



### ACM Transactions on Recommender Systems

*Special Issue on Trustworthy Recommender Systems*

#### Guest Editors:

- **Markus Schedl**, Johannes Kepler University and Linz Institute of Technology, Austria, [markus.schedl@jku.at](mailto:markus.schedl@jku.at)
- **Yashar Deldjoo**, Polytechnic University of Bari, Italy, [yashar.deldjoo@poliba.it](mailto:yashar.deldjoo@poliba.it)
- **Pablo Castells**, Amazon and Autónoma University of Madrid, Spain, [pablo.castells@uam.es](mailto:pablo.castells@uam.es)
- **Emine Yilmaz**, Amazon and University College London, UK, [emine.yilmaz@ucl.ac.uk](mailto:emine.yilmaz@ucl.ac.uk)

The use of recommender systems in our daily lives is a notable example of information-seeking systems in which machines and humans interact to make joint decisions. One important purpose of such systems is to ensure the producers' and providers' long-term economic growth and the consumers' utility, fairness, and privacy alike. These aspects contribute to the research direction of *trustworthy recommender systems*.

Trustworthiness entails various aspects, which have recently also been identified and categorized by the European Commission and institutions in other countries, and turned into policies and regulations, including the Artificial Intelligence Act, the Digital Service Act, and the Digital Market Act. Those regulations demonstrate that recommender systems and their potential damages are also taken seriously from a political and legal perspective.

The primary concerns that trustworthiness summons in a recommender system context includes aspects of *robustness*, *privacy-awareness*, *interpretability*, and *fairness*, with the definitions becoming more and more blurred as we go from the first to the last. Robustness means the model can withstand noisy and adversarial data. Privacy-awareness refers to whether we can trust models to have access to sensitive data, or whether (some part of the) sensitive data could be released during training. Interpretability refers to the ability to understand how the model makes decisions. Fairness means that the model does not systematically discriminate against groups of consumers or producers.

Trustworthy recommender systems are an essential topic of research from a technical, ethical, moral, legal, and utilitarian perspective, and are embedded into the larger area of responsible AI technology. They should, therefore, be researched from a multidisciplinary viewpoint. For this reason, we highly encourage authors of submissions to take such a multidisciplinary perspective, e.g., by establishing interdisciplinary author teams.

#### Topics:

Topics of interest fall into four categories. They include, but are not limited to, the following:

- *Robustness and security*: adversarial machine learning techniques for attacks on/defense of recommender systems, shilling attacks, theoretical and data analytical studies of robustness of recommender systems and the underlying data;
- *Privacy*: differential privacy, federated machine learning and other multi-party computation methods tailored to recommenders, encryption techniques for privacy-awareness in recommender systems;
- *Interpretability*: feature-level explanation, social explanation, discourse and actionable explanation using counterfactual explanation and causal learning, domain-specific explanation approaches, multi-modal explanations;

- *Fairness and biases*: defining fairness, novel bias metrics, critical investigations of existing metrics, multi-sided fairness, algorithms to mitigate or adjust biases and unfairness.

In addition, we solicit works that directly target the trust dimension in recommender systems, in particular, the human perception of the trustworthiness of recommender systems technology. Some important research questions such submissions could strive to answer include: What makes people trust or mistrust a recommendation? Does trust build over time? Do explanations help to establish trust? To which extent is human-perceived trust reflected in computational metrics of, e.g., fairness, robustness, or interpretability?

**Important Dates:**

- Submissions due: April 1, 2023
- Initial reviews: May 30, 2023
- Revisions due: July 30, 2023
- Final notifications: October 15, 2023

**Submission Information:**

We solicit original research, survey articles, and perspective/opinion papers addressing one or more of the above-mentioned topics in the context of trustworthy recommender systems. Submissions must be prepared according to the TORS submission guidelines (<https://dl.acm.org/journal/tors/author-guidelines>) and must be submitted via Manuscript Central (<https://mc.manuscriptcentral.com/tors>). Please note that the limits for paper length are slightly adapted from those imposed for submissions to the main journal, i.e. 20-30 pages for original research and surveys and 10-15 pages for opinion and perspective papers.

We seek original, high-quality contributions that have not been submitted to or accepted by other journals or conferences. Significantly expanded versions of conference or workshop papers (at least 30% original content) will be considered, too. In this case, the authors must ensure that their submission includes a citation to the prior publication and a detailed explanation of the changes between the submitted paper and the previous version(s) must be given in the cover letter. The authors must additionally emphasize, in a dedicated area of the manuscript, the innovative contributions made in comparison to the preliminary version(s), in order to aid the guest editors and reviewers in identifying the changes. All the papers should make explicit the connections of their work with recommender systems or related systems. If authors are in doubt whether their research matches the topics of interest, the guest editors should be contacted.