Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)
Early Career Track

Transforming Recommender Systems: Balancing Personalization, Fairness, and Human Values

Julia Neidhardt
CD Lab for Recommender Systems, TU Wien, Austria
julia.neidhardt@tuwien.ac.at

Abstract

Recent advancements in recommender systems highlight the importance of metrics beyond accuracy, including diversity, serendipity, and fairness. This paper discusses various aspects of modern recommender systems, focusing on challenges such as preference elicitation, the complexity of human decision-making, and multi-domain applicability. The integration of Generative AI and Large Language Models offers enhanced personalization capabilities but also raises concerns regarding transparency and fairness. This work examines ongoing research efforts aimed at developing transparent, fair, and contextually aware systems. Our approach seeks to prioritize user wellbeing and responsibility, contributing to a more equitable and functional digital environment through advanced technologies and interdisciplinary insights.

1 Introduction

Recommender systems (RSs) are software tools and techniques designed to suggest items of interest to users by leveraging various data sources such as user profiles, item attributes, and historical interactions [Ricci *et al.*, 2022]. They are widely used in e-commerce, entertainment, social media, and other domains to enhance user experience and drive business growth by providing personalized suggestions. Traditionally, these systems have focused on maximizing accuracy, aiming to predict the most relevant items for a user [Jannach *et al.*, 2016].

However, focusing solely on accuracy can lead to significant issues, such as the creation of filter bubbles and echo chambers, where users are only exposed to a narrow set of content that reinforces their existing preferences. Although the extent of this effect is uncertain in practice, it can still limit the diversity of information and perspectives available to users [Budak *et al.*, 2024].

Additionally, RSs should address fairness to ensure that all user groups and content creators are treated equitably. Fairness is critical to prevent discrimination and bias, which can lead to unequal access to information and opportunities [Ekstrand *et al.*, 2022].

Beyond accuracy and fairness, other metrics such as diversity, serendipity, and user wellbeing have become increasingly important. Diversity ensures that users are exposed to a wide range of content, enhancing their experience and preventing the narrowing of their informational and cultural horizons [Castells *et al.*, 2015]. Serendipity introduces unexpected and novel items to users, making interactions with the system more engaging and enjoyable [Kaminskas and Bridge, 2017]. Finally, user wellbeing focuses on the ethical and psychological impacts of recommendations, promoting content that supports positive mental health and overall satisfaction [Stray *et al.*, 2022].

The integration of Generative AI and Large Language Models (LLMs) into RSs has brought considerable advancements and new challenges. LLMs have demonstrated substantial performance enhancements in various applications by leveraging their ability to process and generate human-like text [Zhao *et al.*, 2024]. These models, trained on extensive datasets, can predict and generate relevant content, thereby improving the personalization and contextual relevance of recommendations.

However, their deployment also raises concerns regarding explainabilty, transparency, and ethical considerations, particularly in ensuring that generated recommendations are accurate, fair, and unbiased.

The application of LLMs across various domains offers both opportunities and challenges. In the news and media domain, for example, LLMs can enhance user engagement by providing personalized content, but ensuring diversity and addressing bias is crucial to prevent misinformation and echo chambers. In e-commerce, LLMs can help improve the shopping experience with tailored product suggestions, but addressing fairness and bias is essential to ensure fair access to products. In the tourism sector, LLMs offer personalized travel recommendations, enhancing trip planning, but maintaining accuracy and relevance given the dynamic nature of travel information requires innovative and hybrid approaches.

Moreover, ethical considerations in deploying these models call for continuous monitoring and refinement. While LLMs have transformative potential, addressing challenges of accuracy, diversity, fairness, and transparency is fundamental.

2 Our Approach and Contributions

At the Christian Doppler Lab for Recommender Systems[1] at TU Wien, Austria, we address the complex challenges faced by modern RSs by integrating advanced technologies and interdisciplinary insights. Our research aims to enhance both the impact and ethical standards of RSs through the following approaches:

2.1 Enhancing Diversity and Serendipity

Recognizing the importance of moving beyond accuracy to include diversity and serendipity, we aim to:

- Conceptualize and develop diversification algorithms and serendipity metrics to expose users to a broad range of content, introducing unexpected and novel items to make interactions more engaging and enjoyable.

