

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

Ludovico Boratto  
EURECAT, Centre Tecnològic de Catalunya  
*ludovico.boratto@acm.org*

Stefano Faralli  
University of Rome Unitelma Sapienza  
*stefano.faralli@unitelma.it*

Mirko Marras  
University of Cagliari  
*mirko.marras@unica.it*

Giovanni Stilo  
University of L'Aquila  
*giovanni.stilo@univaq.it*

## Abstract

The International Workshop on Algorithmic Bias in Search and Recommendation was held on April 14, 2020 in conjunction with the 42nd European Conference on Information Retrieval (ECIR 2020). The scientific program included paper and demo presentations and a final discussion. The keynote was delivered by Prof Chirag Shah. This report presents an overview of the activities conducted during the workshop and the main topics covered in it.

## 1 Introduction

On April 14, 2020, the International Workshop on Algorithmic Bias in Search and Recommendation<sup>1</sup> [Boratto et al., 2020] was held online, in conjunction with the 42nd European Conference on Information Retrieval (ECIR 2020). The workshop was organized by the Data Science and Big Data Analytics unit at Eurecat (Spain), by the University of Rome Unitelma Sapienza (Italy), by the Department of Mathematics and Computer Science at the University of Cagliari (Italy), and by the Department of Information Engineering, Computer Science and Mathematics at the University of L'Aquila (Italy). The aim of the workshop was to collect novel ideas to detect, measure, characterize, and mitigate bias in the data and the algorithms underlying search and

---

<sup>1</sup><http://bias.disim.univaq.it/>

---

recommendation applications [Hajian et al., 2016, Jannach et al., 2015], to provide a common ground for researchers working in this area.

The workshop had more than 70 participants. The keynote speaker was Prof. Chirag Shah from the University of Washington (USA). The scientific program included demo and paper presentations and a final discussion. The papers covered topics that go from search and recommendation in online dating, education, and social media, over the impact of gender bias in word embeddings, to tools that allow to explore bias and fairness on the Web. The program included also a special session devoted to social aspects and implications. The event concluded with a discussion session to highlight open issues, research challenges, and briefly summarize the outcomes of the workshop.

## 2 Summary of the keynote talk

**“Investigating Bias and Instigating Fairness in Search and Recommendation”, by Chirag Shah.** Bias exists everywhere, from data to algorithms, and from framing a problem to interpreting its solution. In this talk, Dr. Shah highlighted how such bias, especially in relation with search and recommender systems, causes material problems for users, businesses, and society at large. The examples spanned areas of search, education, and healthcare. To find a balance or fairness in the system and address the issue of bias, the concept of a marketplace was introduced, to lead to a more sustainable growth for various industries, governments, and our scientific advancement.

## 3 Summary of the papers

**“Bias Goggles - Exploring the bias of Web Domains through the Eyes of the Users”.** This demo, by Giannis Konstantakis, Ioannis Promponas, Manthos Dretakis, and Panagiotis Papadakos, presented the *bias goggles* system, which computes the bias characteristics of Web domains to user-defined concepts, using the structure of the Web graph. The system is based on propagation models and a variation of PageRank, named Biased-PR, which models various behaviors of biased surfers. Cross-browser plugins let users explore the bias characteristics of domains, along with analytics of user behavior.

**“Facets of Fairness in Search and Recommendation”.** This short paper, by Sahil Verma, Ruoyuan Gao, and Chirag Shah, examines some of the recent works to define relevance, diversity, fairness, and related concepts. Then, it shows how various metrics to measure such concepts are materialized on search results. In doing so, this work highlights comparisons and contrasts among various measures, and gaps in our conceptual and evaluative frameworks.

**“Mitigating Gender Bias in Machine Learning Data Sets”.** This work by Susan Leavy, Gerardine Meaney, Karen Wade, and Derek Greene, mitigates bias on algorithms trained on textual data, considering the complex way gender ideology is embedded in language. This paper proposes a framework for the evaluation of gender bias in machine learning. The work draws upon gender theory and sociolinguistics to systematically indicate levels of bias in textual training data and associated neural word embedding models. It finally highlights pathways for both removing

---

bias from training data and critically assessing its impact in the context of search and recommender systems.

**“Why do we need to be bots? What prevents society from detecting biases in recommendation systems”**. This paper, by Tobias D. Krafft, Marc P. Hauer, and Katharina Anna Zweig, describes the conditions to be met to allow black box analyses of recommendation systems, based on an application on Facebook’s News Feed. Drawing on the insights from this case study and the state of the art of research on algorithmic accountability, the authors formulate several policy demands that need to be met in order to allow monitoring of algorithmic decision-making systems, for their compliance with social values.

**“Matchmaking Under Fairness Constraints: a Speed Dating Case Study”**. This work, by Dimitris Paraschakis and Bengt J. Nilsson, focuses on the ethical aspects of recommendations in the context of dating services. Specifically, the authors assess the capability of an algorithm to respect the user’s expressed sensitive preferences (e.g., racial, religious). To overcome this issue, they introduce the notion of preferential fairness, and propose two approaches for re-ranking the recommendations under these constraints. Experimental results demonstrate that the state of fairness can be reached with minimal accuracy compromises for both binary and non-binary attributes.

**“Effect of Debiasing on Information Retrieval”**. This work, by Emma Gerritse and Arjen de Vries, considers several algorithms for removing bias from word embeddings. The rationale is that debiasing should make the word embeddings fairer in their use, avoiding potential negative effects downstream (e.g., word embeddings with a gender bias that are used in a classification task in a hiring process). Specifically, the authors compare regular and debiased word embeddings in an Information Retrieval task. The outcomes of this study show that the difference between the outputs returned by the two methods is present, but not very big.

