





# International Workshop on Algorithmic Bias in Search and Recommendation (Bias 2020)

Ludovico Boratto<sup>1</sup> , Mirko Marras<sup>2</sup> , Stefano Faralli<sup>3</sup> ,  
and Giovanni Stilo<sup>4</sup> 

<sup>1</sup> EURECAT, Barcelona, Spain  
ludovico.boratto@acm.org

<sup>2</sup> University of Cagliari, Cagliari, Italy  
mirko.marras@unica.it

<sup>3</sup> University of Rome Unitelma Sapienza, Rome, Italy  
stefano.faralli@unitelma.it

<sup>4</sup> University of L'Aquila, L'Aquila, Italy  
giovanni.stilo@univaq.it

**Abstract.** Both search and recommendation algorithms provide results based on their relevance for the current user. In order to do so, such a relevance is usually computed by models trained on historical data, which is biased in most cases. Hence, the results produced by these algorithms naturally propagate, and frequently reinforce, biases hidden in the data, consequently strengthening inequalities. Being able to measure, characterize, and mitigate these biases while keeping high effectiveness is a topic of central interest for the information retrieval community. In this workshop, we aim to collect novel contributions in this emerging field and to provide a common ground for interested researchers and practitioners.

**Keywords:** Bias · Algorithms · Search · Recommendation

## 1 Motivation

*Search* and *recommendation* are getting closer and closer as research areas. Though they require fundamentally different inputs, i.e., the user is asked to provide a query in search, while implicit and explicit feedback is leveraged in recommendation, existing search algorithms are being personalized based on users' profiles and recommender systems are optimizing their output on the ranking quality.

Both classes of algorithms aim to learn patterns from historical data that conveys biases in terms of *unbalances* and *inequalities*. These hidden biases are unfortunately captured in the learned patterns, and often emphasized in the results these algorithms provide to users [2]. When a bias affects a *sensitive attribute* of a user, such as their gender or religion, the inequalities that are

reinforced by search and recommendation algorithms even lead to *severe societal consequences*, like users' discrimination [4].

For this critical reason, being able to *detect, measure, characterize, and mitigate* these biases while keeping high effectiveness is a prominent and timely topic for the IR community. Mitigating the effects generated by popularity bias [1, 5, 6], ensuring results that are fair with respect to the users [3, 7], and being able to interpret why a model provides a given recommendation or search result are examples of challenges that may be important in real-world applications. This workshop aims to collect new contributions in this emerging field and to provide a common ground for interested researchers and practitioners.

The workshop welcomes contributions in all topics related to algorithmic bias in search and recommendation, focused (but not limited) to:

- *Data Set Collection and Preparation:*
  - Managing imbalances and inequalities within data sets.
  - Devising collection pipelines that lead to fair and unbiased data sets.
  - Collecting data sets useful for studying potential biased and unfair situations.
  - Designing procedures for creating synthetic data sets for research on bias and fairness.
- *Countermeasure Design and Development:*
  - Conducting exploratory analysis that uncover biases.
  - Designing treatments that mitigate biases (e.g., popularity bias mitigation).
  - Devising interpretable search and recommendation models.
  - Providing treatment procedures whose outcomes are easily interpretable.
  - Balancing inequalities among different groups of users or stakeholders.
- *Evaluation Protocol and Metric Formulation:*
  - Conducting quantitative experimental studies on bias and unfairness.
  - Defining objective metrics that consider fairness and/or bias.
  - Formulating bias-aware protocols to evaluate existing algorithms.
  - Evaluating existing strategies in unexplored domains
- *Case Study Exploration:*
  - News channels.
  - E-commerce platforms.
  - Educational environments.
  - Entertainment websites.
  - Healthcare systems.
  - Social networks.

## 2 Objectives

The workshop has the following main objectives:

1. Raise awareness on the algorithmic bias problem within the IR community.
2. Identify social and human dimensions affected by algorithmic bias in IR.
3. Solicit contributions from researchers who are facing algorithmic bias in IR.
4. Get insights on existing approaches, recent advances, and open issues.
5. Familiarize the IR community with existing practices from the field.
6. Uncover gaps between academic research and real-world needs in the field.

