# International Workshop on Algorithmic Bias in Search and Recommendation (BIAS 2025)

Alejandro Bellogín
Universidad Autónoma de Madrid
Madrid, Spain
alejandro.bellogin@uam.es

Ludovico Boratto
University of Cagliari
Cagliari, Italy
ludovico.boratto@acm.org

Styliani Kleanthous
Open University of Cyprus
Nicosia, Cyprus
styliani.kleanthous@ouc.ac.cy

Elisabeth Lex
Graz University of Technology
Graz, Austria
elisabeth.lex@tugraz.at

Francesca Maridina Malloci
University of Cagliari
Cagliari, Italy
francescam.malloci@unica.it

Mirko Marras
University of Cagliari
Cagliari, Italy
mirko.marras@acm.org

## Abstract

Designing search and recommendation models that are both efficient and effective has long been a central objective for both industry professionals and academic researchers. Yet, growing evidence highlights how models trained on historical data can reinforce pre-existing biases, potentially leading to harmful outcomes. Addressing these challenges by defining, evaluating, and mitigating bias across development workflows is a crucial step toward the responsible deployment of search and recommendation models in practice. The BIAS 2025 workshop seeks to gather innovative research and foster a shared space for dialogue among researchers and practitioners committed to advancing this fundamental direction. Workshop website: https://biasinrecsys.github.io/sigir2025/.

## CCS Concepts

• **Information systems** → **Information retrieval**; • **Social and professional topics** → **User characteristics**.

## Keywords

Information Retrieval, Search, Recommendation, Bias, Fairness.

**ACM Reference Format:**
Alejandro Bellogín, Ludovico Boratto, Styliani Kleanthous, Elisabeth Lex, Francesca Maridina Malloci, and Mirko Marras. 2025. International Workshop on Algorithmic Bias in Search and Recommendation (BIAS 2025). In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25), July 13–18, 2025, Padua, Italy*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3726302.3730357

## 1 Scope and Relevance

Search and recommendation models play a central role in providing users with suggestions tailored to their preferences and needs. Despite the specificity that distinguishes these two types of model, both fundamentally rely on patterns extracted from historical data, often inheriting and embedding flawed assumptions — referred to as *biases* — that manifest as *imbalances* and *inequalities*.

In practice, the search and recommendation models trained on such biased data frequently reinforce these distortions, i.e., perpetuating biases within the patterns they learn [6]. When such biases target individuals or groups based on protected characteristics such as gender and religion, the effects can be severe, leading to inequalities and forms of discrimination and unfair treatment [8, 14].

Addressing these challenges is critical, particularly in real-world deployments, where it becomes essential, as examples, to mitigate the impact of popularity bias in order to improve perceived result quality [12, 13, 15], to ensure fairness for both consumers and providers in rankings [4, 5, 10, 11], and to improve transparency regarding the origins of biased outcomes [1–3, 7, 9]. Ultimately, the ability to *measure*, *characterize*, and *mitigate* bias, while maintaining high levels of effectiveness, remains a major open challenge.

Given the rapid evolution of search and recommendation models, the BIAS 2025 workshop offers a timely platform for discussing the latest research on bias in these models. Building on the success of five previous editions held at ECIR (2020-2023, which contributed to the creation of the independent track "IR for Good" at ECIR 2024 and 2025) and ACM SIGIR (2024), our workshop aims to foster a vibrant dialogue. The workshop will focus on raising awareness around bias in information retrieval, exploring the various dimensions affected by bias, encouraging contributions that tackle these challenges, gaining insights into recent advancements and outstanding issues, familiarizing the community with emerging practices, and identifying gaps in both academic research and industry practice.

## 2 Organization and Outcomes

The workshop welcomes participants who are interested in examining and addressing the effects of biases within search and recommendation models. It brings together a diverse audience, including researchers from the fields of *information retrieval*, *data mining*, and *machine learning*, as well as *practitioners* from both *academic institutions* and *industry sectors*. Recognizing that bias spans multiple dimensions beyond models alone, the workshop explicitly invites contributions from a broad interdisciplinary community.

Submissions to the workshop include full research papers, reproducibility studies, short papers, and position papers. Contributions focus primarily, though not exclusively, on several critical areas:


*SIGIR '25, Padua, Italy*

ACM ISBN 979-8-4007-1592-1/2025/07
https://doi.org/10.1145/3726302.3730357

*dataset collection and preparation*, such as investigating the interaction between bias, data imbalance, and rare classes, or designing pipelines that promote less biased datasets; *countermeasure development*, including frameworks for bias assessment, exploratory studies identifying new forms of bias, explainability techniques, and novel sampling strategies; *evaluation protocols and metrics*, involving the creation of objective measures for fairness and bias auditing or the establishment of bias-aware evaluation practices; and *case study analyses*, highlighting practical applications across domains such as e-commerce, education, and healthcare.

