

Toward Trustworthy Recommender Systems in the Era of Agentic AI: From Relational User Modeling to Generative Personalization

Wenqi Fan*

Department of Computing (COMP)
Department of Management and Marketing (MM)
The Hong Kong Polytechnic University
wenqifan03@gmail.com

Abstract

Recommender systems are evolving from models that predict user-item matching scores into intelligent systems that learn from relational structure, retrieve evidence, reason over user intent, and support personalized actions. This talk discusses a research agenda for trustworthy recommender systems in the era of agentic AI, organized around a transition from relational user modeling to generative personalization. We will discuss three connected components: relational user modeling with deep neural networks, generative and agentic personalization with large language models, and trustworthy foundation for reliable recommendations. The discussion revisits how social, behavioral, and knowledge relations can be integrated for user modeling; how generative models and personalized agents expand recommendation from ranking to intelligent decision support; and how trustworthiness challenges should be addressed throughout the system. The talk aims to provide a unified view of recommender systems as relational, generative, agentic, and trustworthy AI systems.

1 Introduction

In this era of information proliferation, users are surrounded by overwhelming amounts of online content, products, services, and knowledge [Fan *et al.*, 2019c; Ning *et al.*, 2024]. Recommender systems (**RecSys**) have become an effective and indispensable way to alleviate this information overload problem by helping users locate information that matches their interests [Fan *et al.*, 2022b; Wang *et al.*, 2024a]. They play an increasingly important role in many user-oriented online services, such as e-commerce platforms, social media sites, entertainment services, job matching, education, healthcare, and scientific discovery. The goal of recommender systems is to provide each user with a personalized list of items that are likely to be clicked, purchased, watched, read, or otherwise useful [Fan *et al.*, 2021]. In doing so, recommender systems not only enrich users' online experiences, but

also increase the traffic, engagement, and revenue of service providers [Qu *et al.*, 2024].

The basic idea of recommender systems is to make use of user-item interactions and their associated side information to infer user preferences. Collaborative behaviors between users and items provide the main evidence for learning user and item representations. Meanwhile, side information, such as item descriptions, user profiles, user reviews, social relations, and knowledge graphs, provides valuable signals for understanding why a user might prefer an item [Zhao *et al.*, 2024]. On social media platforms (e.g., WeChat, LinkedIn, Facebook, TikTok, etc.), for instance, users obtain and share information through the people in their social circles, including friends, classmates, relatives, and coworkers. These social relations can help profile user preferences and improve personalized recommendation performance [Fan *et al.*, 2020]. More recently, textual and multimodal information has become increasingly important, as large language models enable recommender systems to better understand user intent, retrieve external evidence, generate explanations, support conversations, and reason over complex recommendation scenarios [Zhao *et al.*, 2024].

Despite their success, traditional recommender systems face several intrinsic limitations. First, user behavior is relational and heterogeneous: user-item interactions, social links, item-item relations, and knowledge-graph facts provide different but related views of preference. Second, personalization is moving beyond simple ranking toward agentic decision support, where the system may retrieve evidence, generate explanations, interact with users, and take personalized actions under user intent. Third, as recommender systems are increasingly deployed for data-driven decision making in a range of safety-critical domains (e.g., finance, healthcare, education, etc.), they face significant trustworthiness risks, such as adversarial attacks, unequal exposure, privacy, explanations, and computational vulnerabilities [Fan *et al.*, 2021; Liu *et al.*, 2021]. Therefore, modern recommender systems are no longer only ranking machines; they are becoming agentic systems that model relations, reason with evidence, personalize decisions, and operate under trustworthiness constraints.

This transition raises a central question: how can we build recommender systems that are accurate, expressive, and trustworthy at the same time? High-quality recommendation per-

*Homepage: <https://wenqifan03.github.io>

formance is necessary but insufficient. A system may achieve strong offline ranking metrics while relying on biased exposure, fragile correlations, privacy-sensitive signals, or misleading generated explanations. These issues become more serious when recommendation is powered by deep learning, generative models, large language models (LLMs), and agentic systems. They motivate three research questions: how relational structure can be used to model users, items, and contexts; how generative and agentic models expand recommendation into reasoning, interaction, and personalized predictions; how trustworthiness can be built into recommender systems?