These techniques ensure users receive a wide spectrum of recommendations that enhance their overall experience and promote media pluralism, helping to avoid the pitfalls of overly personalized content and contributing to a more diverse information ecosystem [Basso *et al.*, 2023; Knees *et al.*, 2023; Modre *et al.*, 2023; Neidhardt and Sertkan, 2022; Nalis *et al.*, 2024; Sertkan *et al.*, 2022; Sertkan and Neidhardt, 2022; Kolb *et al.*, 2023a].

2.2 Fairness and Ethical Considerations

Ensuring fairness and ethical standards is crucial for maintaining user trust and equity. Our strategies include:

- Developing fairness-aware algorithms, implementing bias mitigation techniques, and evaluating community-level fairness using group detection approaches to ensure equitable treatment of all user groups and content creators.

By addressing these ethical considerations, we aim to create RSs that promote inclusivity and fairness across various domains, preventing algorithmic gatekeeping that can lead to discriminatory outcomes. This responsibility is fundamental to fostering a trustworthy and deliberative digital environment [Huebner *et al.*, 2024; Kolb *et al.*, 2022; Modre *et al.*, 2023; Pachinger *et al.*, 2023].

2.3 Integration of Psychological Insights

Incorporating psychological insights allows us to tailor recommendations more effectively, helping and supporting users rather than manipulating them. Our methods include:

- Transitioning from simplistic emotion models to dynamic, contextualized understandings based on the theory of constructed emotions, adjusting recommendations according to users' current emotional states, and incorporating expressed emotions within user behavior and item descriptions to improve accuracy and engagement.

This approach significantly enhances personalization by considering the complexity and context of emotions and the nuances of language [Modre *et al.*, 2023; Nalis and Neidhardt, 2023; Pachinger *et al.*, 2024; Sertkan and Neidhardt, 2022; Sertkan and Neidhardt, 2023; Wagne *et al.*, 2024].

[1]https://recsys-lab.at

2.4 User-Centric and Context-Aware Recommendations

Focusing on the user's current context and long-term preferences, our research emphasizes:

- Utilizing contextual data such as location, time, and social context to continuously adapt recommendations based on changing user preferences. We advance various algorithms to provide relevant suggestions and develop models to classify user roles and understand their behavior.
- Implementing a comprehensive multi-perspective analysis to evaluate user experiences, develop cross-domain recommendations, compare image captioning models, and explore intermediate fusion strategies for session-based recommender systems (SBRS).

This methodology ensures that recommendations remain relevant and personalized, addressing the dynamic nature of user behavior and preferences [Aayesha *et al.*, 2024; Godolja *et al.*, 2024; Scholz *et al.*, 2024].

2.5 Explainability and Transparency

We are committed to enhancing the explainability and transparency of RSs to build user trust. Our efforts involve:

- Integrating methods that provide clear and understandable explanations for recommendations and ensuring transparency mechanisms so users can understand how and why certain recommendations are made.

We believe that transparency is essential for fostering accountability and trust in the recommendation process [Knees *et al.*, 2023; Modre *et al.*, 2023; Kolb *et al.*, 2023a].

2.6 Combining Local and Global Knowledge

To improve the relevance and accuracy of recommendations, we combine the extensive capabilities of LLMs with local contextual data. This hybrid approach allows us to:

- Utilize global knowledge from LLMs while incorporating specific local information to provide recommendations that are both globally informed and locally relevant.

This method addresses issues of data sparsity and hallucination in LLMs, ensuring recommendations are factually accurate and contextually relevant [Kolb *et al.*, 2023b].

2.7 Social-Aware and Conversational Methods

Implementing advanced user modeling techniques, we enhance the recommendation process through:

- Engaging users in interactive dialogues, incorporating social context, integrating LLMs in Conversational Recommender Systems (CRSs), and considering group dynamics to provide more personalized and engaging recommendations.

These methods are applied to improve preference elicitation and user satisfaction and engagement [Aayesha *et al.*, 2024].

