

Yuyue Zhao

University of Amsterdam – the Netherlands

☎ +31 630682933 • ✉ yuyuezha00@gmail.com • 🌐 Yuyue • in Yuyue

Research Interests

Generative recommendation, LLM agents for recommendation, Trustworthy recommendation

Education

University of Amsterdam

Ph.D. in Computer Science

Advisor: Prof. dr. Maarten de Rijke

Amsterdam, the Netherlands

Nov. 2023 - Oct. 2025

University of Science and Technology of China

M.E. Computer Science (Combined M.E.–Ph.D. Program)

Advisor: Prof. dr. Xiangnan He

Hefei, China

Sep. 2019 - Jul. 2023

Southwest Jiaotong University

B.E. in Communication Engineering

Advisor: Prof. dr. Xiangnan He

Chengdu, China

Sep. 2015 - Jul. 2019

Working Experience

Kuaishou

Intern, Recommendation Systems – Re-ranking Team

Developed a user social role modeling framework that leverages MLLM to extract video theme tags and audience information from videos, followed by an LLM for inference-based modeling. This approach enhances recommender systems by better capturing user interest migration and diversity needs. The framework constructs user social role logic graphs from the extracted tags and embeds them as augmentation information into recommendation models, addressing challenges such as static interest modeling and lack of diversity.

Beijing, China

Feb. 2025 - May 2025

Teaching Experience

Teaching assistant at University of Amsterdam

Teaching assistant for Recommender System course

Designed and gave lectures and supervised students to reproduce recommendation methods.

Amsterdam, the Netherlands

Jun. 2024 - Jul. 2024

Teaching assistant at University of Science and Technology of China

Teaching assistant for Data Science course

Worked on exercise/exam designing, grading, and answering students' questions.

Hefei, China

Sep. 2021 - Dec. 2021

List of Publications

Citation metrics (as of Nov. 2025): h-index is 6, with 302 citations.

Unseen Threats: Media Bias-Aware Textual Attacks on News Recommender Systems.

Yuyue Zhao, Jin Huang, Maarten de Rijke

ACM Transactions on Information Systems (under review), 2025.

LANCE: Exploration and Reflection for LLM-based Textual Attacks on News Recommendation.

Yuyue Zhao, Jin Huang, Shuchang Liu, Jiancan Wu, Xiang Wang and Maarten de Rijke.

19th ACM Conference on Recommender Systems, **Spotlight Oral**, 2025: 197-206.

Addressing Attribute Bias in Collaborative Filtering: An Invariant Contrastive Learning Approach.

Yuyue Zhao, Xiang Wang, Xiangnan He

ACM Transactions on Information Systems (under review), 2025.

Revisiting Language Models in Neural News Recommender Systems.

Yuyue Zhao, Jin Huang, David Vos, Maarten de Rijke

47th European Conference on Information Retrieval, Cham: Springer Nature Switzerland, 2025: 161-176.

Let Me Do It For You: Towards LLM Empowered Recommendation via Tool Learning.

Yuyue Zhao, Jiancan Wu, Xiang Wang, Wei Tang, Dingxian Wang, Maarten de Rijke

47th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2024: 1796-1806.

Can LLMs Serve as User Simulators for Rec Sys?

Yuyue Zhao, Jin Huang, Maarten de Rijke

The Search Futures Workshop at 46th European Conference on Information Retrieval, 2024.

Time-aware Path Reasoning on Knowledge Graph for Recommendation.

Yuyue Zhao, Xiang Wang, Jiawei Chen, Yashen Wang, Wei Tang, Xiangnan He, Haiyong Xie

ACM Transactions on Information Systems, vol. 41, no. 2, article 26, 2022: 1-26.

EvoWiki: Evaluating LLMs on Evolving Knowledge.

Wei Tang, Yixin Cao, Yang Deng, Jiahao Ying, Bo Wang, Yizhe Yang, Yuyue Zhao, Qi Zhang

63rd Annual Meeting of the Association for Computational Linguistics, vol. 1, 2025: 948-964.

A+B: A General Generator-Reader Framework for Optimizing LLMs to Unleash Synergy Potential.

Wei Tang, Yixin Cao, Jiahao Ying, Bo Wang, Yuyue Zhao, Yong Liao, Pengyuan Zhou

62nd Annual Meeting of the Association for Computational Linguistics, 2024: 3670-3685.

UniRel: Unified Representation and Interaction for Joint Relational Triple Extraction.

Wei Tang, Benfeng Xu, Yuyue Zhao, Zhendong Mao, Yifeng Liu, Yong Liao, Haiyong Xie

Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, 2022: 7087-7099.

