Beyond Algorithmic Fairness in Recommender Systems

Preprint · July 2021

DOI: 10.1145/3450614.3461685

4 authors, including:

Mehdi Elahi
University of Bergen

Himan Abdollahpouri
Spotify

Masoud Mansoury
University of Amsterdam

Beyond Algorithmic Fairness in Recommender Systems

Mehdi Elahi
University of Bergen
Bergen, Norway
mehdi.elahi@uib.no

Himan Abdollahpouri
Northwestern University
Evanston, USA
himan.abdollahpouri@northwestern.edu

Masoud Mansoury
Eindhoven University of Technology
Eindhoven, Netherlands
m.mansoury@tue.nl

Helma Torkamaan
University of Duisburg-Essen
Duisburg, Germany
h.torkamaan@acm.org

ABSTRACT

Fairness is one of the crucial aspects of modern Recommender Systems which has recently drawn substantial attention from the community. Many recent works have addressed this aspect by studying the fairness of the recommendation through different forms of evaluation methodologies and metrics. However, the majority of these works have mainly concentrated on the recommendation algorithms and hence measured the fairness from the algorithmic viewpoint. While such viewpoint may still play an important role, it does not necessarily project a comprehensive picture of how the users may perceive the overall fairness of a recommender system.

This paper extends the prior works and goes beyond the algorithmic fairness in recommender systems by highlighting the non-algorithmic viewpoint on the fairness in these systems. The paper proposes an evaluation methodology that can be used to assess the fairness of a recommender system perceived by its users. We have adopted a well-known model and re-formulated it to suit the particular characteristics of the recommender systems, and accordingly, their corresponding users. Our proposed methodology can be used in order to elicit the feedback of the users, along with three important dimensions, i.e., *Engagement*, *Representation*, and *Action & Expression*. We have formed a set of survey questions that address the aforementioned dimensions, as a set of examples to assess the fairness in a recommender system.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; **Personalization**; **Recommender systems**; **Personalization**; • **Human-centered computing** → *Accessibility*.

KEYWORDS

recommender systems, fairness, evaluation

Conference'17, July 2017, Washington, DC, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

ACM Reference Format:
Mehdi Elahi, Himan Abdollahpouri, Masoud Mansoury, and Helma Torkamaan. 2021. Beyond Algorithmic Fairness in Recommender Systems. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Nowadays, finding the right product to consume has become a grand challenge for customers due to the growing *volume, variety*, and *velocity* of products online. Recommender Systems can alleviate this issue by supporting users when making decisions on what to consume. These decision support systems can adopt sophisticated AI algorithms in order to build personalized suggestions based on the *specific* tastes and interests of the users that can better match users' needs and constraints rather than suggesting the products based on *generic* mainstream tastes.

Although recommender systems have become essential tools for users in almost any application domain, they generate or -at least-intensify a number of undesired biases. Examples of such biases are over-concentration of the algorithms on already popular items [2], or yielding inconsistent performances across different groups of users according to their gender, race, age, or particular assumptions about the users identity or characteristics [11]. A number of prior works have studied these biases and have proposed solutions to mitigate them [3, 18]. However, the topic has still remained open as no comprehensive solution has been developed for different forms of the problem at hand.

In recent years, there has been a growing interest in *fairness* and relevant ethical considerations of recommender systems [9, 21]. For example, there are numerous works regarding considerations about fair distribution of the recommendations across different groups of users [17], or different suppliers of the recommended items [23]. Furthermore, there exist works that investigated the ethical consideration of whether personalization algorithms can do harm to certain people [22]. A limitation for majority of these works is that they primarily focused on algorithmic aspects of fairness and addressed a rather "specific" form of the problem.

Algorithmic fairness and improving it for the recommender systems is certainly a crucial step towards building a fair recommendation platform, however, it is not enough. There are many non-algorithmic aspects (e.g. feedback collection and recommendation representation) that need to be examined and improved to enable a recommender system to be truly perceived as fair by its users.