The submitted papers undergo a rigorous peer-review process, in which at least *three program committee members* assess each submission. The review considers multiple factors, including the relevance to the workshop's theme, the originality and significance of the work, technical soundness, clarity and quality of presentation, the adequacy of references, and attention to reproducibility. The program committee draws on a wide range of expertise, encompassing academia, applied research, and industry, to ensure comprehensive and balanced evaluations. The accepted contributions are published in the Springer's Communications in Computer and Information Science (CCIS) series, as a post-workshop volume.

The workshop is organized into thematic sessions. Contributions are grouped into clusters of three or four papers based on common topics or application domains. Each session features spot presentations, followed by a structured discussion to reflect on the presented work. These sessions are complemented by two keynote talks: an academic keynote talk, given by Shlomo Berkovsky (Macquarie University, Australia), offering a research perspective on the challenges of bias in search and recommendation, and a regulatory keynote talk given by Erasmo Purificato (European Centre for Algorithmic Transparency, Joint Research Centre, European Commission, Italy), exploring real-world implications and opportunities.

Beyond paper presentations, the workshop emphasizes interaction and collaboration opportunities where authors, participants, and invited experts exchange ideas and critically reflect on the workshop's central themes. Moderated by the organizers, the discussion encourages cross-disciplinary dialogue and the blending of diverse viewpoints. This is followed by a problem-solving session, where participants work together on practical challenges that arise from earlier discussions. This activity fosters creativity, engagement, and the development of new solutions. Finally, the workshop concludes with closing remarks, providing a reflection on key outcomes.

## 3 Organizers' Biography

*Alejandro Bellogín* is Associate Professor at the Department of Computer Engineering of Universidad Autónoma de Madrid (Spain). Previously, he held a post-doctoral research grant associated to the Centrum Wiskunde & Informatica (CWI). His research interests include evaluation of recommender systems, in particular dimensions like fairness and reproducibility, and evaluation of information retrieval systems, such as performance prediction models. He has an extensive experience in conference organization, being general chair for UMAP 2022 as an example, but also workshop organizer at RecSys, and workshop and tutorial chair at RecSys and UMAP.

*Ludovico Boratto* is Associate Professor at the Department of Mathematics and Computer Science of the University of Cagliari (Italy). His research interests focus on recommender systems and their impact on stakeholders, considering (beyond-)accuracy evaluation metrics. He has wide experience in workshop organizations, with 10+ events organized at ECIR, IEEE ICDM, ECML-PKDD, and ACM EICS and given tutorials on recommender systems at UMAP, ICDM, WSDM, ICDE, and ECIR. He is also co-chair of the ECIR 2025' track entitled "IR for Good".

*Styliani Kleanthous* is Assistant Professor at the Open University of Cyprus, Cyprus. She received a Ph.D. in Computer Science from the University of Leeds, UK. Her main research interests and expertise are concentrated in the area of Human-AI Interaction. She specializes in exploiting psychological and social theories for understanding and modeling user perceptions when interacting with technology, and for designing intelligent and adaptive user support. Styliani has organized and chaired several workshops on algorithmic fairness, transparency and explainability at ACM UMAP and ACM IUI conferences, e.g., FairUMAP, TExSS, HAAPIE. She has been the general chair of ACM UMAP 2023.

*Elisabeth Lex* is an Associate Professor at the Department of Computer Science and Biomedical Engineering at Graz University of Technology, Austria. Her research interests include recommender systems, user modeling, data science, and trustworthy AI, with a particular focus on the dimensions of bias, explainability, and privacy. She regularly takes on the role of a track chair, workshop chair, or doctoral consortium chair at distinct conferences such as the ACM Web Conference, ACM UMAP, or ACM RecSys; in addition, she organizes workshops, for example, at ACM RecSys, and has given tutorials at high-profile venues, including SIGIR'22.

*Francesca Maridina Malloci* is Assistant Professor at the Department of Mathematics and Computer Science of the University of Cagliari (Italy). She received her BSc (2016) and MSc (2018) (cum laude) degrees in Computer Science at University of Cagliari. In 2019, she has been visiting scientist at the EURECAT Technology Centre. Her research focuses on predictive analytics, with attention to decision-making systems for multi-stakeholder contexts. She has co-authored papers in international journals and conferences.

*Mirko Marras* is Tenure-Track Assistant Professor at the Department of Mathematics and Computer Science of the University of Cagliari (Italy). His research interests focus on responsible machine learning. He has taken a leading role in chairing the BIAS workshop series and organized workshops held in conjunction with other top-tier venues, such as ECML-PKDD, WSDM, ICCV, and EDM. He has given tutorials at UMAP and ICDM 2020, and WSDM, ICDE, ECIR 2021, RecSys 2022, ECML-PKDD 2023 and 2024, and ECIR 2024. He is also co-chair of the ECIR 2025' track entitled "IR for Good".

