ>.

<sup>7</sup><https://uobevents.eventsair.com/cikm2023/workshops>.

<sup>8</sup><https://dl.acm.org/journal/tois/calls-for-papers>.

<sup>9</sup><https://dl.acm.org/journal/tors/calls-for-papers>.

projects and make contributions to the progress of related domains, such as advertising.

### 3.3 Relevance

This tutorial holds strong relevance to the core domains of SIGIR-AP. Its subject matter is emerging as a crucial and thriving direction within the realm of recommendation, potentially serving as a source of inspiration for other information retrieval applications. Before proposing this tutorial for SIGIR-AP'23, there is a related tutorial [29] at RecSys'23. This RecSys tutorial primarily spotlights technical aspects and mostly covers the studies on medium-size LMs. In contrast, our tutorial focuses more on LLMs for recommendation, and discuss the underlying reasons and advantages of developing LLMs for recommendations. We will also delve into broader, distinct challenges within the realm of LLM4Rec. Besides, there are two workshops<sup>10</sup> covering similar topics scheduled for CIKM 2023. Both the RecSys tutorial and CIKM workshops have garnered huge attention. Given SIGIR-AP's close affiliation with RecSys and CIKM, we anticipate that our tutorial will also attract significant interest from the SIGIR-AP community.

### 3.4 Outline

We present an outline of the topics to be covered with timing:

- Introduction. (15 Min, Xiangnan He)
  - Organization of the tutorial.
  - Background of recommender systems.
  - The development of LMs.
- Progresses of LLM4Rec (60 Min, Keqin Bao, Jizhi Zhang)
  - Development from LMs to LLMs [67].
    - \* Basic prowess of LMs.
    - \* Augmented capabilities of LLMs [2, 39, 67].
  - Benefits of LMs and LLMs for recommendations [37].
  - LMs for recommendation.
    - \* Empowering recommendation with LMs' understanding and representation ability [28, 35, 61].
    - \* Empowering recommendation with LMs' generation ability [34, 64].
    - \* Empowering recommendation with LMs' multi-task ability [14, 23].
  - LLMs for recommendation.
    - \* In-context learning-based recommendation
      - Directly utilizing in-context learning for recommendation [15, 22, 30, 38]
      - Combining in-context learning with conventional recommender models [48, 57]
    - \* Tuning-based LLM4Rec [5, 6, 55, 62].
    - \* Chatting-enhanced LLM4Rec [20, 21]
    - \* Agent-based LLM4Rec [47, 53, 58]
- Q&A. (5 Min)
- Break. (10 Min)
- Open problems and challenges of LLM4Rec. (60 Min, Yang Zhang, Wenjie Wang)
  - Trustworthy LLM4Rec
    - \* Fairness & Bias in LLM4Rec [54, 63]
    - \* Privacy in LLM4Rec [10]

<sup>10</sup><https://uobevents.eventsair.com/cikm2023/workshops>.

- \* OOD & Robustness in LLM4Rec [7, 46, 63]
- \* Safety in LLM4Rec [3, 4]
- \* Explainable LLM4Rec [8, 65]
- \* Environmental Well-being in LLM4Rec [18, 44]
- Efficient LLM4Rec
  - \* Inference/Training costs [16, 17, 32, 40]
  - \* Deployment costs [60]
- LLM4Rec retraining and online training
  - \* Incremental learning for LLM4Rec [12]
  - \* Knowledge injection for LLM4Rec [9]
  - \* Data selection for LLM4Rec [36, 43]
- Recommendation data modeling in LLM4Rec
  - \* Modeling lifelong user interaction sequences [36]
  - \* Injecting collaborative filtering signals [5]
  - \* High-dimensional features & user-item interaction sparsity problems [26]
- Conclusions and future directions. (25 Min, Fuli Feng)
  - Evaluation & Benchmark in LLM4Rec [13, 27, 52]
  - Generative Recommendation with LLMs [49]
- Q&A. (5 Min)

### 3.5 Qualification of presenters

We have been working on recommender systems for a long time with a series of publications [24, 25, 50, 66] that emerged in top-tier conferences and journals. Recently, we have also released several well-known papers about LLM4Rec [5, 6, 63], some of which are published at RecSys'23 [6, 63]. Besides, we organized a relevant workshop titled "Recommendation with Generative Models" at CIKM'23<sup>11</sup> and a special issue on TOIS<sup>12</sup>. We are familiar with this field and will provide a systemic review and prospects of LLM4Rec. Moreover, our team has rich tutorial experience and has conducted more than 10 tutorials at various conferences including SIGIR, WWW, WSDM, CIKM, and RecSys [11, 33, 42, 51, 59]. We thus believe this tutorial would be attractive and insightful.

