ACM Journals

CALL FOR PAPERS

ACM Transactions on Information Systems

Special Section on Causality Representation Learning in LLMs-Driven Recommender Systems

Guest Editors:

- **Lina Yao**, CSIRO's Data61 and University of New South Wales, Australia (lina.yao@unsw.edu.au)
- **Julian McAuley**, University of California San Diego, United States (jmcauley@eng.ucsd.edu)
- **Yongfeng Zhang**, Rutgers University, United States (yongfeng.zhang@rutgers.edu)
- **Kun Zhang**, Carnegie Mellon University, United States, and Mohamed bin Zayed University of Artificial Intelligence, United Arab Emirates (kunz1@cmu.edu)

The realm of recommender systems has witnessed rapid advancements with the application of deep learning and other advanced algorithms. While these models excel at predicting user preferences, understanding the underlying causal relationships has become paramount. The exploration of causal inference in recommender systems aims to offer transparent, reliable, and more interpretable recommendations.

As the digital realm evolves, recommender systems stand at the crossroads of enhancing user experiences by providing tailored content. The advent and prominence of Large Language Models (LLMs) open new avenues to probe deeper into the intricate webs of causality within these systems. LLMs, with their vast knowledge bases and intricate architectures, present an untapped resource for causal discovery. This call for papers seeks pioneering research and reviews that explore, analyze, and showcase the synergy between LLMs and causal inference in recommender systems.

Topics of Interest Include, But Are Not Limited To:

● Foundations of causal inference in recommendation systems.
● Methods for causal effect estimation in user-item interactions.
● Interpretable models based on causal relationships.
● Evaluating the impact of interventions in recommender systems.
● Limitations and challenges in uncovering causalities.
● Trustworthy recommender systems based on causal methods.
● Causal methods for bias and fairness in recommender systems.
● Causal robustness of recommender systems
● Real-world applications and case studies highlighting causal inference benefits.
● Comparison of traditional predictive modelling and causal-based recommendation.
● Theoretical exploration of LLMs in discerning causal relationships within recommender systems.
● Design and development of recommender systems powered by LLM-driven causal insights.
● Leveraging LLMs for enhanced transparency and interpretability in recommendations.
● Addressing inherent biases in LLMs and their impact on causal discovery.
● Case studies demonstrating the practical benefits and challenges of using LLMs for causality in recommendations.
● Comparative evaluation of LLM-integrated causal discovery against conventional methodologies.
● Novel techniques for training LLMs with a focus on causality.

The proposal for a special section on " Causality Representation Learning in LLMs-Driven Recommender Systems " within the ACM Transactions on Information Systems is underpinned by a confluence of academic rigor, timely relevance, and evident momentum within the research community.

● Rapid advancements in machine learning, deep learning, and causality theory have opened new research avenues that require a specialized forum for dissemination and critique has propelled the field of recommender systems forward. The complexity and sophistication of modern recommender systems necessitate a nuanced understanding of the underlying causal mechanisms that inform recommendations, to enhance their transparency, reliability, and interpretability. This special section aims to serve as such a platform, encouraging comprehensive exploration and innovation at the intersection of causality and recommender systems.

● The advent of Large Language Models (LLMs) has significantly broadened the scope for investigating causality within the realm of recommender systems. The expansive knowledge bases and intricate architectures inherent to LLMs present unprecedented opportunities for causal discovery and analysis, positioning them as a focal point of contemporary research endeavors.

● The proposed special section is characterized by its inclusivity and breadth, encompassing a wide array of topics from foundational causal estimation and discovery to the application of LLMs and deep learning techniques in real-world scenarios. This diversity not only reflects the multifaceted nature of the field but also underscores the rich potential for scholarly inquiry across various dimensions of causality in recommender systems.

● This call for papers articulates a balanced approach, soliciting contributions that span practical applications and empirical case studies to theoretical frameworks and methodological explorations. This dual focus is designed to appeal to a broad spectrum of researchers, encompassing both those engaged in applied research and those pursuing more theoretical questions within the context of causality in recommender systems.

● This special section is its emphasis on leveraging causal methods to address critical issues of bias and fairness within recommender systems. This focus is especially pertinent given the escalating concerns surrounding ethical considerations in AI and machine learning. By prioritizing research that confronts these challenges, the special section aims to attract submissions that contribute to the development of more equitable and responsible recommender systems.

