Diversity-aware Recommendations for Social Justice? Exploring User Diversity and Fairness in Recommender Systems

Laura Schelenz
laura.schelenz@uni-tuebingen.de
International Center for Ethics in the Sciences and Humanities (IZEW)
& University of Tübingen
Germany

ABSTRACT

Diversity and fairness are increasingly linked in the field of personalized recommendations. For instance, the diversification of items ("item diversity") is considered key to fairness. Less attention has been paid to "user diversity" and its implications for fairness. In this paper, I problematize the conceptualization and application of user diversity in recommender systems. I argue that the widespread understanding of user diversity as natural, value-neutral, and individual-level categories may accidentally compound historical injustice. To mitigate emerging biases, diversity dimensions need to be contextualized by mapping structural inequalities between users. The paper thus stresses the importance of paying attention to the *structural context* of diversity, whereas the context refers to political and social circumstances surrounding the user's life. The paper makes three contributions: 1) It connects fairness to diversity literature in the field of recommender system, 2) it specifies the tension between item-side and user-side fairness by revealing a bias in the treatment of user diversity, 3) it proposes solutions to mitigate the bias by drawing on Black feminist and critical race theory.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning algorithms**; • **Social and professional topics** → *User characteristics*; • **Information systems** → **Social recommendation**.

KEYWORDS

diversity, recommender systems, user diversity, fairness, Black feminism

ACM Reference Format:
Laura Schelenz. 2021. Diversity-aware Recommendations for Social Justice? Exploring User Diversity and Fairness in Recommender Systems. In *Adjunct Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '21 Adjunct), June 21–25, 2021, Utrecht, Netherlands.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3450614.3463293

UMAP '21 Adjunct, June 21–25, 2021, Utrecht, Netherlands

ACM ISBN 978-1-4503-8367-7/21/06…$15.00
https://doi.org/10.1145/3450614.3463293

1 INTRODUCTION

Diversity and fairness have received increased attention in the field of personalization and recommendation (especially in the FairUMAP workshops). The two concepts are related. For instance, a diversification of recommended items is considered key to fairness. One example is the recommendation of search results. Presenting a diverse range of sources and representations of race and gender will increase justice towards populations that have previously been rendered invisible [47].

Initially, diversity in recommendations was a question of user satisfaction (we don't want to recommend the "same old" items [43]). Plus, diversity was used as a strategy to "optimize the chances that at least some item pleases the user" [9, p. 883] given the uncertainty about users' *actual* preferences. In this vein, methods such as re-ranking were developed "to achieve a balance between diversity and accuracy." Yet today, diversity is increasingly considered in fairness-related efforts [56] and re-ranking is a tool to right wrongs in a recommender model.

This paper deals with the relationship of diversity and fairness from a social justice point of view (not a user satisfaction point of view). Reviewing literature in the field, I investigate how diversity is leveraged in recommender systems and evaluate the implications for fairness from a Black feminist perspective. The paper finds that, while item diversity may be an effective tool to increase fairness, the way researchers in the field currently leverage "user diversity" compromises these efforts. In particular, I argue that a naive employment of user diversity models may *compound* previous injustices. Concern arises especially from the common understanding of diversity categories (gender, age, education, skills, practices, personality) as neutral, individual-level characteristics.

Black feminist theory helps verbalize these concerns. The theory is well suited to this analysis because it provides the vocabulary and methods to reveal *structural inequalities*, which are obscured by mainstream diversity concepts and rhetoric. Section 2 specifies our method and highlights the theoretical underpinnings of Black feminist theory. In Section 3, I tentatively relate diversity and fairness. In Section 4, I provide background information on the differences between item and user diversity. To elaborate the main argument, Section 5.1 provides a critique of "user diversity." A hypothetical use case in 5.2 better illustrates the mentioned concerns for fairness. Following my critique, I offer tentative solutions for the handling of user diversity concepts. Section 6 proposes to contextualize user diversity by mapping users' experiences of privilege and oppression in a given diversity dimension, and section 7 offers tangible recommendations to system designers.

