



Tutorial on Large Language Models for Recommendation

Wenyue Hua
Department of Computer Science
Rutgers University
wenyue.hua@rutgers.edu

Lei Li
Department of Computer Science
Hong Kong Baptist University
cslleili@comp.hkbu.edu.hk

Shuyuan Xu
Department of Computer Science
Rutgers University
shuyuan.xu@rutgers.edu

Li Chen
Department of Computer Science
Hong Kong Baptist University
lichen@comp.hkbu.edu.hk

Yongfeng Zhang
Department of Computer Science
Rutgers University
yongfeng.zhang@rutgers.edu

ABSTRACT

Foundation Models such as Large Language Models (LLMs) have significantly advanced many research areas. In particular, LLMs offer significant advantages for recommender systems, making them valuable tools for personalized recommendations. For example, by formulating various recommendation tasks such as rating prediction, sequential recommendation, straightforward recommendation, and explanation generation into language instructions, LLMs make it possible to build universal recommendation engines that can handle different recommendation tasks. Additionally, LLMs have a remarkable capacity for understanding natural language, enabling them to comprehend user preferences, item descriptions, and contextual information to generate more accurate and relevant recommendations, leading to improved user satisfaction and engagement. This tutorial introduces Foundation Models such as LLMs for recommendation. We will introduce how recommender system advanced from shallow models to deep models and to large models, how LLMs enable generative recommendation in contrast to traditional discriminative recommendation, and how to build LLM-based recommender systems. We will cover multiple perspectives of LLM-based recommendation, including data preparation, model design, model pre-training, fine-tuning and prompting, multi-modality and multi-task learning, as well as trustworthy perspectives of LLM-based recommender systems such as fairness and transparency.

KEYWORDS

Recommendation; Large Language Models; Foundation Models

ACM Reference Format:

Wenyue Hua, Lei Li, Shuyuan Xu, Li Chen, and Yongfeng Zhang. 2023. Tutorial on Large Language Models for Recommendation. In *Seventeenth ACM Conference on Recommender Systems (RecSys '23), September 18–22, 2023, Singapore, Singapore*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3604915.3609494>

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.

RecSys '23, September 18–22, 2023, Singapore, Singapore

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0241-9/23/09...\$15.00

<https://doi.org/10.1145/3604915.3609494>

1 MOTIVATION, RELEVANCE AND OUTLINE

Recently, Large Language Models (LLMs) have emerged as powerful tools and bring significant benefits to recommender systems, making them highly advantageous for personalized recommendations. With exceptional learning and language modeling ability, LLMs make it possible to construct effective, multitask and multimodal recommendation engines. This tutorial aims at sharing knowledge about the development, application and potential benefits of LLMs in recommender systems, which helps to enhance the understanding and adoption of these models among researchers and practitioners. The tutorial also provides an opportunity to address the challenges and considerations specific to LLM-based recommendation systems, such as personalization, data privacy, fairness, and interpretability, fostering a deeper understanding of the implications and responsible use of these models. Finally, the tutorial can also serve as a platform for exchanging ideas, sharing best practices, and encouraging collaborations among experts in the field, ultimately driving advancements and innovation in LLM-based recommendation systems.

The tutorial will introduce LLM-based recommendation from five main perspectives – dataset, model, evaluation, toolkit, and real-world systems. In particular:

- **Datasets:** We introduce datasets that facilitate LLM-based recommendation models. This is particularly important for data-centric machine learning such as LLM-based recommender systems, since the pre-training of LLMs largely determines the ability and utility of LLM-based recommendation.
- **Models:** In this part of the tutorial, we organize and introduce recent LLM-based recommendation models, their relationships, various pre-training, fine-tuning and prompting strategies of LLM-based recommendation models, and possible directions for future improvements.
- **Evaluation:** We introduce evaluation methods for LLM-based recommendation models. Because of the multitask, multimodal, and cross-data nature of LLM-based recommendation models, evaluating the models not only focus on recommendation accuracy, but also many other perspectives such as text quality, efficiency and fluency.
- **Toolkit:** We introduce existing open-source models and platforms to facilitate LLM-based recommendation research, including both LLM backbones such as T5 and LLaMA, and LLM-based recommendation platforms such as OpenP5.