The guest editors, affiliated with renowned institutions, bring a wealth of expertise and a broad professional network to this initiative, further bolstering the likelihood of attracting high-quality submissions. Their previous organization of related events, such as the Special Issue on Responsible Recommender Systems in the ACM Transactions on Intelligent Systems and Technology, tutorial on Large Language Models for Recommendation at RecSys 2023, and workshop on Personalized Generative AI at CIKM 2023, evidences the topic's relevance and the community's engagement. These preliminary events have not only demonstrated the community's interest but have also highlighted the depth and quality of research being conducted in this area.

In light of these considerations, the proposed special section on "Causality Representation Learning in LLMs-Driven Recommender Systems" represents a timely and necessary addition to the scholarly literature. It promises to consolidate the research community by providing a forum for the exchange of ideas, fostering collaboration, and establishing a cohesive body of knowledge that advances our understanding of causality in recommender systems. This proposal stands as a testament to the field's dynamism and the ongoing need for rigorous academic discourse to navigate its complexities and challenges.

Important Dates:
Submission Deadline: September 30, 2024
First Round Notification: November 30, 2024
First Round Revision: February 28, 2025
Notification of Final Decision: April 30, 2025
Tentative Publication: Mid-to-late 2025

Guest Editors

Lina Yao (https://www.linayao.com/), CSIRO's Data61 and University of New South Wales, Australia

Lina Yao holds the position of Senior Principal Research Scientist and Science Lead at CSIRO's Data61, alongside her academic roles as a Conjoint Professor at the University of New South Wales, Honorary Professor at Macquarie University, and Adjunct Professor at the University of Technology Sydney. Her research endeavors are deeply rooted in the development of generalizable, transparent, and data-efficient methodologies within the domains of data mining, machine learning, and deep learning. One of her specialized focuses is deep learning-based recommender systems, where she strives to address challenges of data scarcity, interoperability, and complexity of human patterns. Her research delves into understanding the temporal and contextual nuances of human interactions with digital platforms, enabling the development of recommendation models that perform effectively and evolve in tandem with user needs and societal trends.

Julian McAuley (https://cseweb.ucsd.edu/~jmcauley/), University of California, San Diego, USA

Julian McAuley has been a professor at UCSD since 2014 (Associate in 2019; Professor in 2021), where his lab works on problems in the area of Personalized Machine Learning. Broadly speaking, his lab's research seeks to develop machine learning techniques for settings where differences among individuals explain significant variability in outcomes. A core instance of this problem is that of recommender systems, one of the core areas of his lab's research, where he develop technologies that underlie algorithms like those used for recommendations on Netflix, Amazon, or Facebook.

Yongfeng Zhang (https://www.yongfeng.me/), Rutgers University, USA

Yongfeng Zhang is an Assistant Professor in the Department of Computer Science at Rutgers University. His research interest is in Machine Learning and Data Mining, Information Retrieval and Recommender Systems, Natural Language Processing, and Trustworthy AI. His research works appear in top-tier computer science conferences and journals such as SIGIR, WWW, KDD, NeurIPS, ACL, TOIS, TORS, etc. He has been frequently serving as area chair or senior program committee member at SIGIR, KDD, WWW, RecSys, CIKM, AAAI, etc. He also serves as Associate Editor for ACM Transactions on Information Systems (TOIS), ACM Transactions on Recommender Systems (TORS), and Frontiers in Big Data. He is a Siebel Scholar of the class 2015 and an NSF career awardee in 2021.

Kun Zhang, Carnegie Mellon University & Mohamed bin Zayed University of Artificial Intelligence, USA and UAE

Kun Zhang (http://www.andrew.cmu.edu/user/kunz1/) is currently on leave from Carnegie Mellon University (CMU), where he is an associate professor of philosophy and an affiliate faculty in the machine learning department; he is working as a professor of machine learning, the acting chair of the machine learning department, and the director of the Center for Integrative AI at Mohamed bin Zayed University of Artificial Intelligence (MBZUAI). He has been actively developing methods for automated causal discovery from various kinds of data and investigating machine learning problems including transfer learning, concept learning, and deep learning from a causal perspective. Dr. Zhang co-authored a best student paper for UAI and a best finalist paper for CVPR, and received an ICML 2022 Test of Time Award Honorable Mention and the best benchmark award of the causality challenge. He has been frequently serving as a senior area chair, area chair, or senior program committee member for flagship conferences in machine learning or artificial intelligence, including NeurIPS, ICML, UAI, IJCAI, AISTATS, and ICLR. He was a co-founder and general and program co-chair of the first Conference on Causal Learning and Reasoning (CLeaR 2022), a program co-chair of the 38th Conference on Uncertainty in Artificial Intelligence (UAI 2022), and a general co-chair of UAI 2023.