Users may come from different backgrounds, have differences in terms of physical or cognitive abilities, or even their beliefs. A recommender system should take these differences into account when interacting with users and serving them with personalized recommendation. Take a user with an impaired vision as an example: she may feel much of difficulty interacting with a "system" that provides recommendation for her, regardless of what these recommendations are.

These non-algorithmic aspects need to be properly defined and evaluated to determine how fair is a recommender system from the users' perspective. For doing so, we propose a novel mechanism for evaluating the perceived fairness of a given recommender system from non-algorithmic point of view. The proposed mechanism relies on the popular theory of **Universal Design for Learning (UDL)** [26] and stands on three principle pillars of the theory by assessing the perceived fairness of the recommender system in terms of *Engagement*, *Representation*, and *Action & Expression*.

In summary, our contribution is two-fold:

- We extend the notion of fairness in recommender systems to cover other aspects of the recommendations that are widely non-algorithmic.
- We propose a novel framework based on a well-known theoretical foundation that can be used to evaluate different aspects of the fairness in recommender systems.

2 FAIRNESS IN RECOMMENDER SYSTEMS

In general, the concept of fairness in user decision making is not new and can be traced back to well over 50 years [15]. However, only recently, it has become a crucial aspect of research in popular topics, relevant to decision support tools such as recommender systems or more generally, Artificial Intelligence (AI)[1].

Research on recommender systems has drawn a growing interest surrendering the fairness topic. For example, Abdollahpouri et al. provide a taxonomy for different classes of fairness in recommender systems where the fairness is categorized based on various stakeholders: *C-fairness* where the focus is on the perspective of those who receive the recommendations (aka users), *P-Fairness* for those who provide the items or content (aka providers), and *S-fairness* for those who neither receive nor provide the recommendations yet are impacted by the outcome of a recommender system (aka side stakeholders) [1].

A recommendation algorithm can be biased towards any of the aforementioned stakeholders. For example, Ekstrand et al. [10, 20] show that some recommendation algorithms give consistently lower accuracy to certain user groups than the others. In addition, Abdollahpouri et al. [2, 16] demonstrate that many recommendation algorithms are biased towards popular items leading to unsatisfactory outcomes for users with an interest towards the niche and less popular items.

Work on fairness beyond the users' perspective is a relevant topic that received growing attention over the recent years. For example, Mehrotra et al. [19] show that, a small portion of artists (i.e., providers of the songs) on Spotify, get an overwhelmingly large number of streams on the entire platform leading to an unfair outcome for the other less popular artists. Patro et al. [23] propose an algorithm that can improve the fairness of the recommendations for both users and suppliers by imposing a minimum number of exposure to different items that ensures different suppliers getting a fair share of the recommended items.

Almost all the prior works, including the above-mentioned ones, are mainly concentrated on the *outcome* of the recommender systems in terms of what items are being recommended and from whom. For example, Smith et al. [27] explored users' opinion about the fairness of the recommendations via a user study, though it was still mainly related to the algorithmic aspect of recommendation. While algorithm is indeed an important aspect of the fairness of recommender systems, it is not the only one. In a typical recommendation process, users interact with the *system* through an interface to initially express their preferences. Then, the system analyzes the preference data and utilizes a recommendation algorithm to learn the patterns within these user preferences. Finally, the algorithm exploits the preferences to generate the personalized recommendations for users that are presented to them through the interface. The unfairness issues could arise in any of these steps.

The existing works on fairness in recommendation have been mostly focusing on core recommender algorithms. However, we argue that fairness needs considerations beyond algorithmic fairness. In this paper, we address this gap by providing a more comprehensive view towards fairness and focusing on non-algorithmic aspects, described later. In the next section, we provide more details on different aspects (or dimensions) of fairness in recommender systems and propose qualitative measures that can be utilized in order to evaluate how fair the recommender system is from the perspectives of its users.