To address these research questions, in this paper, we summarize a research program toward trustworthy recommender systems in the era of agentic AI from relational user modeling to generative personalization. The first component is **relational user modeling**. Users are not independent samples in recommender systems. They are connected by social ties, item interactions, sequential behaviors, and knowledge relations. To model such relational information for personalized recommendations, a powerful approach is to leverage deep learning methods, in particular graph neural networks (GNNs) [Fan *et al.*, 2019b; Fan *et al.*, 2019a; Fan *et al.*, 2020; Wu *et al.*, 2022; Derr *et al.*, 2020; Fan *et al.*, 2022a]. The second component is **generative personalization and agentic recommendation**. Generative models—especially Large Language Models (LLMs), which represent the most advanced AI methods—offer powerful abilities in understanding human natural language, reasoning about complex tasks, and generating content [Liu *et al.*, 2023; Zhao *et al.*, 2024; Wang *et al.*, 2025a; Qu *et al.*, 2025; Huang *et al.*, 2026a; Qu *et al.*, 2026]. The third component is **trustworthy foundations**. Recommender systems face attacks, bias, privacy leakage, unfair exposure, unreliable generated outputs, and computational risks in agentic settings [Liu *et al.*, 2021; Fan *et al.*, 2021; Chen *et al.*, 2022; Zhang *et al.*, 2022; Liu *et al.*, 2020; Ning *et al.*, 2026c; Ning *et al.*, 2026b; He *et al.*, 2026].

2 Relational User Modeling

2.1 Heterogeneous User Behavior

Relational user modeling begins from the observation that user behavior is heterogeneous. In social recommendation, a user is involved in the item domain and the social relation domain. The item domain records interactions between users and items, while the social relation domain records relations between users. These two domains are correlated but not interchangeable. A user may share interests with friends, but social relations can be weak, noisy, asymmetric, or context dependent. A straightforward solution is to incorporate user’s social relations by learning their social embedding knowledge based on the network embedding techniques [Fan *et al.*, 2018] and leveraging the sequential models (e.g., LSTM and RNNs) [Fan *et al.*, 2019c]. Moreover, to capture these heterogeneous user behaviors, GraphRec is proposed to model social recommendation with graph neural networks (GNNs) [Fan *et al.*, 2019b; Fan *et al.*, 2020]. Particularly, GraphRec jointly captured user–item interactions, asso-

ciated opinions, user–user social relations, and heterogeneous relation strengths. This formulation shows that graph neural networks are not simply a replacement of matrix factorization based on deep neural networks. They provide a mechanism for propagating information across users, items, and relations while preserving the structure that generates user preferences. DASO further studied the domain heterogeneity problem [Fan *et al.*, 2019a]. Rather than forcing a unified user representation for both social and item domains, it learned domain-specific representations and transferred information through bidirectional adversarial mapping. The lesson is that relational signals should be shared when they are informative, but separated when they describe different behaviors.

2.2 Selective and Scalable Relation Use

Relational information is useful but not always reliable. Some social links are weak; some interactions are accidental; some temporal signals reflect exposure rather than preference. This motivates models that are not only expressive but also selective. Disentangled contrastive learning separates different factors in social recommendation [Wu *et al.*, 2022]; graph trend filtering studies smoother and more robust graph signals [Fan *et al.*, 2022a]; AutoLoss and AutoEmb investigate how to search loss functions and embedding dimensions for recommendation [Zhao *et al.*, 2021b; Zhao *et al.*, 2021a]. These works share one principle: recommendation models should adapt their learning process to the structure and quality of data. Scalability is another challenge. Real-world recommendation graphs are large, dynamic, and noisy [Zhang *et al.*, 2024]. Graph reduction, sparsification, coarsening, and condensation provide potential tools for making graph learning efficient and reliable [Hashemi *et al.*, 2024; Wang *et al.*, 2023]. For trustworthy recommendation, scalability is not merely a computational concern. It also affects whether a model can be audited, updated, forgotten, and deployed under resource constraints.

