



### ACM Transactions on Information Systems

*Special Issue on Pre-trained Models for Search and Recommendation*

#### Guest Editors:

- **Wenjie Wang**, National University of Singapore, Singapore, [wangwenjie@u.nus.edu](mailto:wangwenjie@u.nus.edu) (contact person for recommendation track)
- **Zheng Liu**, Beijing Academy of Artificial Intelligence, China, [zhengliu1026@gmail.com](mailto:zhengliu1026@gmail.com) (contact person for search track)
- **Fuli Feng**, University of Science and Technology of China, China, [fulifeng93@gmail.com](mailto:fulifeng93@gmail.com)
- **Zhicheng Dou**, Renmin University, China, [dou@ruc.edu.cn](mailto:dou@ruc.edu.cn)
- **Qingyao Ai**, Tsinghua University, China, [aiqy@tsinghua.edu.cn](mailto:aiqy@tsinghua.edu.cn)
- **Grace Hui Yang**, Georgetown University, USA, [huiyang@cs.georgetown.edu](mailto:huiyang@cs.georgetown.edu)
- **Defu Lian**, University of Science and Technology of China, China, [liandefu@ustc.edu.cn](mailto:liandefu@ustc.edu.cn)
- **Lu Hou**, Huawei, China, [lhouab@connect.ust.hk](mailto:lhouab@connect.ust.hk)
- **Aixin Sun**, Nanyang Technological University, Singapore, [axsun@ntu.edu.sg](mailto:axsun@ntu.edu.sg)
- **Hamed Zamani**, University of Massachusetts Amherst, USA, [zamani@cs.umass.edu](mailto:zamani@cs.umass.edu)
- **Donald Metzler**, Google Research, USA, [metzler@google.com](mailto:metzler@google.com)
- **Maarten de Rijke**, University of Amsterdam, The Netherlands, [m.derijke@uva.nl](mailto:m.derijke@uva.nl)

The boom of pre-trained models has paved the way for significant advances in information retrieval domains. Especially, pre-trained models for search and recommendation have received wide attention from the research and industry communities. The powerful capabilities of pre-trained models, such as large language models, offer novel and generalizable solutions for many search and recommendation tasks. This special issue will serve as an invaluable platform for researchers to contribute the latest ideas, advances, and breakthroughs in this rapidly evolving field, delving into influential aspects such as neural ranking, universal retrievers, user modeling, content generation, trustworthiness, evaluation, and industry practices.

The primary aim of this special issue is to foster innovative research centered around the integration of pre-trained models with search and recommendation tasks. In the scope of the search track, this special issue focuses on four main issues. First of all, it is observed that generic pre-trained models are not well suited for the purpose of information retrieval. Therefore, new models and pre-training algorithms are expected to be proposed, which may better support tasks like dense retrieval and neural ranking. Secondly, it is still challenging to learn high-quality neural retrievers in many situations, such as with limited training data or shifting domains. As a result, we look forward to innovative utilization of pre-trained models, like advanced fine-tuning and domain adaptation algorithms, that handle the demanding conditions of real-world scenarios. Thirdly, information retrieval becomes particularly important considering its unique role in extending large language models with external knowledge. Thus, it's necessary to explore new approaches so as to facilitate the corresponding collaborations. Last but not least, given the wide application of pre-trained models in different IR problems, it is the right time to make a systematic analysis of their impact. To this end, we encourage the sharing of empirical studies, evaluation resources, and practical experiences, which will help derive a deeper understanding of current techniques.

In the realm of recommendation, we highlight five key aspects. First, we encourage leveraging pre-trained models, especially generative models such as large language models and diffusion models, to improve recommender algorithms for better user modeling. Second, we foster exploring the possibility of using generative models to produce more diverse content, to supplement human-generated content for meeting users' wide-ranging

information needs. Third, we welcome major innovations in enhancing various recommender tasks by using generative models, e.g., new user-system interaction paradigms based on large language models. Fourth, the special issue will emphasize the importance of trustworthiness when using generative models for recommendations. Last but not least, the special issue also encourages researchers to design diverse evaluation methodologies to examine the usage of pre-trained models in recommender systems.

## Topics

We welcome original submissions on harnessing pre-trained models for search and recommendation, including but not limited to the following topics:

### Search Track

- Pre-training for dense retrieval, sparse retrieval, and neural ranking
- Universal retrievers for cross-domain, cross-lingual, and cross-modality scenarios
- Retrieval augmentations for large language models (with applications like question answering and conversations)
- Generative models for information retrieval (e.g., auto-regressive search engines, query generation)
- Evaluations, benchmarking, empirical studies, and industry practice about pre-trained methods for search engines

### Recommendation Track

- Integrating large language models and other pre-trained generative models to enhance recommender algorithms for user modeling (with applications in finance, streaming platforms, social networks, etc.)
- Generative recommendation by using generative AI to help create personalized item content in some scenarios, such as advertisements, images, and micro-videos.
- Improving traditional recommender tasks by using generative models, e.g., new user-system interaction paradigms, and combining content generation and retrieval for personalized recommendations.
- Trustworthy recommendation with generative models, e.g., developing standards and technologies to inspect the recommended content from the aspects of bias, fairness, privacy, safety, authenticity, and identifiability.
- Evaluations of recommender systems with pre-trained models, especially generative models (e.g., new metrics, standards, and human evaluation paradigms/interfaces)

## Important Dates

- Open for Submissions: Aug 1, 2023
- Submissions deadline: Feb 1, 2024
- First-round review decisions: May 1, 2024
- Deadline for minor revision submissions: Jun 1, 2024
- Deadline for major revision submissions: Aug 1, 2024
- Notification of final decisions: Oct 1, 2024
- Final manuscripts due: Nov 1, 2024
- Tentative publication: Dec 2024 or early 2025

## Submission Information

Authors can submit their manuscripts via <https://dl.acm.org/journal/tois>. Submissions to this special issue will follow the regular TOIS submission guidelines (<https://dl.acm.org/journal/tois/author-guidelines>). Submissions must be accompanied by a cover letter containing all of the following: (1) Confirm that the paper is not currently under submission at another journal or conference. (2) Confirm that the paper is substantially different from any previously published work. (3) Confirm that none of the co-authors is a Guest Editor for this special session. (4) Disclose possible conflicts of interest with Guest Editors. The review process will be single-blind. Strict policies will be followed for plagiarism, submission confidentiality, reviewer anonymity, prior and concurrent paper submission based on the guidelines.

Papers with a “Major Revision” decision should be resubmitted within 3 months, and with a “Minor Revision” decision should be resubmitted within 1 month. Revised submissions must be accompanied with a detailed response to reviewers explaining what revisions were implemented. The editors will conduct a second-round review process and give the decision (accept or reject or need further revision) in one month.

For questions and further information, please contact Wenjie Wang at [wangwenjie@u.nus.edu](mailto:wangwenjie@u.nus.edu) (recommendation track) or Zheng Liu at [zhengliu1026@gmail.com](mailto:zhengliu1026@gmail.com) (search track).