2 METHODOLOGY AND THEORY: A BLACK FEMINIST APPROACH TO FAIRNESS

For the critical analysis of the relationship between diversity and fairness in recommendations, I first conducted a literature review of how researchers and designers in the field understand "diversity." The papers were selected by a keyword search in the ACM Digital Library and Google Scholar, combining the keywords "diversity," "personalization," and "recommendation." From there, I searched the references of the papers for further material. To ensure that the analysis is up to date, I further included 73 full papers from the ACM conference on User Modelling and Adaptive Technologies (UMAP) 2020 and the ACM conference on Recommender Systems (RecSys) 2019.

Then I reviewed the concept "diversity" from a (US-centric) Black feminist perspective. Black feminism is a critical social theory [11] that has long dealt with questions of diversity from a social justice perspective. Contrary to other theories of social justice, which attend to fairness from a single axis lens, Black feminism advances an *intersectional* perspective. Intersectionality renders visible the existence of double or multiple converging forms of oppression that shape the lived reality of Black women [18]. These intersecting forms of oppression exist not just in the physical world. Algorithmic bias and discrimination disproportionally affects Black women [5, 7, 47].

According to a Black feminist vision of "fairness," societal and technological systems must render visible and dismantle oppression and structural inequalities between social groups [5, 14]. Contrary to common perception, this does not necessarily mean equality of opportunity or equal treatment [17, 1346]. Due to historically grown discrimination, society does not represent a level playing field. Algorithms may compound prior injustice "by carrying it forward into another domain" [33, p.828]. Fairness thus means that a system dismantles oppression *and* prevents the compounding of prior injustice.

3 DIVERSITY AND FAIRNESS

Diversity has both a conceptual and a normative component. Conceptually, diversity refers to the difference of many "things." Normatively, diversity is linked to discussions of multiculturalism and pluralism [64, p. 39] but also calls for inclusion and justice: here, diversity is debated in terms of belonging, equal access to resources or matters of recognition [27].

Fairness refers to the equal treatment of human beings. Rawls 1971 [51] stresses that people should have equal opportunities, while inequalities resulting from different levels of talent and capabilities must be mitigated. However, this account of fairness may not hold in real life because it neglects (historical) discrimination that leads to an asymmetric playing field. New accounts of fairness should thus correct for historical injustice [17, 33].

In the field of recommender systems, diversity and fairness are linked. On the one hand, researchers invoke diversity's *normative* quality by referring to diversity as a moral imperative. Ekstrand, Burke, and Diaz 2019 [22] stress that the legacies of historical discrimination may influence recommendations, which has implications for different groups of users. Definitions of fairness in recommendations thus build on notions of moral diversity, such as inclusion, non-discrimination, and justice. According to Sacharidis 2020 [53], "fairness means that the system exhibits certain desirable ethical traits, such as being non-discriminatory, diversity-aware, and bias-free" [p. 313].

On the other hand, *conceptual* diversity is leveraged to develop fairness-aware methods. The diversification of items (through re-ranking) is a prominent approach to increase fairness in recommender systems [57]. Diversity from a fairness perspective wants to increase the number of different items that are recommended to a user, especially in a way that benefits protected classes as producers of items [22].

Finally, some works consider diversity among users in fairness endeavors [8, 29, 40]. The focus here is on the performance of recommendations for different groups of users. This is an important new avenue. Especially Burke's 2017 [8] "multi-sided" fairness hints to a tension between item-side fairness and user-side fairness, which is further explored in this paper.

4 ITEM DIVERSITY VS. USER DIVERSITY IN RECOMMENDER SYSTEMS

In order to further clarify the tension between fairness on the item side and fairness on the user side, let us first consider the diversity of items and users in detail. Research that *explicitly* deals with diversity in the field of recommender systems mostly refers to a) item diversity and b) personalizing the level of diversity in item recommendations. "Diversity" is considered the dissimilarity between items in an item pool [37]. Castells, Hurley, and Vargas 2015 [9] provide a formal definition: "Diversity generally applies to a set of items and "pieces," and has to do with how different the items or pieces are with respect to each other" [p. 884].