• **Real-world systems:** Finally, we introduce existing industrial LLM systems that support recommender functionality and their advantages and problems to improve. Examples include ChatGPT, Microsoft Bing and Google Bard.

Based on the above rich set of materials, this tutorial will greatly help researchers from both academia and industry who are interested in LLM-based recommendation.

2 TUTORIAL LOGISTICS

Tutorial length: The length of the tutorial is 90 minutes.

Targeted audience: The tutorial will be mainly targeting on recommender system researchers and practitioners from both academia and industry since we introduce LLM for recommendation. The tutorial will also attract researchers who work in broader AI/ML communities such as NLP since we will introduce how large language models can be personalized and tailored for recommendation tasks. Audiences are not required to have preliminary knowledge on NLP and LLMs, since we will introduce the basic NLP and LLM knowledge.

Relevant Tutorials: To the best of our knowledge, this is the first tutorial on large language models for recommendation.

3 BRIEF BIO OF ORGANIZERS

Wenyue Hua is a PhD student in the Department of Computer Science at Rutgers University under the supervision of Prof. Yongfeng Zhang. Her research interest focuses on the intersection of Natural Language Processing and Recommender Systems. Her current research focuses on LLM and its application on recommendation. In the previous she did her BA in Linguistics and BS in Mathematics at UCLA. Her research appears on ACL, EMNLP, ICLR, etc. and she is actively serving as a reviewer for conferences such as RecSys, WWW, SIGIR, ACL, and EMNLP.

Lei Li is a post-doc at the Department of Computer Science, Hong Kong Baptist University (HKBU), and a visiting researcher at the Department of Computer Science, Rutgers University. He has been working on large pre-trained language models for recommender systems. He did PhD at the same department at HKBU and worked on large language models for explainable recommendation. His research appears on SIGIR, WWW, ACL, CIKM, TOIS, etc. and he regularly serves as PC member or reviewer for conferences and journals such as WWW, RecSys, TKDE, TOIS, etc.

Shuyuan Xu is a PhD student in the Department of Computer Science at Rutgers University supervised by Prof. Yongfeng Zhang. His research interest lies in the intersection of Machine Learning and Information Retrieval. His current research focuses on large language models for recommendation. He has published on RecSys, SIGIR, WWW, CIKM, ICTIR, TORS, IJCAI, WSDM, etc. and he is actively serving as reviewer for conferences or journals such as AAAI, SIGIR, WWW, CIKM, RecSys, KDD, ACM TOIS, ACM TORS, IEEE TKDE. He co-organized the Tutorial on Advances in Simulation Technology for Web Applications at WWW 2023.

Li Chen is Associate Head (Research) and Associate Professor at the Department of Computer Science, Hong Kong Baptist University.

Her research focus is on developing trustworthy and responsible data-driven personalization systems, which integrate research in artificial intelligence, recommender systems, user modeling, and user behavior analytics for the application in various domains including social media, e-commerce, online education, and mental health. She has authored and co-authored over 100 publications, most of which appear in high-impact journals (such as IJHCS, TOCHI, UMUAI, TIST, TIS, KNOSYS, Behavior & Information Technology, AI Magazine, and IEEE Intelligent Systems), and key conferences in the areas of data mining (SIGKDD, SDM), artificial intelligence (IJCAI, AAAI), recommender systems (ACM RecSys), user modeling (UMAP), and intelligent user interfaces (CHI, IUI, Interact). She is now an ACM senior member, co-editor-in-chief of ACM Transactions on Recommender Systems (TORS), editorial board member of User Modeling and User-Adapted Interaction Journal (UMUAI), and associate editor of ACM Transactions on Interactive Intelligent Systems (TiiS). She has also been serving in a number of journals and conferences as guest editor, general co-chair, program co-chair, and senior PC member.

