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

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,
CYENS Centre of Excellence
Nicosia, Cyprus
s.kleanthous@cyens.org.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

Creating efficient and effective search and recommendation algorithms has been the main objective of industry practitioners and academic researchers over the years. However, recent research has shown how these algorithms trained on historical data lead to models that might exacerbate existing biases and generate potentially negative outcomes. Defining, assessing, and mitigating these biases throughout experimental pipelines is a primary step for devising search and recommendation algorithms that can be responsibly deployed in real-world applications. This workshop aims to collect novel contributions in this field and offer a common ground for interested researchers and practitioners. More information about the workshop is available at https://biasinrecsys.github.io/sigir2024/.

CCS CONCEPTS

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

KEYWORDS

Bias, Algorithms, Search, Recommendation, Fairness.

ACM Reference Format:
Alejandro Bellogín, Ludovico Boratto, Styliani Kleanthous, Elisabeth Lex, Francesca Maridina Malloci, and Mirko Marras. 2024. International Workshop on Algorithmic Bias in Search and Recommendation (BIAS). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24), July 14–18, 2024, Washington, DC, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3626772.3657990

1 INTRODUCTION

Search and recommendation algorithms provide users with suggestions that match their needs and interests. Notwithstanding the unique features of each algorithm category in terms of personalization, both derive patterns from historical data, embedding erroneous assumptions, namely, *biases*, manifested as *imbalances* and *inequalities*.

In many instances, the trained algorithms, and consequently their final output, entrench the inherent biases in the data, by reinforcing the cognitive bias in the learned patterns [3]. When bias manifests against individuals or groups with legally protected characteristics (e.g., gender, religion, race, or ideology), search and recommendation algorithms can exacerbate inequalities, leading to *profound societal consequences* as discrimination and unfairness [5, 12].

In real-world applications, there are notable challenges that demand particular focus, including mitigating the impacts of popularity bias to enhance user perceptions of result quality [1, 10, 11, 13], guaranteeing fair rankings for both consumers and providers [2, 7–9], and transparently explaining why a model yields a biased result [4, 6]. Consequently, the capacity to *measure*, *characterize*, and *mitigate* while preserving effectiveness remains a key challenge.

With the rapidly-changing methods for search and recommendation, Bias 2024 represents a timely workshop on the recent advances in bias within search and recommendation models. Specifically, BIAS 2024 come after four very successful runs between 2020 and 2023 at ECIR[1], which led to an independent track, entitled "IR for Good", at ECIR 2024. Specifically, the goal of BIAS 2024 is to favor a dialogue aimed to foster the following objectives:

(1) Increase awareness of the algorithmic bias problem in IR.
(2) Identify dimensions influenced by algorithmic bias in IR.
(3) Solicit contributions addressing algorithmic bias in IR.
(4) Gain insights into recent advances and open issues in IR.
(5) Familiarize the IR community with current field practices.
(6) Uncover gaps in academic and industry research in IR.

The event is also supported by the ACM Conference on Fairness, Accountability, and Transparency (FAccT) Network.

2 WORKSHOP PLAN

This workshop is aimed at individuals interested in studying and mitigating the effects generated by biases in search and recommendation, from *information retrieval*, *data mining*, and *machine*

[1] https://biasinrecsys.github.io/ecir2023/

SIGIR '24, July 14–18, 2024, Washington, DC, USA

ACM ISBN 979-8-4007-0431-4/24/07
https://doi.org/10.1145/3626772.3657990

learning researchers to *practitioners* from both *academic institutions* and *industry*. One aspect emerging from the aims of our workshop is that bias is a highly interdisciplinary topic, touching on several dimensions beyond algorithms. Hence, our workshop is of interest to an interdisciplinary audience, with different backgrounds.

Submitted papers falling into the following categories: full papers, reproducibility papers, short papers, and position papers. Their topics center on, though are not restricted to, *data set collection and preparation* (e.g., examining how bias and imbalanced data or rare classes interact, developing collection pipelines that result in fairer and less biased datasets), *countermeasure design and development* (e.g., bias and fairness assessing frameworks, exploratory analyses that uncover novel types of bias, explainable bias methods, sampling techniques on dataset composition), *evaluation protocol and metric formulation* (e.g., objective metrics definition that consider bias an fairness, auditing studies with respect to bias and fairness, establishing bias-aware evaluation protocols), and *case study exploration* (e.g., in e-commerce, education and healthcare).