Eskandanian and Mobasher 2020 [24] differentiate between individual diversity and aggregate diversity. Whereas aggregate diversity promotes a wide coverage of different items from an item pool (and thus mitigating a possible popularity bias that favors the recommendation of popular items), individual diversity describes the variability of items recommended to a user. On the individual level, diversity is also adapted to user's individual diversity tolerance (maximum variety vs. feeling overwhelmed [25]).

User diversity is considered rather *implicitly* in recommender systems and refers to information about the user. Since information about users is mostly gathered through implicit feedback (user-item interaction), the interactions of users with items can be diverse. Diversity in user-item interaction can be leveraged to build sub profiles or style profiles of users [36]. However, Burke 2017 [8] as well as Sacharidis, Mukamakuza, and Werthner 2020 [54] raised awareness about the difficulty to predict diverse user preferences based solely on implicit feedback. This difficulty is amplified in cold start scenarios.

Research increasingly deals with the diversity of users beyond implicit feedback. Frolov and Oseledets 2019 [28] propose combining user-item interaction and additional user attributes to improve the quality of recommendations. Costa and Dolog 2019 [16] propose a context-aware recommendation model that focuses on temporal aspects to contextualize user preferences. Dudzik et al. 2020 [21] also stress the importance of contextualizing the user, e.g. observing their individual emotional reactions to videos, which may be triggered by personal experiences and activated memories.

Hence, we see a shift in attention from item diversity or user-item diversity to user diversity. While item diversity as a tool to build recommender systems and increase their fairness still dominates the field, emerging perspectives are interested in "What's in a user?" [44]. This is a welcome shift as recommender systems may become more human-centric. However, some risks emerge for the fairness of a system.

5 HOW "FAIR" RECOMMENDER SYSTEMS RISK COMPOUNDING PREVIOUS DISCRIMINATION

By now, it has been established that technology can reinforce existing structures of inequality (see works on data and algorithmic bias [5, 7, 47]). This section highlights a similar bias, which results from the way that user diversity categories are used in the design of recommender systems.

5.1 User Diversity in Recommender Systems - Challenges

The field of personalized recommendation draws on a series of diversity dimensions that describe types or groups of users. A preliminary mapping of diversity in two ACM conferences (RecSys 2019 and UMAP 2020) reveals 7 dimensions of diversity (see Table 1). Users are considered diverse in their demographics (e.g. age, gender, occupation), psychology (including personality and affect), physiology, culture, skills, social practices, and relationships. An additional area of diversity relates to specific functional preferences or preferences in the realm of user-computer interaction (e.g. how much transparency or control users wish to have over a system [44]). This area, however, was not documented because it appeared less relevant to the research question.

A bias may emerge from the way that these categories are understood and leveraged in the development of recommender systems. In particular, three problems have implications for fairness. First, the diversity categories are treated as self-evident, natural, and value-neutral. Second, diversity categories are employed with little theoretical foundation and reflection of the *social construction* of diversity categories. In other words, what is missing in respective research is a written section that not only specifies the diversity concept but briefly lays out its origins and premises, and reflects possible implications of the diversity concept when applied within a recommender model. Third, the diversity of users is considered on an individual rather than a structural level. These three shortcomings usually occur simultaneously and are connected to each other. A number of examples exist:

- [13] study the effects of culture on users' interaction with picture passwords. The authors make transparent why they believe that "culture" is relevant to the subject but the authors take for granted that culture can be divided into "Eastern vs. Western"; in this case, the readers (and authors) may benefit from further theorizing the meaning of culture and how culture can be operationalized in a system [13, p. 44].
- [23] explore the "diversity" of students and speak of "student-level characteristics" and "differences in prior knowledge" [p. 66] yet it remains unclear what "student characteristics" and "knowledge" entail; how are these concepts operationalized to produce (fair) recommendations?
- [44] explore the relationship of "personal characteristics" and preferences for explanations in music recommender systems. Although the authors define personal characteristics [p. 174], it remains unclear why this particular *combination* of diversity features is more relevant than others
- [3] and [41] optimize services for users on the autism spectrum by drawing on users' (dis)abilities as a diversity dimension; while these efforts are important to increase access and fairness for impaired users, the authors remain vague as to their conceptualization of disability and thus, in an unfortunate scenario, may risk reinforcing static ideas of "disability"; the cause could benefit from further reflection on disability, e.g. a specification whether authors adopt a social or biological/medical model of disability [30] and the implications of such a decision for the recommender model
- Another example is the common use of the Big Five personality model without reflecting its origins, premises, and implications for the particular use case [1, 55, 59]; while the personality model may be popular and widely respected, it may nevertheless require contextualization in a particular use case
- In one example [59], the authors nicely lay out the demographic and psychological user characteristics (age, gender, personality) in their study on the perception of serendipity [p. 270]. While this extensive reflection of diversity dimensions can be considered a best practice case, concerns arise from a lack of awareness of users' differences due to social status. How do societal structural dynamics (such as gender inequalities) shape the experiences and preferences of users? Might financially strained elderly female users have different preferences for serendipity than female users who are comfortably situated? Here, social context matters and even within diversity categories, intersectional experiences may result in divergent preferences

The above examples were chosen to illustrate the difficulty of picking, operationalizing, and justifying diversity concepts in the development of computer models. Although attending to the diversity of users in the first place is commendable, risks emerge from a superficial treatment of diversity. Diversity concepts do not originate in a vacuum. They are socially constructed through discourse and practices [32]. They are embedded in a given *political and social context* that determines what these diversity categories "do." For instance, diversity concepts (e.g. race, gender) can be used to establish hierarchies and a particular social order, where some groups are privileged and some are oppressed [11]. Hanna et al. 2020 [31] have argued that the widespread understanding of race as an individual descriptor rather than a system impedes fairness efforts in machine learning. I pick up this line of argumentation and argue that the way we currently leverage user diversity in recommender systems may impede fairness. The following hypothetical use case better illustrates this argument.

Table 1: Concepts of user diversity in the field of personalized recommendations

Diversity dimensions	Diversity of users	Sources
Demographics	Geographical location, age, gender, occupation	[28, 46, 50, 58, 59, 66]
Psychology	Personality, psychology, cognition, emotions, affect	[1, 39, 55, 59, 60]
Physiology	Physiology, heart rate, stress levels, physiological reactions	[38, 65]
Culture	Sociocultural background, culture, language	[13, 48]
Skills/Experience	Skills, abilities, cognitive abilities, performance of different tasks, mental health, mental capabilities, knowledge, experience	[3, 4, 6, 23, 34, 41, 49, 52, 62]
Social practices	Interests, activities, practices	[35, 58]
Relationships	Social relations, social interactions	[16, 26]

5.2 Hypothetical Use Case: Educational Diversity and Historical Discrimination

A public broadcasting website suggests documentaries to users. The pool of documentaries ranges from sports, nature, and travel to history, science, and politics. It includes diverse producers and featured groups within its pool. Rather than relying solely on user-item interaction, the system takes into account the "real" diversity of users. Diversity is understood as educational level (high school diploma, associate's degree, bachelor's degree, master's degree and doctorate - which is asked when users create an account). The designers of the system assume that education is an indicator of preferences for documentaries: highly educated users are considered more interested in scientific documentaries rather than sports.

In this example, user diversity is explicitly considered in the process of designing a recommender system, contrary to classic collaborative filtering. Whereas the system "only" recommends documentaries, it adds to extra-institutional education for viewers and can inspire them to adopt interests or role models from the documentaries. In the absence of formal educational opportunities for marginalized groups (e.g. due to costs), recommendations for scientific documentaries can become an issue of social justice.