2.8 Digital Humanism (DigHum)

Our research aligns with Digital Humanism principles, emphasizing human-centric values in digital technology development [Werthner, 2024]. This approach ensures that RSs promote human dignity, autonomy, and inclusivity by:

- Designing algorithms that prioritize fairness, transparency, user wellbeing, diversity, and accountability.

Additionally, the Christian Doppler Lab for Recommender Systems at TU Wien is connected with the UNESCO Chair on Digital Humanism[2], with Julia Neidhardt serving as co-chairholder, underscoring our commitment to these values. This collaboration aims to advance RSs that are responsible and supportive of user autonomy, fostering a more equitable and inclusive digital environment in line with Digital Humanism principles [Knees *et al.*, 2023; Neidhardt *et al.*, 2022; Werthner, 2024].

3 Challenges

Despite significant advancements, integrating advanced technologies and interdisciplinary insights into RSs brings forth several challenges that need to be addressed to enhance both effectiveness and ethical standards.

3.1 User Modeling

Creating accurate and comprehensive user models is a complex task. These models need to understand users on multiple levels. Achieving a balance between latent user models, which use deep learning to map users into a latent vector space, and domain models, which rely on expert knowledge, is essential. Combining these models into a hybrid approach is particularly challenging, especially when it comes to distinguishing between long-term and short-term user preferences [Zhang *et al.*, 2021].

3.2 Preference Elicitation and Cold-Start Problem

Effective user modeling requires accurately representing user preferences, which is difficult due to the cold-start problem, variability in user intents, and the complexity of expressing preferences.

Innovative methods such as gamified and conversational elicitation techniques are needed to engage users without overwhelming them. Adaptive user models that refine preferences during interactions can also address these issues. Additionally, LLMs can assist in mitigating cold-start problems by generating initial user preferences and content recommendations based on minimal input data, leveraging their extensive training on diverse datasets [Schein *et al.*, 2002; Amatriain and Basilico, 2015; Sanner *et al.*, 2023].

3.3 Human Decision-Making Complexity

Human decision-making is influenced by various factors, including emotions and personality traits, often deviating from purely logical processes. RSs must balance the trade-off between effort and accuracy in user decisions, considering hidden decision factors through techniques like clustering and dimensionality reduction. Leveraging LLMs' global knowledge can help, but capturing these dimensions and effectively integrating them presents a significant challenge [Jannach *et al.*, 2010; Chen *et al.*, 2013].

[2]https://informatics.tuwien.ac.at/digital-humanism/

3.4 Multi-Domain Applicability

RSs typically operate within specific domains, but integrating insights across varied domains presents a challenge. Each domain has unique characteristics that influence recommendation effectiveness. Designing systems that recognize and adapt to these nuances, such as cultural and regional variations, is crucial. Systematically defining and categorizing domains and exploring multi-domain scenarios is essential to avoid overly generalized models that miss nuanced differences [Ricci *et al.*, 2022; Dacrema *et al.*, 2022].

3.5 Resource Availability

Sufficient data and computational resources are crucial for training, tuning, and testing models. As RSs become more sophisticated, the demand for GPU power and scalability increases. Ensuring rapid response times while balancing computational efficiency, model accuracy, and resource availability is essential [Zhang *et al.*, 2022; Zhang *et al.*, 2019].

3.6 Generative AI and LLMs

The deployment of LLMs in RSs raises ethical concerns, particularly related to hallucination (even when using local knowledge [Chen *et al.*, 2024]), toxic language generation, and biases in training data. These models, while powerful, can generate content that is not always factually accurate or free from harmful biases. Ensuring that LLM-generated recommendations adhere to ethical standards and promote positive user experiences is a key challenge.

This involves implementing robust bias mitigation techniques and transparency mechanisms to provide users with clear and understandable explanations for the recommendations they receive. Additionally, LLMs can sometimes manipulate user preferences or disseminate fake information, which requires continuous monitoring and adjustment to maintain trustworthiness and reliability [Fan *et al.*, 2023].