3 NON-ALGORITHMIC ASPECTS OF FAIRNESS

As described before, the majority of the related work has mainly focused on fairness of the recommender systems from an algorithmic point of view [4, 6, 24]. Indeed, a recommender system may benefit from a sophisticated AI algorithm that is fairness-driven, i.e. enhanced to take fairness into consideration. However, fairness has a non-algorithmic aspect as well. A system may not necessarily be perceived as fair by its users even when the generated recommendations are considered as fair from an algorithmic point of view. In particular, individual expectations, physical or cognitive abilities, and cultural backgrounds can heavily influence the fairness perception of users towards a recommender system.

In this paper, we adopt a well-known framework, Universal Design for Learning (UDL) [26], which includes three aspects that are essential for a successful learning experience: Engagement, Representation, and Action & Expression. Engagement refers to how students can be motivated and engaged in a learning process [13]. Representation deals with the considerations regarding how the content and the learning materials are represented to the students [29]. And finally, Action & Expression refers to the fact that students should be offered different methods and tools to respond and interact with the learning system [25]. We believe, these aspects can be directly linked to the recommender systems and hence be adopted to address fairness of these systems. Hence, a successful

[1]A number of recent conferences has been entirely dedicated to these topics including FaccT Conference and AIES.

and fair recommender system should take all three aspects into account. In the following, each aspect will be discussed.

In the recommender systems, fairness of the core recommender algorithm mainly focuses on ensuring that the recommended items for different users meet certain fairness criteria. For instance, if a job recommender system suggests only low-paying jobs to one group of users while consistently provide recommendations of high-paying jobs to another group, regardless of their actual preferences, we consider that recommender system to be *algorithmically unfair*. Therefore, to some extent, the algorithmic fairness in recommender systems is more related to the *Engagement* aspect of the *UDL* framework where the goal is to make sure that the recommendations are relevant and engaging to all users, regardless of their race, age, gender, etc., and not just a certain group of users.

The non-algorithmic aspect of fairness in recommender systems, on the other hand, relates to *Representation* and *Action & Expression*. While these may play a crucial role on the fairness perception by the users, they have not received deserved attention from a fairness point of view in recommender systems. In this paper, we intend to describe the importance of these aspects and how fairness can indeed go beyond algorithmic fairness.

3.1 Fairness in User Engagement

In recommender systems, *Engagement* aspect focuses on how different users are engaged and motivated to use the system and interact with the recommendations. There are a variety of sources that can influence individual variation in engagement including culture, personal relevance, subjectivity, background knowledge, and even neurology, along with a variety of other factors. Some users are highly engaged when novel recommendations are provided to them while some may become disengaged, even dissatisfied. As mentioned earlier, fairness in engagement is somehow related to the algorithmic fairness in recommendation where the algorithm designer wants to ensure the recommendations are engaging for different users or groups of users (see motivating example 1).

Motivating Example 1: User Engagement
Lisa is mainly interested in niche movies that are not very popular these days. She enjoys watching old Japanese movies and some Polish and French movies. She joins an online movie streaming platform and starts watching some movies. However, after a while she realizes the recommendations are hugely skewed towards mainstream movies and her preferences are not well-represented in the recommendations. His friend, Jack, who is into blockbusters, however, is satisfied with his experience since the recommendations he receives match his preferences. Lisa feels disappointed and disengaged and leaves the platform. ***A fair recommendation should have been able to engage her and provide her recommendations that matched her interest.***

3.2 Fairness in Representation

Representation dimension is about providing multiple means of representation for users to empower different users, who have their own specific physical and cognitive capabilities, to better comprehend the presented recommendations. The representations can include the recommendations, their explanation, as well as the content of the items. Summarizing the key features of the recommended items for users and highlighting them by supportive materials are approaches that can contribute to fairness in representation. A way to realize this dimension is through informing and providing explanation over recommendation generation process, and specifically considering transparency, trustworthiness, and honesty of the system. As a result of a fair representation, the users would perceive that their individual differences, sensory disabilities (e.g., blindness or dyslexia), beliefs, and cultures are taken into account. Consequently, user emotional perception of fairness would be positive toward the system (see motivating example 2).