## 4 ORGANIZATION DETAILS AND MATERIALS

- **Duration.** The tutorial is planned as a 3-hour tutorial.
- **Interaction style.** This is a lecture-style tutorial.
- **Tutorial materials.** The slides will be released on the tutorial website. Organizers can obtain copyright permission.
- **Organization details.** The tutorial will be delivered in an on-site format, with a majority of the tutorial presenters to physically conducting and overseeing the tutorial. Additionally, we can prepare pre-recorded lectures if deemed necessary. Moreover, subject to approval, we are open to live-streaming the tutorial via popular video-streaming platforms.

## REFERENCES

- [1] Qingyao Ai et al. 2023. Information Retrieval Meets Large Language Models: A Strategic Report from Chinese IR Community. *AI Open* 4 (2023), 80–90.
- [2] Rohan Anil et al. 2023. Palm 2 technical report. *arXiv preprint arXiv:2305.10403* (2023).
- [3] Yuntao Bai et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073* (2022).
- [4] Yuntao Bai et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862* (2022).

<sup>11</sup><https://rgm-cikm23.github.io/>.

<sup>12</sup><https://dl.acm.org/journal/tois/calls-for-papers>.

- [5] Keqin Bao, Jizhi Zhang, Wenjie Wang, Yang Zhang, Zhengyi Yang, Yancheng Luo, Fuli Feng, Xiangnan He, and Qi Tian. 2023. A Bi-Step Grounding Paradigm for Large Language Models in Recommendation Systems. *CoRR* abs/2308.08434 (2023). <https://doi.org/10.48550/arXiv.2308.08434> arXiv:2308.08434
- [6] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. TALLRec: An Effective and Efficient Tuning Framework to Align Large Language Model with Recommendation. In *RecSys*. ACM, 1007–1014. <https://doi.org/10.1145/3604915.3608857>
- [7] Lukas Berglund et al. 2023. The Reversal Curse: LLMs trained on  $\tilde{A}$  is  $\tilde{B}$  fail to learn  $\tilde{B}$  is  $A$ . *arXiv preprint arXiv:2309.12288* (2023).
- [8] Steven Bills et al. 2023. Language models can explain neurons in language models. URL <https://openaipublic.blob.core.windows.net/neuron-explainer/paper/index.html>. (Date accessed: 14.05. 2023) (2023).
- [9] Boxi Cao et al. 2023. The Life Cycle of Knowledge in Big Language Models: A Survey. *arXiv preprint arXiv:2303.07616* (2023).
- [10] Aldo Gael Carranza et al. 2023. Privacy-Preserving Recommender Systems with Synthetic Query Generation using Differentially Private Large Language Models. *arXiv preprint arXiv:2305.05973* (2023).
- [11] Jiawei Chen, Xiang Wang, Fuli Feng, and Xiangnan He. 2021. Bias Issues and Solutions in Recommender System. In *RecSys*. 825–827.
- [12] Wuyang Chen et al. 2023. Lifelong language pretraining with distribution-specialized experts. In *ICML*. PMLR, 5383–5395.
- [13] Xu Chen et al. 2023. REASONER: An Explainable Recommendation Dataset with Multi-aspect Real User Labeled Ground Truths Towards more Measurable Explainable Recommendation. *arXiv preprint arXiv:2303.00168* (2023).
- [14] Zeyu Cui et al. 2022. M6-rec: Generative pretrained language models are opened recommender systems. *arXiv preprint arXiv:2205.08084* (2022).
- [15] Sunhao Dai et al. 2023. Uncovering ChatGPT’s Capabilities in Recommender Systems. In *RecSys*, Jie Zhang, Li Chen, Shlomo Berkovsky, Min Zhang, Tommaso Di Noia, Justin Basilico, Luiz Pizzato, and Yang Song (Eds.). ACM, 1126–1132.
- [16] Tri Dao et al. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems* 35 (2022), 16344–16359.
- [17] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314* (2023).
- [18] Nan Du et al. 2022. Glam: Efficient scaling of language models with mixture-of-experts. In *ICML*. PMLR, 5547–5569.