The accepted contributions are selected with a rigorous peer-review process involving at least *three programme committee members*. Decisions consider the relevance for the workshop, novelty /originality, significance, technical correctness, clarity of presentation, quality of references, and reproducibility. The committee consists of experts with backgrounds spanning academia, research, and industry, ensuring a wide range of viewpoints during the evaluation process. The proceedings will be published as volumes into the Springer's CCIS.

The workshop is structured as a sequence of spot thematic and interactive sessions, lasting 45 minutes each. Specifically, the accepted contributions are split into groups of 3/4 papers, based on the topic or application domain. Each session combines spot paper talks (6 minutes each, 24 minutes in total) and spot discussion on their outcomes (16 minutes). These novel thematic sessions are paired with two keynote talks. An academic keynote speaker gives a talk on the state of the art on a novel perspective of algorithmic bias in IR research. Moreover, an industry keynote speaker opens up on the challenges and opportunities in these areas for the industry.

The workshop also fosters interaction and collaborative engagement among participants through an interdisciplinary round table, where paper authors, attendees, and possibly invited experts, from diverse backgrounds engage in a comprehensive discussion, exchanging insights and perspectives on the workshop's themes. This round table, moderated by the organizers, aims to bridge disciplinary boundaries, encouraging the synthesis of ideas and methodologies. Following the round table, the workshop features a group problem-solving session lasting 60 minutes, providing participants with a platform to collaboratively address challenges or explore potential solutions that emerged in the previous discussion. This hands-on activity is designed to stimulate creativity and collective problem-solving, fostering a dynamic and participatory environment. Finally, the workshop ends with remarks, offering a reflective summary and key takeaways of the day's proceedings.

3 WORKSHOP ORGANIZERS

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. 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 various editions of RecSys or UMAP.

Ludovico Boratto is Assistant 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 2024' track entitled "IR for Good".

Styliani Kleanthous is a Research Team Leader of the Fairness and Ethics in AI-Human Interaction (fAIre) group at CYENS CoE, 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 regularly organizes workshops, for example, at ACM RecSys, and 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 has been visiting scientist at the EURECAT Technology Centre. Her research focuses on predictive analytics, with attention to decision-making systems, such as recommender systems, for multi-stakeholder contexts. She has co-authored papers in international journals and conferences. Francesca has chaired workshop on knowledge discovery at ACM UMAP and ECML-PKDD.

Mirko Marras is Assistant Professor at the Department of Mathematics and Computer Science of the University of Cagliari (Italy). His research interests focus on responsible machine learning, with particular attention to educational environments. He has taken a leading role when chairing the first editions of the ECIR BIAS workshop (2020-23) and workshops held in conjunction with other

top-tier venues, such as WSDM, ICCV, and EDM. He is giving tutorials on bias and explainability in recommender systems at UMAP and ICDM 2020, and WSDM, ICDE, ECIR 2021, RecSys 2022, ECML-PKDD 2023, and ECIR 2024. He is also co-chair of the ECIR 2024' track entitled "IR for Good".

REFERENCES

[1] Ludovico Boratto, Gianni Fenu, and Mirko Marras. 2021. Connecting user and item perspectives in popularity debiasing for collaborative recommendation. *Inf. Process. Manag.* 58, 1 (2021), 102387.

[2] Ludovico Boratto, Gianni Fenu, and Mirko Marras. 2021. Interplay between upsampling and regularization for provider fairness in recommender systems. *User Model. User Adapt. Interact.* 31, 3 (2021), 421–455.

[3] 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

[4] 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. https://doi.org/10.1016/j.ipm.2021.102662

[5] 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

[6] 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 2020, Virtual Event, China, July 25-30, 2020.* ACM, 69–78.

[7] 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 *WSDM '21, The Fourteenth ACM International Conference on Web Search and Data Mining, Virtual Event, Israel, March 8-12, 2021.* ACM, 445–453.

[8] Elizabeth Gómez, Carlos Shui Zhang, Ludovico Boratto, Maria Salamó, and Mirko Marras. 2021. The Winner Takes it All: Geographic Imbalance and Provider (Un)fairness in Educational Recommender Systems. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021.* ACM, 1808–1812.

[9] 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.

[10] 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. https://doi.org/10.1007/S10618-022-00913-5

[11] 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 *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, Singapore, August 14-18, 2021.* ACM, 1791–1800.

[12] 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

[13] Ziwei Zhu, Yun He, Xing Zhao, Yin Zhang, Jianling Wang, and James Caverlee. 2021. Popularity-Opportunity Bias in Collaborative Filtering. In *WSDM '21, The Fourteenth ACM International Conference on Web Search and Data Mining, Virtual Event, Israel, March 8-12, 2021.* ACM, 85–93.