## Acknowledgements

Financial support is acknowledged under the National Recovery and Resilience Plan (NRRP), Mission 4 Component 2 Investment 1.5 - Call for tender No.3277 published on December 30, 2021 by the Italian Ministry of University and Research (MUR) funded by the European Union – Next Generation EU. Project Code ECS0000038 – Project Title eINS Ecosystem of Innovation for Next Generation Sardinia – CUP F53C22000430001- Grant Assignment Decree No. 1056 adopted on June 23, 2022 by the MUR.

## References


[1] Andrea Atzori, Gianni Fenu, and Mirko Marras. 2022. Explaining bias in deep face recognition via image characteristics. In *Proceedings of the IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 1–10.

[2] Andrea Atzori, Gianni Fenu, and Mirko Marras. 2023. Demographic bias in low-resolution deep face recognition in the wild. *IEEE Journal of Selected Topics in Signal Processing* 17, 3 (2023), 599–611.

[3] Giacomo Balloccu, Ludovico Boratto, Christian Cancedda, Gianni Fenu, and Mirko Marras. 2023. Knowledge is power, understanding is impact: Utility and beyond goals, explanation quality, and fairness in path reasoning recommendation. In *Proceedings of the European Conference on Information Retrieval*. 3–19.

[4] Ludovico Boratto, Francesco Fabbri, Gianni Fenu, Mirko Marras, and Giacomo Medda. 2023. Counterfactual Graph Augmentation for Consumer Unfairness Mitigation in Recommender Systems. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM 2023, Birmingham, United Kingdom, October 21-25, 2023*. ACM, 3753–3757.

[5] Ludovico Boratto, Gianni Fenu, Mirko Marras, and Giacomo Medda. 2023. Practical perspectives of consumer fairness in recommendation. *Inf. Process. Manag.* 60, 2 (2023), 103208.

[6] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2020. Bias and Debias in Recommender System: A Survey and Future Directions. *CoRR* abs/2010.03240 (2020). arXiv:2010.03240 https://arxiv.org/abs/2010.03240

[7] Yashar Deldjoo, Alejandro Bellogín, and Tommaso Di Noia. 2021. Explaining recommender systems fairness and accuracy through the lens of data characteristics. *Inf. Process. Manag.* 58, 5 (2021), 102662. doi:10.1016/j.ipm.2021.102662

[8] Michael D. Ekstrand, Anubrata Das, Robin Burke, and Fernando Diaz. 2021. Fairness and Discrimination in Information Access Systems. *CoRR* abs/2105.05779 (2021). arXiv:2105.05779 https://arxiv.org/abs/2105.05779

[9] Zuohui Fu, Yikun Xian, Ruoyuan Gao, Jieyu Zhao, Qiaoying Huang, Yingqiang Ge, Shuyuan Xu, Shijie Geng, Chirag Shah, Yongfeng Zhang, and Gerard de Melo. 2020. Fairness-Aware Explainable Recommendation over Knowledge Graphs. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR*. ACM, 69–78.

[10] Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, and Yongfeng Zhang. 2021. Towards Long-term Fairness in Recommendation. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM)*. ACM, 445–453.

[11] Mesut Kaya, Derek G. Bridge, and Nava Tintarev. 2020. Ensuring Fairness in Group Recommendations by Rank-Sensitive Balancing of Relevance. In *RecSys 2020: Fourteenth ACM Conference on Recommender Systems, Virtual Event, Brazil, September 22-26, 2020*. ACM, 101–110.

[12] Pablo Sánchez, Alejandro Bellogín, and Ludovico Boratto. 2023. Bias characterization, assessment, and mitigation in location-based recommender systems. *Data Min. Knowl. Discov.* 37, 5 (2023), 1885–1929.

[13] Tianxin Wei, Fuli Feng, Jiawei Chen, Ziwei Wu, Jinfeng Yi, and Xiangnan He. 2021. Model-Agnostic Counterfactual Reasoning for Eliminating Popularity Bias in Recommender System. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, 1791–1800.

[14] Meike Zehlike, Ke Yang, and Julia Stoyanovich. 2021. Fairness in Ranking: A Survey. *CoRR* abs/2103.14000 (2021). arXiv:2103.14000 https://arxiv.org/abs/2103.14000

[15] Ziwei Zhu, Yun He, Xing Zhao, Yin Zhang, Jianling Wang, and James Caverlee. 2021. Popularity-Opportunity Bias in Collaborative Filtering. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM)*. ACM, 85–93.