Motivating Example 2: Representation
Alex is a university student with partial visual impairment who constantly uses the recommendation service offered by a popular online shop to search product and order the desired one. Alex loves checking the reviews and compare the available products. However, he finds it no-easy to read the customer reviews in the current format and needs to personally adjust different properties of the browser to properly figure out what has been written. ***A fair recommender system can address this by offering multiple channels for different groups of users with different physical abilities.***

3.3 Fairness in Action & Expression

This aspect of fairness in recommendation is two-fold: on the one hand, the system should give control to the user in terms of how she wants her recommendations to be changed (see motivating example 3). On the other hand, the system should also provide multiple means of interaction for the users. Users differ in the ways that they can navigate through a recommendation platform and express what they want. For example, individuals with significant movement impairments (e.g., cerebral palsy), those who struggle with strategic and organizational abilities (executive function disorders), those who have language barriers, and so forth often may desire to interact with the recommendation system very differently. Some may be able to express their feedback to the recommendations via pressing a button and some via speech (see motivating example 4).

Motivating Example 3: Action & Expression
Nadi is a vegetarian sportsman who has a strong religious background. Nadi started using a mobile application that generates food recommendations for him, based on his explicit and implicit feedback. Nadi is currently disappointed since he constantly receives recommendation of food recipes containing meat-rich ingredients. Nadi wanted to inform the system to remove meat from his recommendation but couldn't do so. ***A fair recommender system could have given a full control to Nadi over expressing his preferences in a variety of ways allowing him to remove or modify, completely or partially, the data that the system has collected from him.***

Motivating Example 4: Action & Expression
Naghi suffers from Acromegaly. As a result, he often finds it difficult to interact with tiny buttons on a mobile screen. Naghi owns a video blog and regularly checks video recommendations provided for him. However, he continuously encounters uneasiness when expressing his feedback on a given recommendation. ***A fair recommender system in terms of Action & Expression should provide other means of expression such as voice so that people like Naghi can easily interact with the recommendations.***

Figure 1 summarizes these aspects in recommender systems and describes the related problems and solutions within each aspect. Despite the subjective nature of perceived fairness and correspondingly, the need for its assessment via self-reports, one can hardly find any scale for the evaluation of user perceived fairness in interactive systems. In this paper, we have proposed a short survey that can be adopted by academic scholars and industry practitioners in order to evaluate a recommender system in terms of fairness. The survey is presented in Table 1.

It is worth noting that evaluation of perceived fairness has been discussed in a variety of domains, such as organizational justice [5, 14], education [7, 26], legal [28], health [12], etc. Inspired by user perceived fairness from well-established measures of fairness in education and justice systems, we propose an adaptation of user perceived fairness in recommender systems. An overall fair user perception of the recommender systems can be achieved by addressing three dimensions: fairness in user engagement, fairness in representation, and fairness in action & expression.

4 DISCUSSION AND LIMITATIONS

We consider fairness as a procedural and integrative component that has to be applied in various system levels. From user level to algorithms, various aspects of fairness need to be considered so that the resulting system is perceived as fair and bias free. A recommender system perceived as unfair by the user could bring user discomfort, distrust, negative emotions, and reduced engagement resulting even in loss of revenue, reputation, or customer loyalty for the system. Three dimensions of fairness (i.e. fairness in engagement, representation, and action & expression) together provide a primary framework for evaluating user-perceived fairness.

In a typical commercial recommender system, the recommender component functions together with the other elements of the ecosystem. For example, the recommendation interface may only occupy a small portion of the user interface, designed for the users [8]. The overall system may integrate users in a variety of ways using user-centered design approaches which can overlap with user perceived fairness. Nonetheless, user perceived fairness would serve as an emphasizing layer for better user integration in the whole system and in particular, it provides an overall assessment of the perceived fairness in the recommender system component.