3 Generative Personalization and Agentic Recommendation

3.1 From Ranking to Generative Decision Support

Generative personalization is motivated by a limitation of classical recommendation: a ranking score describes preference strength, but it does not directly model the distribution, uncertainty, or rationale behind user decisions. Generative models provide a different view by synthesizing representations and model distribution. This makes generation useful for data augmentation, uncertainty modeling, and interactive personalization.

Diffusion models are a representative family of generative models. Instead of generating an output in one step, they define a forward process that gradually corrupts data with noise and learn a reverse denoising process that recovers high-quality samples from noise [Liu *et al.*, 2023]. This principle has led to strong results in image synthesis, video generation, molecular design, time-series modeling, and graph generation [Liu *et al.*, 2023]. For recommendation, it is useful because user feedback is sparse, preference signals are

noisy, and the target output may be a representation, a sequence of items, an explanation, or a decision trace. Recent work adapts diffusion to recommendation in complementary ways. SGSR addresses noisy social recommendation by treating raw social relations as imperfect signals and using stochastic-differential-equation-based score models to generate social user representations aligned with collaborative behavior, supported by curriculum training and self-supervised alignment [Liu *et al.*, 2025]. CDRec targets the mismatch between continuous diffusion and discrete interactions: it builds a continuous-time discrete diffusion process over historical interactions, uses masking-style transitions and a popularity-aware noise schedule, and accelerates sampling with consistency parameterization and multi-hop contrastive learning [Liu *et al.*, 2026b]. Continuous-token diffusion further formulates recommendation as generation over continuous item tokens in the context of LLM [Qu *et al.*, 2026].

3.2 LLM-Empowered and Agentic Recommendation

LLMs have generated a substantial impact on AI because they are pre-trained on massive text corpora (i.e., world knowledge) and scaled to billions of parameters. This training paradigm gives them strong language understanding and generation abilities, and, more importantly, allows them to exhibit generalization and reasoning capabilities across unseen tasks and domains. Instead of training a separate model for every recommendation scenario, an LLM can adapt through instructions, demonstrations, prompting, and in-context learning. These capabilities are highly relevant to recommender systems. User preferences are not only reflected in clicks and ratings; they are also expressed through reviews, item descriptions, search queries, conversations, and external knowledge. LLMs can therefore connect collaborative signals with semantic evidence: they can encode item content, infer intent from sparse feedback, explain why an item fits a user, resolve ambiguity through dialogue, and generate step-by-step reasoning for complicated decision-making processes. In our survey on recommender systems in the era of LLMs, we organized existing work around pre-training, fine-tuning, prompting, and in-context learning [Zhao *et al.*, 2024]. This view clarifies that LLMs can act as feature encoders, preference reasoners, simulators, and conversational interfaces, making them a promising foundation for next-generation recommender systems.

ID Tokenization for Generative Recommendation. The first story concerns how a recommender can speak the language of an LLM. Traditional recommenders represent users and items as IDs, while LLMs operate over tokens. TokenRec bridges this gap by learning to tokenize user and item IDs with collaborative knowledge: it quantizes masked user/item representations into discrete tokens and uses generative retrieval to avoid expensive auto-regressive decoding for top- K recommendation [Qu *et al.*, 2025]. This tokenization view is further extended by diffusion generative recommendation with continuous tokens, which moves beyond discrete ID prediction and models recommendation as generation over a continuous item-token space [Qu *et al.*, 2026].

Recommendation Assistants with Reasoning. The second

story moves from item generation to interactive assistance. Real users often ask complex queries with implicit constraints, incomplete descriptions, and multi-step intentions. RecBench+ studies this setting by building a benchmark for personalized recommendation assistants and showing that LLMs still struggle with reasoning-intensive or misleading queries [Huang *et al.*, 2026b]. ReRec then improves this capability through reinforcement fine-tuning: it designs dual-graph enhanced rewards, reasoning-aware advantage estimation, and an online curriculum scheduler to train LLM-based assistants for complex recommendation reasoning [Huang *et al.*, 2026a].