### 3 Organizers Biography

*Ludovico Boratto* is senior research scientist in the Data Science and Big Data Analytics research group at Eurecat. His research interests focus on Data Mining and Machine Learning approaches, mostly applied to recommender systems and social media analysis. The results of his research have been published in top-tier conferences and journals. His research activity also brought him to give talks and tutorials at top-tier conferences (e.g., ACM RecSys 2016, IEEE ICDM 2017) and research centers (Yahoo! Research). He is editor of the book “Group Recommender Systems: An Introduction”, published by Springer. He is editorial board member of the “Information Processing & Management” journal (Elsevier) and guest editor of several journal’s special issues. He is regularly part of the program committee of the main Data Mining and Web conferences, such as RecSys, KDD, SIGIR, WSDM, ICWSM, and TheWebConf. In 2012, he got a Ph.D. at the University of Cagliari (Italy), where he was research assistant until May 2016. In 2010 and 2014 he spent 10 months at Yahoo! Research in Barcelona as a visiting researcher. He is member of the ACM and of the IEEE.

*Mirko Marras* is a PhD student in Computer Science at the Department of Mathematics and Computer Science of the University of Cagliari (Italy). He received the MSc Degree in Computer Science (summa cum laude) from the same University in 2016. His research interests focus on algorithmic bias in machine learning for educational platforms, specifically in the context of semantic-aware systems, recommender systems, biometric systems, and opinion mining systems. He has co-authored papers in top-tier international journals, such as Pattern Recognition Letters (Elsevier), Computers in Human Behavior (Elsevier), and IEEE Cloud Computing. He has given talks and demonstrations at several international conferences and workshops, such as The Web Conference 2018, ECIR 2019, ESWC2017, INTERSPEECH 2019. He is student member in several national and international associations, including CVPL, AIXIA, IEEE, and ACM.

*Stefano Faralli* is an assistant professor at University of Rome Unitelma Sapienza, Rome, Italy. His research interests include Ontology Learning, Distributional Semantics, Word Sense Disambiguation/Induction, Recommender Systems, Linked Open Data. He co-organized the International Workshop: Taxonomy Extraction Evaluation (TexEval) Task 17 of Semantic Evaluation (SemEval-2015) and the International Workshop on Social Interaction-based Recommendation (SIR 2018).

*Giovanni Stilo* is an Assistant Professor in the Department of Information Engineering, Computer Science and Mathematics at the University of L’Aquila. He received his PhD. in Computer Science in 2013, and in 2014 he was a visiting researcher at Yahoo! Labs in Barcelona. Between 2015 and 2018, he was a researcher in the Computer Science Department at La Sapienza University, in Rome. His research interests are in the areas of machine learning and data mining, and specifically temporal mining, social network analysis, network medicine, semantics-aware recommender systems, and anomaly detection. He is a member

of the steering committee of the Intelligent Information Mining research group (<http://iim.disim.univaq.it/>). He has organized several international workshops, held in conjunction with top-tier conferences (ICDM, CIKM, and ECIR), and he is involved as editor and reviewer of top-tier journals, such as TITS, TKDE, DMKD, AI, KAIS, and AIIM.

## References

1. Abdollahpouri, H., Burke, R., Mobasher, B.: Controlling popularity bias in learning-to-rank recommendation. In: Proceedings of the Eleventh ACM Conference on Recommender Systems, pp. 42–46. ACM (2017)
2. Boratto, L., Fenu, G., Marras, M.: The effect of algorithmic bias on recommender systems for massive open online courses. In: Azzopardi, L., Stein, B., Fuhr, N., Mayr, P., Hauff, C., Hiemstra, D. (eds.) ECIR 2019. LNCS, vol. 11437, pp. 457–472. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-15712-8\\_30](https://doi.org/10.1007/978-3-030-15712-8_30)
3. Burke, R., Sonboli, N., Ordonez-Gauger, A.: Balanced neighborhoods for multi-sided fairness in recommendation. In: Conference on Fairness, Accountability and Transparency, pp. 202–214 (2018)
4. Hajian, S., Bonchi, F., Castillo, C.: Algorithmic bias: from discrimination discovery to fairness-aware data mining. In: Krishnapuram, B., Shah, M., Smola, A.J., Aggarwal, C.C., Shen, D., Rastogi, R. (eds.) Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016, pp. 2125–2126. ACM (2016). DOI: <https://doi.org/10.1145/2939672.2945386>
5. Jannach, D., Lerche, L., Kamehkhosh, I., Jugovac, M.: What recommenders recommend: an analysis of recommendation biases and possible countermeasures. *User Modeling and User-Adapted Interaction* **25**(5), 427–491 (2015). <https://doi.org/10.1007/s11257-015-9165-3>
6. Kamishima, T., Akaho, S., Asoh, H., Sakuma, J.: Correcting popularity bias by enhancing recommendation neutrality. In: RecSys Posters (2014)
7. Zheng, Y., Dave, T., Mishra, N., Kumar, H.: Fairness in reciprocal recommendations: a speed-dating study. In: Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization, pp. 29–34. ACM (2018)