Academic Activities

- **Organizer:** Hackathon - 15th European Summer School on Information Retrieval (ESSIR 2024)
- **Journal Reviewer:** TOIS 2021/2022/2023/2024/2025, TKDE 2022/2023/2024/2025, IRM 2023
- **Reviewer:** SIGIR 2025, ACL 2025, KDD 2021/2022, TheWebConf 2024
- **Sub-reviewer:** SIGIR 2024, TheWebConf 2023, NeurIPS 2022

Presentations

- **Conference talks:** RecSys 2025, ECIR 2025, SIGIR 2024, ECIR 2024
- **Invited talk** at Recommender System Course in *University of Amsterdam*: Talk on the research and impact of advanced recommender systems, Jul. 2024
- **Invited talk at Soos:** IRLab Amsterdam, Jan. 2024, Sep. 2024

Projects and Open Source Contributions

A Method for Controlling Media Bias in News Articles using LLMs (BALANCE)

Controllable Media Bias Rewriting for News Recommendation Attacks

Apr. 2025

Developed a novel approach to control media bias in news articles using large language models (LLMs) to influence news recommendation systems. The method rewrites articles to adjust their ideological stance while preserving core content, increasing their recommendation likelihood without detection. This work addresses bias inconsistency in existing textual attack methods, which can cause detection or reduced user engagement. Using an MoE-LoRA finetuning framework, we demonstrated that controlling bias direction while improving news exposure is essential for attacking news recommender systems.

<https://github.com/Go0day/BALANCE>

An LLM-Based News Content Rewriting Framework (LANCE)

Exploration and Reflection for LLM-based Textual Attacks on News Recommender Systems

Feb. 2025

Developed LANCE, a framework to explore vulnerabilities in news recommender systems through textual attacks. LANCE uses a two-component approach: an explorer generates rewritten news with varied styles, sentiments, and personas, and a reflector fine-tunes an open-source LLM to enhance attack effectiveness. LANCE manipulates news rankings by exploiting vulnerabilities in recommender systems, with negative and neutral rewrites proving most effective. It generalizes to unseen systems, exposing shared weaknesses and the need for stronger defenses.

<https://github.com/Go0day/LANCE>

An LM Enhanced News Recommendation Framework (LM4newsRec)

Revisiting Language Models in Neural News Recommender Systems

Oct. 2024

Developed a comprehensive framework for news recommendation that incorporates a variety of language models, ranging from traditional embeddings like GloVe to state-of-the-art transformers such as BERT, RoBERTa, and LLaMA, into established recommendation architectures (NAML, NRMS, LSTUR). This project enables researchers to benchmark and advance the field of news recommendation using the language models.

<https://github.com/Go0day/LM4newsRec>

An Invariant Contrastive Learning Framework for Collaborative Filtering (ICE)

Addressing Attribute Bias in Collaborative Filtering: An Invariant Contrastive Learning Approach June 2024

Developed a novel recommendation paradigm to mitigate attribute bias in collaborative filtering (CF) recommender systems using Invariant Contrastive Learning (ICE). The approach addresses inter- and intra-attribute biases, where certain attributes or items within attributes are over-recommended, leading to skewed performance and reduced diversity. By enhancing recommendation objectives with constraints for inter-attribute distinction and intra-attribute invariance, ICE improves attribute-level robustness. ICE achieves leading performance in matching user preferences (original test) while reducing attribute bias (attribute-balanced test).

UnderReview

An LLM as an agent using tools for recommendation (ToolRec)

Development of this LLM-based recommender

Feb. 2024

Developed ToolRec, a framework leveraging LLMs as surrogate users to enhance recommender systems by capturing fine-grained user preferences. ToolRec integrates LLMs with attribute-oriented tools (ranking and retrieval) to address challenges like hallucinations and semantic-behavior misalignment. By formulating recommendations as an attribute-driven exploration process, the framework invokes external tools to probe item pools and generate personalized, nuanced recommendations.

<https://github.com/Go0day/ToolRec-Code>

A KG-based framework for recommendation in the time-aware scenario (TPRec)

Development of the pipeline of KG-based recommendation

Apr. 2022

Developed TPreC, a novel framework for time-aware path reasoning on knowledge graphs to enhance explainable recommendation systems. By incorporating temporal information, such as purchase and recommendation times, TPreC addresses limitations in existing knowledge graph-based methods that lack temporal context. The framework constructs a temporal collaborative knowledge graph (TCKG) and employs a time-aware path reasoning method to deliver recommendations with contextual explanations. TPreC can achieve performance gains and further improve the quality of explanations.

<https://github.com/Go0day/TPRec>

A Tuning Framework to Align LLM with Recommendation (TALLRec)

Participation in the development of this recommendation framework

Jun. 2023

A novel framework (TALLRec) that enables the efficient and effective adaptation of LLMs for recommendation tasks, even with a limited dataset of fewer than 100 samples.

<https://github.com/SAI990323/TALLRec>

Skills

- **Technologies:** \LaTeX , Python, Pytorch
- **Language:** English, Mandarin (Mother tongue)

References

- **Prof. dr. Maarten de Rijke**
Distinguished University Professor, University of Amsterdam
Email: m.derijke@uva.nl
- **Prof. dr. Xiangnan He**
Professor / Vice Dean, University of Science and Technology of China
Email: xiangnanhe@gmail.com