Dimensions of user perceived fairness proposed here could potentially have similar or contributing factors to user-centered evaluation of recommender systems. For example, factors such as effectiveness, confidence, trust, perceived novelty, transparency, and overall user satisfaction can be impacted by perceived fairness. It would be however intriguing to determine the relationship between these factors in a future work.

5 CONCLUSION AND FUTURE WORK

Fairness in recommender systems has been of great attention in recent years. In this paper, we extended the fairness in recommender systems from being algorithmic-oriented to be more comprehensive and observing non-algorithmic aspects as well. This is done by proposing a novel evaluation methodology to be used to assess the fairness of a recommender system perceived by the users. We adapted a well-known framework, Universal Design for Learning (UDL), to evaluate the fairness in recommender systems in three different aspects: Engagement, Representation, and Action & Expression. We described how unfairness could arise in any of these aspects. In particular, we formed a short survey including a number of (example) questions that addresses these aspects of fairness. The questions have been carefully formulated in order to cover different aspects of users' needs for fairness when interacting with a recommender system. The proposed survey can be adopted to acquire user feedbacks and aggregate them to obtain an indication of the overall fairness of the system, perceived by users.

As a future work, we plan to build a recommender system and perform a large scale user study to extensively assess the effectiveness of the proposed methodology in order to form a better understanding of the fairness in recommender systems. While, in this paper, we have mainly focused on non-algorithmic aspects of the fairness in recommender systems, however, our user study will go beyond that and create a more comprehensive view which sees the fairness from both algorithmic and non-algorithm perspectives. In addition to that, we will study the differences among users when it comes to their perception of fairness in the recommender system.

ACKNOWLEDGMENTS

This work was supported by industry partners and the Research Council of Norway with funding to MediaFutures: Research Centre for Responsible Media Technology and Innovation, through The Centres for Research-based Innovation scheme, project number 309339.

REFERENCES

[1] Himan Abdollahpouri, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato. 2020. Multistakeholder recommendation: Survey and research directions. *User Modeling and User-Adapted Interaction* 30, 1 (2020).

[2] Himan Abdollahpouri, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. 2019. The unfairness of popularity bias in recommendation. In *RecSys Workshop on Recommendation in Multistakeholder Environments (RMSE)*.

[3] Gediminas Adomavicius, Jesse Bockstedt, Shawn Curley, and Jingjing Zhang. 2014. De-Biasing User Preference Ratings in Recommender Systems. In *Proceedings of the Joint Workshop on Interfaces and Human Decision Making for Recommender Systems co-located with ACM Conference on Recommender Systems (RecSys 2014)*. http://ceur-ws.org/Vol-1253/

[4] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Li Wei, Yi Wu, Lukasz Heldt, Zhe Zhao, Lichan Hong, Ed H Chi, et al. 2019. Fairness in recommendation ranking through pairwise comparisons. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2212–2220.

[5] Joel Brockner, Mary Konovsky, Rochelle Cooper-Schneider, Robert Folger, Christopher Martin, and Robert J. Bies. 1994. Interactive Effects of Procedural Justice and Outcome Negativity on Victims and Survivors of Job Loss. *Academy of Management Journal* 37, 2 (April 1994), 397–409. https://doi.org/10.5465/256835 Publisher: Academy of Management.