Recommendation with Retrieval-Augmented Generation. Generative personalization should not rely only on the parametric memory of LLMs [Fan *et al.*, 2024; Yuan *et al.*, 2025; Jiang *et al.*, 2025], as updating the parameters in LLMs needs huge computational resources and time cost. In order to efficiently incorporate external knowledge into LLM-based recommender systems, Retrieval-Augmented Generation technique is proposed to retrieve informative knowledge to enhance the understanding, reasoning, and generation of LLM [Fan *et al.*, 2024]. In addition, users, items, attributes, and external knowledge form graphs, and these graphs provide the evidence needed for faithful recommendation. A promising solution is graph-enhanced generative personalization, where LLMs reason over social graphs, interaction graphs, item graphs, and knowledge graphs while preserving structural provenance [Wang *et al.*, 2025a; Wang *et al.*, 2025b]. Thus, KG-RAG for LLM-based recommendation studies how retrieved knowledge subgraphs can reduce hallucination and improve evidence-grounded recommendation [Wang *et al.*, 2025a]. MixRAGRec further develops a mixture-of-experts KG-RAG framework for multi-agent LLM-based recommendation, where different retrieval experts handle queries with different graph granularities and agents align structured knowledge with user preference [Wang *et al.*, 2026].

4 Trustworthy Foundations

Trustworthiness is not a single property. In our computational perspective on trustworthy AI, we discussed robustness, fairness, explainability, privacy, accountability, and environmental well-being [Liu *et al.*, 2021]. Recommender systems connect these dimensions because they operate on sensitive user behavior, shape information access, and interact with strategic users and platforms.

One line of work studies vulnerabilities of recommender systems. In the cross-domain setting, CopyAttack shows that black-box recommender systems can be attacked by copying transferable cross-domain user profiles [Fan *et al.*, 2021]. KGAttack studies how knowledge graph-enhanced attacks can manipulate recommendation outcomes from item side [Chen *et al.*, 2022]. In social media, attackers can create fake user behaviors on items and social relations to manipulate the recommendation performance [Wang *et al.*, 2024b]. As recommender systems become agentic and retrieval-augmented, new vulnerabilities appear, including recommendation [Ning *et al.*, 2024], computational cost at-

tacks on WebAgents [Ning *et al.*, 2026c] and inference-cost attacks on retrieval-augmented LLMs [Liu *et al.*, 2026a]. These works indicate that recommendation security involves not only prediction accuracy but also model access, knowledge sources, and resource consumption. Defense should be designed together with model architecture [Ning *et al.*, 2025a]. Trend Filtering technique is employed to enhance the robustness of GNN-based recommendation [Fan *et al.*, 2022a]. Retrieval-augmented purification studies how to filter and ground external evidence for robust LLM-empowered recommendation [Ning *et al.*, 2026b]; graph defense diffusion explores generative defense mechanisms based on diffusion techniques for graph learning [He *et al.*, 2026]. To study the vulnerabilities of LLM-based RecSys under backdoor attacks, BadRec is developed to perturb the items' titles with triggers for backdoor attack and P-Scanner as a universal defense strategy is designed to detect the poisoned items by leveraging the powerful LLMs [Ning *et al.*, 2026a].

Fairness is particularly important in recommendation because the system allocates exposure and attention. A small bias in ranking can be amplified through feedback loops. Fairness reprogramming studies how models can be adapted toward fairer behavior [Zhang *et al.*, 2022]; fairness in dialogue systems examines whether interactive models behave differently across gender groups [Liu *et al.*, 2020]; fairly adaptive negative sampling studies how training signals affect fairness in recommendation [Chen *et al.*, 2023]. Another new vulnerabilities in LLM have been explored on computational cost [Ning *et al.*, 2026c; Liu *et al.*, 2026a].