3.7 Explainability and Transparency

Enhancing the explainability and transparency of RSs is crucial for building user trust. Methods that provide clear and understandable explanations for recommendations, as well as transparency mechanisms that allow users to understand how and why certain recommendations are made, are essential. This fosters accountability and trust in the recommendation process [Wang *et al.*, 2023].

4 Conclusion

RSs have evolved to include metrics such as diversity, fairness, and user wellbeing, beyond just accuracy. The integration of Generative AI and LLMs offers the potential for improved performance but also brings new challenges. At the Christian Doppler Lab for Recommender Systems at TU Wien, we address these challenges by integrating advanced technologies and interdisciplinary insights.

We focus on enhancing diversity and serendipity to enrich user experiences, ensuring fairness to treat all users and content creators adequately, and incorporating psychological insights to develop recommendations that truly support users. Our user-centric and context-aware methods ensure relevance, while our commitment to explainability and transparency builds trust.

By combining global LLM knowledge with local data, we improve recommendation accuracy and address data sparsity. In line with Digital Humanism principles, we promote fairness, transparency, and inclusivity. Despite ongoing challenges, our research aims to create RSs that users find both useful and enjoyable.

Acknowledgments

I would like to thank my team at the Christian Doppler Lab for RSs at TU Wien for their excellent work and dedication. The financial support by the Austrian Federal Ministry of Labour and Economy, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association is gratefully acknowledged.

References

[Aayesha *et al.*, 2024] Aayesha Aayesha, Muhammad Afzaal, and Julia Neidhardt. User experience of recommender system: A user study of social-aware fashion recommendations system. In *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*, pages 356–361, 2024.

[Amatriain and Basilico, 2015] Xavier Amatriain and Justin Basilico. Recommender systems in industry: A netflix case study. In *Recommender systems handbook*, pages 385–419. Springer, 2015.

[Basso *et al.*, 2023] Linda Basso, Irina Nalis, and Julia Neidhardt. News diversity and well-being – an experimental exploration of diversity-aware recommender systems (forthcoming). In *5th FAccTRec Workshop on Responsible Recommendation at RecSys 2023*, 2023.

[Budak *et al.*, 2024] Ceren Budak, Brendan Nyhan, David M Rothschild, Emily Thorson, and Duncan J Watts. Misunderstanding the harms of online misinformation. *Nature*, 630(8015):45–53, 2024.

[Castells *et al.*, 2015] Pablo Castells, Neil J Hurley, and Saul Vargas. Novelty and diversity in recommender systems. In *Recommender Systems Handbook*, pages 881–918. Springer, 2015.

[Chen *et al.*, 2013] Li Chen, Marco de Gemmis, Alexander Felfernig, Pasquale Lops, Francesco Ricci, and Giovanni Semeraro. Human decision making and recommender systems. *ACM Trans. Interact. Intell. Syst.*, 3(3), oct 2013.

[Chen *et al.*, 2024] Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. Benchmarking large language models in retrieval-augmented generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17754–17762, 2024.

[Dacrema *et al.*, 2022] Maurizio Ferrari Dacrema, Iván Cantador, Ignacio Fernández-Tobías, Shlomo Berkovsky, and Paolo Cremonesi. *Design and Evaluation of Cross-Domain Recommender Systems*, pages 485–516. Springer US, New York, NY, 2022.

[Ekstrand *et al.*, 2022] Michael D. Ekstrand, Anubrata Das, Robin Burke, and Fernando Diaz. *Fairness in Recommender Systems*, pages 679–707. Springer US, 2022.

[Fan *et al.*, 2023] Wenqi Fan, Zihuai Zhao, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Zhen Wen, Fei Wang, Xiangyu Zhao, Jiliang Tang, and Qing Li. Recommender systems in the era of large language models (llms), 2023.

[Godolja *et al.*, 2024] Dante Godolja, Thomas E. Kolb, and Julia Neidhardt. Unlocking the potential of content-based restaurant recommender systems. In *Information and Communication Technologies in Tourism 2024*, page 239–244, Cham, 2024. Springer Nature Switzerland.