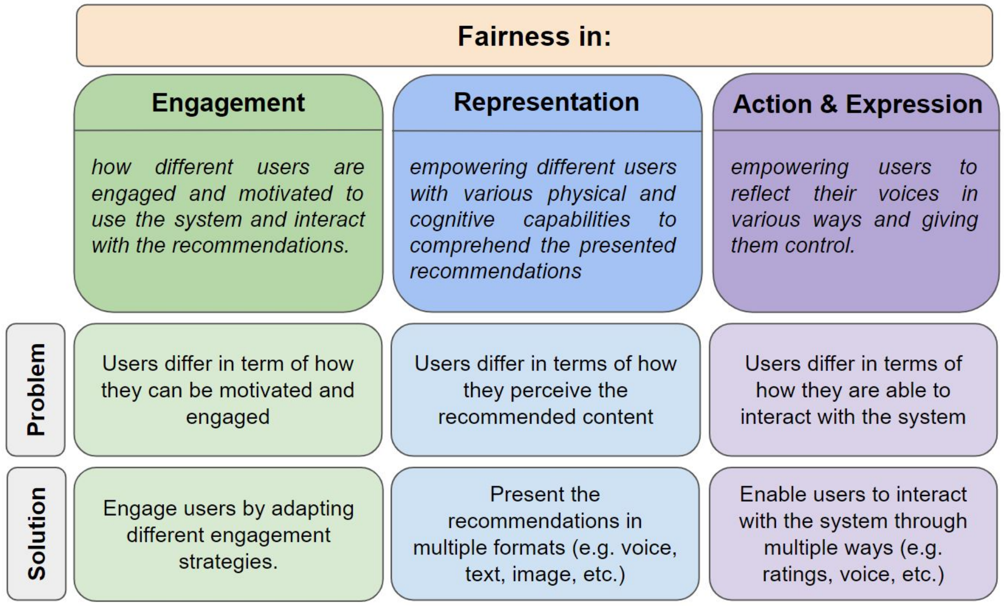

Figure 1: Different dimensions of fairness consideration in recommender systems.

Table 1: Proposed survey, addressing different dimensions of fairness in recommender systems, i.e., fairness in user engagement, fairness in representation, fairness in action & expression. This survey could be used with a five point Likert-type scale.

Dimensions	Questions
Fairness in User engagement	1) Do users feel engaged and motivated when interacting with the system? 2) Does the Recommender System present items description in multiple formats and in an accessible way? 3) Does the Recommender System offer an interaction mechanism where users diversity and individual differences is respected (considering user context, such as channel, modal, social, and physical context) 4) Is Recommender System interface easily approachable and available to users?
Fairness in Representation	5) Does the system respect culture and beliefs of the users? 6) Does the recommender system summarize key features of the items and presented them to users in an accessible, organized, and easy-to-use way? 7) Does the Recommender System explain how the recommendation is generated? 8) Is user informed and aware of fairness considerations by the system?
Fairness in Action and expression	9) Does the system provide the right of being forgotten and reversibility to user's behavior and choices? 10) Is user able to control the fairness considerations incorporated in his/her own recommendations? 11) Does the system provide control to user over decision making process? 12) Does the Recommender System offers opportunities for users to express preferences on the items in various ways?

[6] Abhijnan Chakraborty, Gourab K Patro, Niloy Ganguly, Krishna P Gummadi, and Patrick Loiseau. 2019. Equality of voice: Towards fair representation in crowdsourced top-k recommendations. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 129–138.

[7] Rebecca M. Chory-Assad and Michelle L. Paulsel. 2004. Classroom justice: student aggression and resistance as reactions to perceived unfairness. *Communication Education* 53, 3 (July 2004), 253–273. https://doi.org/10.1080/0363452042000265189 Publisher: Routledge _eprint: https://doi.org/10.1080/0363452042000265189.

[8] Paolo Cremonesi, Mehdi Elahi, and Franca Garzotto. 2017. User interface patterns in recommendation-empowered content intensive multimedia applications. *Multimedia Tools and Applications* 76, 4 (2017), 5275–5309.

[9] Yashar Deldjoo, Vito Walter Anelli, Hamed Zamani, Alejandro Bellogin, and Tommaso Di Noia. 2020. A flexible framework for evaluating user and item fairness in recommender systems. *User Modeling and User-Adapted Interaction* (2020), 1–47.