- [19] Wenqi Fan, Zihuai Zhao, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Jiliang Tang, and Qing Li. 2023. Recommender systems in the era of large language models (llms). *arXiv preprint arXiv:2307.02046* (2023).
- [20] Yue Feng et al. 2023. A Large Language Model Enhanced Conversational Recommender System. *arXiv preprint arXiv:2308.06212* (2023).
- [21] Luke Friedman et al. 2023. Leveraging Large Language Models in Conversational Recommender Systems. *arXiv preprint arXiv:2305.07961* (2023).
- [22] Yunfan Gao et al. 2023. Chat-REC: Towards Interactive and Explainable LLMs-Augmented Recommender System. *CoRR* abs/2303.14524 (2023).
- [23] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In *RecSys*. 299–315.
- [24] Xiangnan He et al. 2017. Neural Collaborative Filtering. In *WWW*, Rick Barrett, Rick Cummings, Eugene Agichtein, and Evgeniy Gabrilovich (Eds.). ACM, 173–182.
- [25] Xiangnan He et al. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *SIGIR*. 639–648.
- [26] Xiangnan He and Tat-Seng Chua. 2017. Neural Factorization Machines for Sparse Predictive Analytics. In *SIGIR*. ACM, 355–364.
- [27] Zhankui He et al. 2023. Large language models as zero-shot conversational recommenders. *arXiv preprint arXiv:2308.10053* (2023).
- [28] Yupeng Hou, Shanlei Mu, Wayne Xin Zhao, Yaliang Li, Bolin Ding, and Ji-Rong Wen. 2022. Towards universal sequence representation learning for recommender systems. In *SIGKDD*. 585–593.
- [29] Wenyue Hua, Lei Li, Shuyuan Xu, Li Chen, and Yongfeng Zhang. 2023. Tutorial on Large Language Models for Recommendation. In *RecSys*. 1281–1283.
- [30] Wang-Cheng Kang et al. 2023. Do LLMs Understand User Preferences? Evaluating LLMs on User Rating Prediction. *CoRR* abs/2305.06474 (2023).
- [31] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *ICDM*. IEEE, 197–206.
- [32] Woosuk Kwon et al. 2023. Efficient Memory Management for Large Language Model Serving with PagedAttention. *arXiv preprint arXiv:2309.06180* (2023).
- [33] Wenqiang Lei, Xiangnan He, Maarten de Rijke, and Tat-Seng Chua. 2020. Conversational recommendation: Formulation, methods, and evaluation. In *SIGIR*. 2425–2428.
- [34] Jinming Li et al. 2023. GPT4Rec: A generative framework for personalized recommendation and user interests interpretation. *arXiv preprint arXiv:2304.03879* (2023).
- [35] Jiacheng Li et al. 2023. Text Is All You Need: Learning Language Representations for Sequential Recommendation. In *SIGKDD*. ACM, 1258–1267.
- [36] Jianghao Lin et al. 2023. ReLLa: Retrieval-enhanced Large Language Models for Lifelong Sequential Behavior Comprehension in Recommendation. *arXiv preprint arXiv:2308.11131* (2023).
- [37] Jianghao Lin, Xinyi Dai, Yunjia Xi, Weiwen Liu, Bo Chen, Xiangyang Li, Chenxu Zhu, Huifeng Guo, Yong Yu, Ruiming Tang, et al. 2023. How Can Recommender Systems Benefit from Large Language Models: A Survey. *arXiv preprint arXiv:2306.05817* (2023).
- [38] Junling Liu et al. 2023. Is ChatGPT a Good Recommender? A Preliminary Study. *CoRR* abs/2304.10149 (2023).
- [39] OpenAI. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [40] Bo Peng et al. 2023. RWKV: Reinventing RNNs for the Transformer Era. *arXiv preprint arXiv:2305.13048* (2023).
- [41] Zhaopeng Qiu et al. 2021. U-BERT: Pre-training User Representations for Improved Recommendation. In *AAAI*. AAAI Press, 4320–4327.
- [42] Zhaochun Ren, Xiangnan He, Dawei Yin, and Maarten de Rijke. 2018. Information Discovery in E-commerce: Half-day SIGIR 2018 Tutorial. In *SIGIR*. 1379–1382.
- [43] Zhiqiang Shen et al. 2023. SlimPajama-DC: Understanding Data Combinations for LLM Training. *arXiv preprint arXiv:2309.10818* (2023).