[Huebner *et al.*, 2024] Blake Huebner, Thomas Elmar Kolb, and Julia Neidhardt. Evaluating group fairness in news recommendations: A comparative study of algorithms and metrics. In *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*, pages 337–346, 2024.

[Jannach *et al.*, 2010] Dietmar Jannach, Markus Zanker, Alexander Felfernig, and Gerhard Friedrich. *Online consumer decision making*, page 234–252. Cambridge University Press, 2010.

[Jannach *et al.*, 2016] Dietmar Jannach, Paul Resnick, Alexander Tuzhilin, and Markus Zanker. Recommender systems—beyond matrix completion. *Communications of the ACM*, 59(11):94–102, 2016.

[Kaminskas and Bridge, 2017] Marius Kaminskas and Derek Bridge. Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 7(1):2, 2017.

[Knees *et al.*, 2023] Peter Knees, Julia Neidhardt, and Irina Nalis. *Recommender Systems: Techniques, Effects, and Measures Towards Pluralism and Fairness*. Springer, 2023.

[Kolb *et al.*, 2022] Thomas E. Kolb, Irina Nalis, Mete Sertkan, and Julia Neidhardt. The role of bias in news recommendation in the perception of the covid-19 pandemic. In *5th FAccTRec Workshop on Responsible Recommendation at RecSys 2022 Seattle, WA, USA, September 17, 2022,*, 2022.

[Kolb *et al.*, 2023a] Thomas E. Kolb, Irina Nails, and Julia Neidhardt. Like a skilled dj - an expert study on news recommendations beyond accuracy. In *Proceedings of the 11th International Workshop on News Recommendation and Analytics (INRA) in conjunction with ACM RecSys 2023*. CEUR-WS, 2023.

[Kolb *et al.*, 2023b] Thomas E. Kolb, Ahmadou Wagne, Mete Sertkan, and Julia Neidhardt. Potentials of combining local knowledge and llms for recommender systems.

In *Proceedings of the 5th Edition of Knowledge-aware and Conversational Recommender Systems (KaRS) in conjunction with ACM RecSys 2023*. CEUR-WS, 2023.

[Modre *et al.*, 2023] Laura Modre, Julia Neidhardt, and Irina Nalis. Value-based nudging in news recommender systems – results from an experimental user study (forthcoming). In *Proceedings of the First Workshop on the Normative Design and Evaluation of Recommender Systems in conjunction with ACM RecSys 2023*. CEUR-WS, 2023.

[Nalis and Neidhardt, 2023] Irina Nalis and Julia Neidhardt. Not facial expression, nor fingerprint – acknowledging complexity and context in emotion research for human-centered personalization and adaptation. In *Adjunct Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization*, page 325–330, New York, 2023. Association for Computing Machinery.

[Nalis *et al.*, 2024] Irina Nalis, Tobias Sippl, Thomas Elmar Kolb, and Julia Neidhardt. Navigating serendipity-an experimental user study on the interplay of trust and serendipity in recommender systems. In *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*, pages 386–393, 2024.

[Neidhardt and Sertkan, 2022] Julia Neidhardt and Mete Sertkan. Towards an approach for analyzing dynamic aspects of bias and beyond-accuracy measures. In *Advances in Bias and Fairness in Information Retrieval: 3rd International Workshop, BIAS 2022, Stavanger, Norway, April 10, 2022*, pages 35–42. Springer, 2022.

[Neidhardt *et al.*, 2022] Julia Neidhardt, Hannes Werthner, and Stefan Woltran. It is simple, it is complicated. *Perspectives on Digital Humanism*, page 335, 2022.

[Pachinger *et al.*, 2023] Pia Pachinger, Allan Hanbury, Julia Neidhardt, and Anna Planitzer. Toward disambiguating the definitions of abusive, offensive, toxic, and uncivil comments. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 107–113. Association for Computational Linguistics, May 2023.

[Pachinger *et al.*, 2024] Pia Pachinger, Janis Goldzycher, Anna Maria Planitzer, Wojciech Kusa, Allan Hanbury, and Julia Neidhardt. Austrotox: A dataset for target-based austrian german offensive language detection. *arXiv preprint arXiv:2406.08080*, 2024.