5 Future Directions

5.1 From Recommenders to Personalized Agents

Agentic AI makes active assistance a natural extension of recommendation. LLM-based WebAgents can perceive web environments, plan multi-step actions, reason over instructions, and execute tasks for users [Ning *et al.*, 2025b]. This development expands the role of recommender systems: a future system may not only suggest a paper, product, route, or service, but also compare evidence, negotiate constraints, complete forms, schedule actions, or monitor outcomes. This shift requires a new theory of user modeling. The model must represent long-term preference, short-term intent, task context, risk tolerance, and delegation boundaries. It must know when to recommend, when to ask, when to act, and when to stop. In addition, as the advanced wearable device, AI smart glasses provide a concrete setting for this direction. A personal recommendation assistant on smart glasses observes the user's egocentric visual context, understands spoken or implicit intent, retrieves external knowledge, and recommends actions in real time. For example, it may identify a product, landmark, food item, sign, or object in front of the user, retrieve relevant information, and suggest what to buy, visit, avoid, or do next. SUPERGLASSES studies this emerging setting by benchmarking vision-language models as intelligent agents for AI smart glasses and showing the need for task-specific perception, retrieval, and reasoning pipelines [Jiang *et al.*, 2026b]. For recommendation, smart glasses make personalization embodied: the system must reason about what the

user sees, where the user is, what the user wants, and what action is appropriate now.

5.2 Trustworthy Recommendation Infrastructure

Trustworthiness should become infrastructure, not a post-hoc checklist. This is especially important for agentic and retrieval-augmented recommendation, where errors may arise from biased data, malicious manipulation, noisy relations, weak retrieval, unsupported generation, or excessive resource consumption. Efficiency is a central part of this infrastructure. LLM-based recommenders must serve large user traffic under strict latency constraints, so direct online LLM invocation is often impractical. For example, AIR investigates this problem in industrial cross-domain recommendation by moving expensive LLM reasoning offline and composing compact intent representations online, achieving efficient deployment while preserving semantic consistency [Jiang *et al.*, 2026a]. However, efficiency also creates new security risks. Retrieval-augmented LLM systems can be attacked through poisoned external knowledge that increases token consumption during inference [Liu *et al.*, 2026a]. WebAgents can also be attacked through adversarial webpage prompts that induce unnecessarily long reasoning traces and excessive computational cost [Ning *et al.*, 2026c]. These studies show that computational cost is not only an engineering metric; it is an attack surface. Future trustworthy recommender systems therefore need cost-aware retrieval, budgeted reasoning, attack detection, and graceful degradation when generated reasoning becomes unnecessarily long or resource intensive.

6 Conclusion

This paper discusses trustworthy recommender systems in the era of agentic AI, focusing on the transition from relational user modeling to generative personalization. Relational user modeling provides the foundation for learning user preferences from heterogeneous evidence, including interactions, social links, item attributes, and knowledge graphs. Generative and agentic models further expand recommendation from ranking to explanation, reasoning, conversation, tool use, and multi-step personal assistance. These new capabilities also introduce new risks. Trustworthiness should therefore be built into recommender systems from the beginning, rather than added after deployment. Future systems need to be robust to manipulation, fair in exposure and interaction, privacy-aware in user modeling, faithful in generated outputs, and efficient for real-world use. This calls for methods that connect graph learning, LLM-based reasoning, personalized agents, and cost-aware infrastructure. The long-term goal is to develop recommender systems that not only predict what users may like, but also know when to ask, act under user control, and remain reliable in open environments.

References

[Chen *et al.*, 2022] Jingfan Chen, Wenqi Fan, Guanghui Zhu, Xiangyu Zhao, Chunfeng Yuan, Qing Li, and Yihua Huang. Knowledge-enhanced black-box attacks for recommendations. In *The ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022.