Yongfeng Zhang is an Assistant Professor in the Department of Computer Science at Rutgers University. His research interest is in Information Retrieval, Recommender Systems, Machine Learning, Data Mining and Natural Language Processing. In the previous he was a postdoc in the Center for Intelligent Information Retrieval (CIIR) at UMass Amherst, and did his PhD and BE in Computer Science at Tsinghua University, with a BS in Economics at Peking University. He is a Siebel Scholar of the class 2015. He has been consistently working on recommender system research including explainable recommendation, fairness-aware recommendation, conversational recommendation, and large language models for recommendation. His research appears on related conferences such as RecSys, SIGIR, WWW, CIKM, WSDM, KDD, TOIS, TORS, etc. He serves as Associate Editor for ACM Transactions on Information Systems (TOIS) and ACM Transactions on Recommender Systems (TORS) and senior PC member or area chair for RecSys, SIGIR, WWW, KDD, CIKM, AAAI, etc. He has co-organized Tutorial on Explainable Recommendation and Search (WWW, SIGIR, ICTIR), Tutorial on Conversational Recommendation Systems (RecSys, WSDM, IUI), and Tutorial on Fairness of Machine Learning in Recommender Systems (SIGIR, CIKM).

4 RELEVANT PUBLICATIONS BY THE PRESENTERS

- [1] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. Recommendation as Language Processing (RLP): A Unified Pretrain, Personalized Prompt & Predict Paradigm (P5). In Proceedings of the 16th ACM Conference on Recommender Systems, pp. 299-315. 2022.
- [2] Wenyue Hua, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. How to Index Item IDs for Recommendation Foundation Models. arXiv:2305.06569 (2023).
- [3] Shuyuan Xu, Wenyue Hua, and Yongfeng Zhang. OpenP5: Benchmarking Foundation Models for Recommendation. arXiv:2306.11134 (2023).

- [4] Lei Li, Yongfeng Zhang, and Li Chen. Personalized Prompt Learning for Explainable Recommendation. *ACM Transactions on Information Systems* 41, no. 4 (2023): 1-26.
- [5] Wenyue Hua, Yingqiang Ge, Shuyuan Xu, Jianchao Ji, and Yongfeng Zhang. UP5: Unbiased Foundation Model for Fairness-aware Recommendation. *arXiv:2305.12090* (2023).
- [6] Shijie Geng, Juntao Tan, Shuchang Liu, Zuohui Fu, and Yongfeng Zhang. VIP5: Towards Multimodal Foundation Models for Recommendation. *arXiv:2305.14302* (2023).
- [7] Lei Li, Yongfeng Zhang, and Li Chen. Personalized Transformer for Explainable Recommendation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 4947-4957. 2021.
- [8] Shijie Geng, Zuohui Fu, Yingqiang Ge, Lei Li, Gerard de Melo, and Yongfeng Zhang. Improving Personalized Explanation Generation through Visualization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 244-255. 2022.
- [9] Shijie Geng, Zuohui Fu, Juntao Tan, Yingqiang Ge, Gerard De Melo, and Yongfeng Zhang. Path Language Modeling over Knowledge Graphs for Explainable Recommendation. In *Proceedings of the ACM Web Conference 2022*.
- [10] Lei Li, Yongfeng Zhang, and Li Chen. Generate Neural Template Explanations for Recommendation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pp. 755-764. 2020.
- [11] Shijie Geng, Jianbo Yuan, Yu Tian, Yuxiao Chen, Yongfeng Zhang. HiCLIP: Contrastive Language-Image Pretraining with Hierarchy-aware Attention. In *The Eleventh International Conference on Learning Representations*. 2023.
- [12] Guo Lin and Yongfeng Zhang. Sparks of Artificial General Recommender (AGR): Early Experiments with ChatGPT. *arXiv:2305.04518* (2023).
- [13] Lei Li, Li Chen, and Yongfeng Zhang. Towards Controllable Explanation Generation for Recommender Systems via Neural Template. In *Proceedings of the Web Conference 2020*.