[Ricci *et al.*, 2022] Francesco Ricci, Lior Rokach, and Bracha Shapira. *Recommender Systems: Techniques, Applications, and Challenges*, pages 1–35. Springer US, New York, NY, 2022.

[Sanner *et al.*, 2023] Scott Sanner, Krisztian Balog, Filip Radlinski, Ben Wedin, and Lucas Dixon. Large language models are competitive near cold-start recommenders for language-and item-based preferences. In *Proceedings of the 17th ACM conference on recommender systems*, pages 890–896, 2023.

[Schein *et al.*, 2002] Andrew I. Schein, Alexandrin Popescul, Lyle H. Ungar, and David M. Pennock. Methods and metrics for cold-start recommendations. In *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '02, page 253–260, New York, NY, USA, 2002. Association for Computing Machinery.

[Scholz *et al.*, 2024] Felix Scholz, Thomas Elmar Kolb, and Julia Neidhardt. Classifying user roles in online news forums: A model for user interaction and behavior analysis. In *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*, pages 240–249, 2024.

[Sertkan and Neidhardt, 2022] Mete Sertkan and Julia Neidhardt. Exploring expressed emotions for neural news recommendation. In *Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization*, UMAP '22 Adjunct, page 22–28, New York, NY, USA, 2022. Association for Computing Machinery.

[Sertkan and Neidhardt, 2023] Mete Sertkan and Julia Neidhardt. On the effect of incorporating expressed emotions in news articles on diversity within recommendation models. In *Proceedings of the 11th International Workshop on News Recommendation and Analytics (INRA) in conjunction with ACM RecSys 2023*. CEUR-WS, 2023.

[Sertkan *et al.*, 2022] Mete Sertkan, Sophia Althammer, Sebastian Hofstätter, and Julia Neidhardt. Diversifying sentiments in news recommendation. In *Perspectives on the Evaluation of Recommender Systems Workshop (PERSPECTIVES 2022)*, 2022.

[Stray *et al.*, 2022] Jonathan Stray, Alon Halevy, Parisa Assar, Dylan Hadfield-Menell, Craig Boutilier, Amar Ashar, Chloe Bakalar, Lex Beattie, Michael Ekstrand, Claire Leibowicz, et al. Building human values into recommender systems: An interdisciplinary synthesis. *ACM Transactions on Recommender Systems*, 2022.

[Wagne *et al.*, 2024] Ahmadou Wagne, Julia Neidhardt, and Thomas E. Kolb. Popaut: An annotated corpus for populism detection in austrian news comments. In Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue, editors, *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, page 12879–12892, Torino, Italy, 2024. ELRA and ICCL.

[Wang *et al.*, 2023] Shoujin Wang, Xiuzhen Zhang, Yan Wang, and Francesco Ricci. Trustworthy recommender systems. *ACM Trans. Intell. Syst. Technol.*, oct 2023.

[Werthner, 2024] Hannes Werthner. Digital transformation, digital humanism: What needs to be done. *Hannes Werthner· Carlo Ghezzi· Jeff Kramer· Julian Nida-Rümelin· Bashar Nuseibeh· Erich Prem·*, page 115, 2024.

[Zhang *et al.*, 2019] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys (CSUR)*, 52(1):1–38, 2019.

[Zhang *et al.*, 2021] Shuai Zhang, Yi Tay, Lina Yao, Aixin Sun, and Ce Zhang. Deep learning for recommender systems. In *Recommender systems handbook*, pages 173–210. Springer, 2021.

[Zhang *et al.*, 2022] Shuai Zhang, Yi Tay, Lina Yao, Aixin Sun, and Ce Zhang. *Deep Learning for Recommender Systems*, pages 173–210. Springer US, New York, NY, 2022.

[Zhao *et al.*, 2024] Zihuai Zhao, Wenqi Fan, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Zhen Wen, Fei Wang, Xiangyu Zhao, Jiliang Tang, and Qing Li. Recommender systems in the era of large language models (llms). *IEEE Transactions on Knowledge and Data Engineering*, pages 1–20, 2024.