- [Chen *et al.*, 2023] Xiao Chen, Wenqi Fan, Jingfan Chen, Haochen Liu, Zitao Liu, Zhaoxiang Zhang, and Qing Li. Fairly adaptive negative sampling for recommendations. In *Proceedings of the ACM Web Conference 2023*, 2023.
- [Derr *et al.*, 2020] Tyler Derr, Yao Ma, Wenqi Fan, Xiaorui Liu, Charu Aggarwal, and Jiliang Tang. Epidemic graph convolutional network. In *Proceedings of the 13th International Conference on Web Search and Data Mining (WSDM)*, 2020.
- [Fan *et al.*, 2018] Wenqi Fan, Qing Li, and Min Cheng. Deep modeling of social relations for recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- [Fan *et al.*, 2019a] Wenqi Fan, Tyler Derr, Yao Ma, Jianping Wang, Jiliang Tang, and Qing Li. Deep adversarial social recommendation. In *28th International Joint Conference on Artificial Intelligence (IJCAI-19)*. International Joint Conferences on Artificial Intelligence, 2019.
- [Fan *et al.*, 2019b] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. Graph neural networks for social recommendation. In *The world wide web conference*, 2019.
- [Fan *et al.*, 2019c] Wenqi Fan, Yao Ma, Dawei Yin, Jianping Wang, Jiliang Tang, and Qing Li. Deep social collaborative filtering. In *Proceedings of the 13th ACM Conference on Recommender Systems*, 2019.
- [Fan *et al.*, 2020] Wenqi Fan, Yao Ma, Qing Li, Jianping Wang, Guoyong Cai, Jiliang Tang, and Dawei Yin. A graph neural network framework for social recommendations. *IEEE Transactions on Knowledge and Data Engineering*, 2020.
- [Fan *et al.*, 2021] Wenqi Fan, Tyler Derr, Xiangyu Zhao, Yao Ma, Hui Liu, Jianping Wang, Jiliang Tang, and Qing Li. Attacking black-box recommendations via copying cross-domain user profiles. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 2021.
- [Fan *et al.*, 2022a] Wenqi Fan, Xiaorui Liu, Wei Jin, Xiangyu Zhao, Jiliang Tang, and Qing Li. Graph trend filtering networks for recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022.
- [Fan *et al.*, 2022b] Wenqi Fan, Xiangyu Zhao, Xiao Chen, Jingran Su, Jingtong Gao, Lin Wang, Qidong Liu, Yiqi Wang, Han Xu, Lei Chen, et al. A comprehensive survey on trustworthy recommender systems. *arXiv preprint arXiv:2209.10117*, 2022.
- [Fan *et al.*, 2024] Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. A survey on rag meeting llms: Towards retrieval-augmented large language models. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024.
- [Hashemi *et al.*, 2024] Mohammad Hashemi, Shengbo Gong, Juntong Ni, Wenqi Fan, B. Aditya Prakash, and Wei Jin. A comprehensive survey on graph reduction: Sparsification, coarsening, and condensation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2024.
- [He *et al.*, 2026] Xin He, Wenqi Fan, Yili Wang, Chengyi Liu, Rui Miao, Xin Juan, and Xin Wang. Graph defense diffusion model. In *The ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2026.
- [Huang *et al.*, 2026a] Jiani Huang, Shijie Wang, Liangbo Ning, Wenqi Fan, and Qing Li. Rerac: Reasoning-augmented llm-based recommendation assistant via reinforcement fine-tuning. In *The Annual Meeting of the Association for Computational Linguistics*, 2026.
- [Huang *et al.*, 2026b] Jiani Huang, Shijie Wang, Liangbo Ning, Wenqi Fan, Shuaiqiang Wang, Dawei Yin, and Qing Li. Towards next-generation recommender systems: A benchmark for personalized recommendation assistant with llms. In *Proceedings of the ACM International Conference on Web Search and Data Mining*, 2026.
- [Jiang *et al.*, 2025] Zhuohang Jiang, Pangjing Wu, Xu Yuan, Wenqi Fan, and Qing Li. Qa-dragon: Query-aware dynamic rag system for knowledge-intensive visual question answering. *arXiv preprint arXiv:2508.05197*, 2025.
- [Jiang *et al.*, 2026a] Zhuohang Jiang, Yuxin Chen, Shijie Wang, Haoqiao Qu, Jindong Zhou, Wenqi Fan, Qing Li, Dongxu Liang, and Jun Wang. Atomic intent reasoning: Bringing llm semantics to industrial cross-domain recommendations. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2026.
- [Jiang *et al.*, 2026b] Zhuohang Jiang, Xu Yuan, Haoqiao Qu, Shanru Lin, Kanglong Liu, Wenqi Fan, and Li Qing. Superglasses: Benchmarking vision language models as intelligent agents for ai smart glasses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026.
- [Liu *et al.*, 2020] Haochen Liu, Jamell Dacon, Wenqi Fan, Hui Liu, Zitao Liu, and Jiliang Tang. Does gender matter? towards fairness in dialogue systems. In *Proceedings of the 28th International Conference on Computational Linguistics*, 2020.
- [Liu *et al.*, 2021] Haochen Liu, Yiqi Wang, Wenqi Fan, Xiaorui Liu, Yaxin Li, Shaili Jain, Yunhao Liu, Anil K Jain, and Jiliang Tang. Trustworthy ai: A computational perspective. *arXiv preprint arXiv:2107.06641*, 2021.
- [Liu *et al.*, 2023] Chengyi Liu, Wenqi Fan, Yunqing Liu, Jiatong Li, Hang Li, Hui Liu, Jiliang Tang, and Qing Li. Generative diffusion models on graphs: Methods and applications. In *32nd International Joint Conference on Artificial Intelligence*, 2023.
- [Liu *et al.*, 2025] Chengyi Liu, Jiahao Zhang, Shijie Wang, Wenqi Fan, and Qing Li. Score-based generative diffusion models for social recommendations. *IEEE Transactions on Knowledge and Data Engineering*, 2025.
- [Liu *et al.*, 2026a] Chengliang Liu, Liangbo Ning, Yujuan Ding, and Wenqi Fan. Inference cost attacks for retrieval-augmented large language models. In *Proceedings of The Web Conference*, 2026.