- [44] Giuseppe Spillo et al. 2023. Towards Sustainability-aware Recommender Systems: Analyzing the Trade-off Between Algorithms Performance and Carbon Footprint. In *RecSys*. 856–862.
- [45] Fei Sun et al. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 1441–1450.
- [46] Jindong Wang et al. 2023. On the robustness of chatgpt: An adversarial and out-of-distribution perspective. *arXiv preprint arXiv:2302.12095* (2023).
- [47] Lei Wang et al. 2023. RecAgent: A Novel Simulation Paradigm for Recommender Systems. *arXiv preprint arXiv:2306.02552* (2023).
- [48] Lei Wang and Ee-Peng Lim. 2023. Zero-Shot Next-Item Recommendation using Large Pretrained Language Models. *CoRR* abs/2304.03153 (2023).
- [49] Wenjie Wang et al. 2023. Generative recommendation: Towards next-generation recommender paradigm. *arXiv preprint arXiv:2304.03516* (2023).
- [50] Xiang Wang et al. 2019. Neural Graph Collaborative Filtering. In *SIGIR*, Benjamin Piwowarski, Max Chevalier, Éric Gaussier, Yoelle Maarek, Jian-Yun Nie, and Falk Scholer (Eds.). ACM, 165–174.
- [51] Xiang Wang, Xiangnan He, and Tat-Seng Chua. 2020. Learning and reasoning on graph for recommendation. In *WSDM*. 890–893.
- [52] Xiaolei Wang, Xinyu Tang, Wayne Xin Zhao, Jingyuan Wang, and Ji-Rong Wen. 2023. Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models. *arXiv preprint arXiv:2305.13112* (2023).
- [53] Yancheng Wang et al. 2023. Recmind: Large language model powered agent for recommendation. *arXiv preprint arXiv:2308.14296* (2023).
- [54] Yifan Wang et al. 2023. A survey on the fairness of recommender systems. *ACM Transactions on Information Systems* 41, 3 (2023), 1–43.
- [55] Likang Wu et al. 2023. Exploring Large Language Model for Graph Data Understanding in Online Job Recommendations. *CoRR* abs/2307.05722 (2023).
- [56] Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, et al. 2023. A Survey on Large Language Models for Recommendation. *arXiv preprint arXiv:2305.19860* (2023).
- [57] Yunjia Xi et al. 2023. Towards Open-World Recommendation with Knowledge Augmentation from Large Language Models. *CoRR* abs/2306.10933 (2023).
- [58] Zhiheng Xi et al. 2023. The Rise and Potential of Large Language Model Based Agents: A Survey. *arXiv preprint arXiv:2309.07864* (2023).
- [59] Jun Xu, Xiangnan He, and Hang Li. 2018. Deep learning for matching in search and recommendation. In *SIGIR*. 1365–1368.
- [60] Minrui Xu et al. 2023. Unleashing the power of edge-cloud generative ai in mobile networks: A survey of aigc services. *arXiv preprint arXiv:2303.16129* (2023).
- [61] Zheng Yuan et al. 2023. Where to Go Next for Recommender Systems? ID-vs. Modality-based Recommender Models Revisited. In *SIGIR*, Hsin-Hsi Chen, Wei-Jou (Edward) Duh, Hen-Hsen Huang, Makoto P. Kato, Josiane Mothe, and Barbara Pobleto (Eds.). ACM, 2639–2649.
- [62] Junjie Zhang et al. 2023. Recommendation as instruction following: A large language model empowered recommendation approach. *arXiv preprint arXiv:2305.07001* (2023).
- [63] Jizhi Zhang, Keqin Bao, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Is ChatGPT Fair for Recommendation? Evaluating Fairness in Large Language Model Recommendation. In *RecSys*. ACM, 993–999. <https://doi.org/10.1145/3604915.3608860>
- [64] Yuhui Zhang et al. 2021. Language models as recommender systems: Evaluations and limitations. (2021).
- [65] Yongfeng Zhang, Xu Chen, et al. 2020. Explainable recommendation: A survey and new perspectives. *Foundations and Trends® in Information Retrieval* 14, 1 (2020), 1–101.
- [66] Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021. Causal Intervention for Leveraging Popularity Bias in Recommendation. In *SIGIR*. 11–20.
- [67] Wayne Xin Zhao et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023).