**“Using String-Comparison measures to Improve and Evaluate Collaborative Filtering Recommender Systems”**. This work, by Luiz Mário Lustosa Pascoal, Hugo Alexandre Dantas Do Nascimento, Thierson Couto Rosa, Edjalma Queiroz da Silva and Everton Lima Aleixo, proposes the usage of traditional string comparison approaches for defining new similarity measures between pairs of users, taking into consideration time-ordered data in Collaborative Filtering tasks. Experiments conducted in different application domains demonstrate that the measures based on a string-comparison approach can improve accuracy.

**“Recommendation filtering à la carte for intelligent tutoring systems”**. This work, by Wesley Perreira, Marcos Spalenza, Jean-Rémi Bourguet and Elias de Oliveira, considers collaborative and content-based filters as solutions of a data-centric approach to uncover the student abilities in computerized adaptive testing. Given that past lecturer recommendations can induce possible bias during the warm starts, the authors try to reduce this issue by proposing a hybrid recommendation filtering. The approach is supported by the Item Response Theory and clustering techniques that make it possible to output purely objective recommendation filters. This is

---

done by selecting activities and building an evaluation path based on historical evolution of past students.

**“Social aspects and implications” track.** Issues like algorithmic bias and fairness clearly have a societal impact. For this reason, the workshop continued with a special track on “Social aspects and implications”, to address the impact of behavioral data of online social users in recommendation and search applications. This was done through the following spot talks:

- Guilherme Ramos and Carlos Caleiro. A Novel Similarity Measure for Group Recommender Systems with Optimal Time Complexity.
- Tiago de Melo, Altigran S. da Silva, Edleno S. de Moura and Pável Calado. Enriching Product Catalogs with User Opinions.
- Alessandro Celi, Romina Eramo, Alejandro Piad and Jòsvald Diaz Blanco. Analyzing the Interaction of Users with News Articles to Create Personalization Services.
- Lorena Recalde and Ricardo Baeza-Yates. What kind of content are you prone to tweet? Multi-topic Preference Model for Tweeters.
- Xiao Bai, B. Barla Cambazoglu, Francesco Gullo, Amin Mantrach and Fabrizio Silvestri. Improving News Personalization through Search Logs.
- Seyed Amin Mirlohi Falavarjani, Hawre Hosseini and Ebrahim Bagheri. The Impact of Foursquare Checkins on Users’ Emotions on Twitter.
- Javier Sanz-Cruzado and Pablo Castells. Beyond Accuracy in Link Prediction.
- Sihem Amer-Yahia, Anh Tho Le and Eric Simon. Data Pipelines for Personalized Exploration of Rated Datasets.
- Mohammad Aliannejadi and Fabio Crestani. Venue Suggestion Using Social-Centric Scores.

As these studies show, the analysis of social behavior provides a complementary perspective that can feed the algorithms and provide support in the identification and mitigation of possible societal issues.

## 4 Conclusions

Overall, the International Workshop on Algorithmic Bias in Search and Recommendation was a success, both in terms of number of participants and of interests that emerged during the presentations and the final discussion, creating new relationships and novel ideas in this area. Tentative plans to organize the second workshop next year were formed. Additionally, we serve as guest editors of a special issue<sup>2</sup> in the Information Processing & Management journal on these

---

<sup>2</sup><https://www.journals.elsevier.com/information-processing-and-management/call-for-papers/special-issue-on-algorithmic-bias-and-fairness-in-search>

---

topics. Furthermore, a proceedings volume is currently under production and should be published soon via Springer’s “Communications in Computer and Information Science” series.

The organizers would like to thank the authors of the submitted papers, the reviewers for allowing us to shape an interesting program, and the attendants for their active participation in the activities. We are open to share experiences and arrange joint initiatives in this prominent research area. We would like to invite interested researchers and practitioners to follow our website<sup>3</sup> and stay in touch via email, for ongoing and future initiatives in this area.

## References

- Ludovico Boratto, Mirko Marras, Stefano Faralli, and Giovanni Stilo. International workshop on algorithmic bias in search and recommendation (bias 2020). In Joemon M. Jose, Emine Yilmaz, João Magalhães, Pablo Castells, Nicola Ferro, Mário J. Silva, and Flávio Martins, editors, *Advances in Information Retrieval - 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14-17, 2020, Proceedings, Part II*, volume 12036 of *Lecture Notes in Computer Science*, pages 637–640. Springer, 2020. doi: 10.1007/978-3-030-45442-5\\_84. URL [https://doi.org/10.1007/978-3-030-45442-5\\_84](https://doi.org/10.1007/978-3-030-45442-5_84).
- Sara Hajian, Francesco Bonchi, and Carlos Castillo. Algorithmic bias: From discrimination discovery to fairness-aware data mining. In Balaji Krishnapuram, Mohak Shah, Alexander J. Smola, Charu C. Aggarwal, Dou Shen, and Rajeev Rastogi, editors, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, pages 2125–2126. ACM, 2016. doi: 10.1145/2939672.2945386. URL <https://doi.org/10.1145/2939672.2945386>.
- Dietmar Jannach, Lukas Lerche, Iman Kamehkhosh, and Michael Jugovac. What recommenders recommend: an analysis of recommendation biases and possible countermeasures. *User Model. User-Adapt. Interact.*, 25(5):427–491, 2015. doi: 10.1007/s11257-015-9165-3. URL <https://doi.org/10.1007/s11257-015-9165-3>.

---

<sup>3</sup><http://bias.disim.univaq.it/>