- [Liu et al., 2026b] Chengyi Liu, Xiao Chen, Shijie Wang, Wenqi Fan, and Qing Li. Continuous-time discrete-space diffusion model for recommendation. In *Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining*, 2026.
- [Ning et al., 2024] Liang-bo Ning, Shijie Wang, Wenqi Fan, Qing Li, Xin Xu, Hao Chen, and Feiran Huang. Cheatagent: Attacking llm-empowered recommender systems via llm agent. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024.
- [Ning et al., 2025a] Liangbo Ning, Wenqi Fan, and Qing Li. Retrieval-augmented purifier for robust llm-empowered recommendation. *ACM Transactions on Information Systems*, 2025.
- [Ning et al., 2025b] Liangbo Ning, Ziran Liang, Zhuohang Jiang, Haohao Qu, Yujuan Ding, Wenqi Fan, Xiao-yong Wei, Shanru Lin, Hui Liu, Philip S Yu, et al. A survey of webagents: Towards next-generation ai agents for web automation with large foundation models. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, 2025.
- [Ning et al., 2026a] Liangbo Ning, Wenqi Fan, and Qing Li. Exploring backdoor attack and defense for llm-empowered recommendations. *IEEE Transactions on Knowledge and Data Engineering*, 2026.
- [Ning et al., 2026b] Liangbo Ning, Wenqi Fan, and Qing Li. Retrieval-augmented purifier for robust llm-empowered recommendation. *ACM Transactions on Information Systems*, 2026.
- [Ning et al., 2026c] Liangbo Ning, Yuchen Zhu, Heqing Huang, Xin Wang, Yi Chang, Qing Li, and Wenqi Fan. When efficiency becomes a vulnerability: Computational cost attacks on webagents. In *The Annual Meeting of the Association for Computational Linguistics*, 2026.
- [Qu et al., 2024] Haohao Qu, Yifeng Zhang, Liangbo Ning, Wenqi Fan, and Qing Li. Ssd4rec: A structured state space duality model for efficient sequential recommendation. *arXiv preprint arXiv:2409.01192*, 2024.
- [Qu et al., 2025] Haohao Qu, Wenqi Fan, Zihuai Zhao, and Qing Li. Tokenrec: Learning to tokenize id for llm-based generative recommendations. *IEEE Transactions on Knowledge and Data Engineering*, 2025.
- [Qu et al., 2026] Haohao Qu, Shanru Lin, Yujuan Ding, Yiqi Wang, and Wenqi Fan. Diffusion generative recommendation with continuous tokens. In *Proceedings of The Web Conference*, 2026.
- [Wang et al., 2023] Lin Wang, Wenqi Fan, Jiatong Li, Yao Ma, and Qing Li. Fast graph condensation with structure-based neural tangent kernel. *arXiv preprint arXiv:2310.11046*, 2023.
- [Wang et al., 2024a] Hanbing Wang, Xiaorui Liu, Wenqi Fan, Xiangyu Zhao, Venkataramana Kini, Devendra Yadav, Fei Wang, Zhen Wen, Jiliang Tang, and Hui Liu. Rethinking large language model architectures for sequential recommendations. *arXiv preprint arXiv:2402.09543*, 2024.