[10] Michael D. Ekstrand, Mucun Tian, Ion Madrazo Azpiazu, Jennifer D. Ekstrand, Oghenemaro Anuyah, David McNeill, and Maria Soledad Pera. 2018. All The Cool Kids, How Do They Fit In?: Popularity and Demographic Biases in Recommender Evaluation and Effectiveness. In *In Conference on Fairness, Accountability and Transparency*. 172–186.

[11] Michael D Ekstrand, Mucun Tian, Mohammed R Imran Kazi, Hoda Mehrpouyan, and Daniel Kluver. 2018. Exploring author gender in book rating and recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems*.

ACM, 242–250.
[12] Nieke A. Elbers, Alex Collie, Sheilah Hogg-Johnson, Katherine Lippel, Keri Lockwood, and Ian D. Cameron. 2016. Differences in perceived fairness and health outcomes in two injury compensation systems: a comparative study. *BMC Public Health* 16, 1 (July 2016), 658. https://doi.org/10.1186/s12889-016-3331-3
[13] Miriam Evans and Alyssa R Boucher. 2015. Optimizing the power of choice: Supporting student autonomy to foster motivation and engagement in learning. *Mind, Brain, and Education* 9, 2 (2015), 87–91.
[14] Stephen W. Gilliland. 1993. The Perceived Fairness of Selection Systems: An Organizational Justice Perspective. *The Academy of Management Review* 18, 4 (1993), 694–734. https://doi.org/10.2307/258595 Publisher: Academy of Management.
[15] Ben Hutchinson and Margaret Mitchell. 2019. 50 years of test (un) fairness: Lessons for machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 49–58.
[16] Dominik Kowald, Markus Schedl, and Elisabeth Lex. 2020. The unfairness of popularity bias in music recommendation: A reproducibility study. In *European Conference on Information Retrieval*. Springer, 35–42.
[17] Jurek Leonhardt, Avishek Anand, and Megha Khosla. 2018. User fairness in recommender systems. In *Companion Proceedings of the The Web Conference 2018*. 101–102.
[18] Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. Feedback loop and bias amplification in recommender systems. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2145–2148.
[19] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *Proceedings of the 27th acm international conference on information and knowledge management*. 2243–2251.
[20] Alessandro B Melchiorre, Eva Zangerle, and Markus Schedl. 2020. Personality Bias of Music Recommendation Algorithms. In *Fourteenth ACM Conference on Recommender Systems*. 533–538.
[21] Silvia Milano, Mariarosaria Taddeo, and Luciano Floridi. 2020. Recommender systems and their ethical challenges. *AI & SOCIETY* 35, 4 (2020), 957–967.
[22] Dimitris Paraschakis. 2016. Recommender systems from an industrial and ethical perspective. In *Proceedings of the 10th ACM conference on recommender systems*. 463–466.
[23] Gourab K Patro, Arpita Biswas, Niloy Ganguly, Krishna P Gummadi, and Abhijnan Chakraborty. 2020. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proceedings of The Web Conference 2020*. 1194–1204.
[24] Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy. 2020. Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 469–481.
[25] Kavita Rao, Sean J Smith, and K Alisa Lowrey. 2017. UDL and intellectual disability: What do we know and where do we go? *Intellectual and developmental disabilities* 55, 1 (2017), 37–47.
[26] David Rose. 2000. Universal design for learning. *Journal of Special Education Technology* 15, 3 (2000), 45–49.
[27] Jessie Smith, Nasim Sonboli, Casey Fiesler, and Robin Burke. 2020. Exploring User Opinions of Fairness in Recommender Systems. *arXiv preprint arXiv:2003.06461* (2020).
[28] Tom R. Tyler. 1984. The Role of Perceived Injustice in Defendants' Evaluations of Their Courtroom Experience. *Law & Society Review* 18, 1 (1984), 51–74. https://doi.org/10.2307/3053480 Publisher: [Wiley, Law and Society Association].
[29] Kathryn Kavanagh Webb and Jeanne K Hoover. 2015. Universal design for learning (UDL) in the academic library: A methodology for mapping multiple means of representation in library tutorials. (2015).