- [Wang et al., 2024b] Shijie Wang, Wenqi Fan, Xiao-Yong Wei, Xiaowei Mei, Shanru Lin, and Qing Li. Multi-agent attacks for black-box social recommendations. *ACM Transactions on Information Systems*, 2024.
- [Wang et al., 2025a] Shijie Wang, Wenqi Fan, Yue Feng, Lin Shanru, Xinyu Ma, Shuaiqiang Wang, and Dawei Yin. Knowledge graph retrieval-augmented generation for llm-based recommendation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2025.
- [Wang et al., 2025b] Shijie Wang, Jiani Huang, Zhikai Chen, Yu Song, Wenzhuo Tang, Haitao Mao, Wenqi Fan, Hui Liu, Xiaorui Liu, Dawei Yin, et al. Graph machine learning in the era of large language models (llms). *ACM Transactions on Intelligent Systems and Technology*, 2025.
- [Wang et al., 2026] Shijie Wang, Chengyi Liu, Yujuan Ding, Shanru Lin, See-Kiong Ng, Xu Xin, and Wenqi Fan. Mixture-of-experts knowledge graph retrieval-augmented generation for multi-agent llm-based recommendation. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2026.
- [Wu et al., 2022] Jiahao Wu, Wenqi Fan, Jingfan Chen, Shengcai Liu, Qing Li, and Ke Tang. Disentangled contrastive learning for social recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2022.
- [Yuan et al., 2025] Xu Yuan, Liangbo Ning, Wenqi Fan, and Qing Li. mkg-rag: Multimodal knowledge graph-enhanced rag for visual question answering. *arXiv preprint arXiv:2508.05318*, 2025.
- [Zhang et al., 2022] Guanhua Zhang, Yihua Zhang, Yang Zhang, Wenqi Fan, Qing Li, Sijia Liu, and Shiyu Chang. Fairness reprogramming. In *Thirty-sixth Conference on Neural Information Processing Systems*, 2022.
- [Zhang et al., 2024] Jiahao Zhang, Rui Xue, Wenqi Fan, Xin Xu, Qing Li, Jian Pei, and Xiaorui Liu. Linear-time graph neural networks for scalable recommendations. In *Proceedings of the ACM on Web Conference 2024*, 2024.
- [Zhao et al., 2021a] Xiangyu Zhao, Haochen Liu, Wenqi Fan, Hui Liu, Jiliang Tang, and Chong Wang. Autoloss: Automated loss function search in recommendations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021.
- [Zhao et al., 2021b] Xiangyu Zhao, Haochen Liu, Wenqi Fan, Hui Liu, Jiliang Tang, Chong Wang, Ming Chen, Xudong Zheng, Xiaobing Liu, and Xiwang Yang. Autoemb: Automated embedding dimensionality search in streaming recommendations. In *IEEE International Conference on Data Mining (ICDM)*. IEEE, 2021.
- [Zhao et al., 2024] Zihuai Zhao, Wenqi Fan, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Zhen Wen, Fei Wang, Xiangyu Zhao, Jiliang Tang, et al. Recommender systems in the era of large language models (llms). *IEEE Transactions on Knowledge and Data Engineering*, 2024.