On the item side, diversity/fairness is achieved by including documentaries produced by or featuring diverse groups. On the user side, however, previous race and gender disparities in educational opportunities [20] may result in recommendations of scientific documentaries to mostly privileged groups. The bias has emerged because designers considered user diversity as individual-level categories that are detached from larger social relations. They ignored the fact that members of different groups have different starting points that determine their chances to access the same resources. Because of historical discrimination (e.g. Jim Crow in the United States), marginalized groups (especially Black women) have not gained the same level of education as Whites and males. The legacies of such discrimination materialize today in a smaller number of Black women in science and higher education [10].

Although the scenario itself is set up to incorporate biases (e.g. designers are said to treat education as indicator of preferences), the potential for unfairness emerges primarily from the treatment of diversity dimensions as individual-level characteristics rather than structural ones. The point of the scenario is then to illustrate a particular form of bias, namely the compounding of previous discrimination. The hypothetical use case reveals that fairness goals require closer attention to the way we conceptualize and utilize diversity categories. While diversity as a tool to achieve fairness on the item side may be effective, unfairness on the user side compromises item fairness.

6 CONTEXTUALIZING USER DIVERSITY: LOOKING AT PRIVILEGE AND OPPRESSION

In order to avoid compounding previous injustices in society via recommender systems, user diversity must be leveraged in a way that considers and dismantles *structural* inequalities (cf. Section 2). The first step to achieve this goal is to contextualize user diversity dimensions. Any diversity dimension (gender, race, education, disability, social practices etc.) involves a *structural* context that mainstream diversity discourses often neglect. This context relates to political and social circumstances that determine how different members of society fare in this society. Black feminist and critical race scholars have pointed particularly to structural differences (inequalities) in society, which materialize in people's different experiences of oppression and privilege [2, 11, 17, 18]. Contextualizing user diversity dimensions then means understanding and attending to users' different levels of privilege and oppression in a given political and social context.

Oppression refers to a situation, where "systematically and over a long period of time, one group denies another group access to the resources of society" [11, p.4]. The Black Feminist concept intersectionality highlights the multiple, converging forms of oppression that for instance Black women face [15, 18]: they are affected by gender configurations and race hierarchies, as well as heteronormative and classist value systems.

Privileges are the often unnoticed or taken for granted advantages that a group holds. McIntosh 1988 [42] offers an account of "white privilege" which renders visible the unearned advantages that white people experience, such as being sure that neighbors will be neutral or pleasant to them or that media productions feature people who look like them [p.4]. Oppression and privilege are dynamic. Whether a social group is oppressed depends on the so-called "matrix of domination" [11, p. 227f], which describes the particular organization of oppression in a society. This can be different from one location to another.

In order to reveal the structural context of a diversity dimension, I propose to map users' experiences of privilege and oppression within a diversity category. Wong-Villacres et al. 2018 [61] have conducted such an exercise in a development project using ICT in the Global South. They unpacked the degrees and forms of "penalties" and "privileges" that their beneficiaries experience. First, they reviewed their (ethnographic) data and found common themes or subjects that shape users reality. Then, they produced narrations of the life circumstances and socio-economic dynamics experienced by the users. Finally, they identified points of intervention, where the technology can alleviate penalties.

In the hypothetical use case (see Section 5.2), our diversity dimension was education. Fairness efforts were compromised because the designers picked educational "degree" as a diversity dimension. They did not take into account that the educational system in the USA does not afford everyone the same opportunity to gain a degree of higher education. To mitigate the bias, designers could map the privileges and oppression of groups of users in primary, secondary, and tertiary education by studying critical race and feminist works on (higher) education, e.g. [19, 45]. This contextualization can help reiterate the choice of diversity dimensions and better reflect implications for the recommender model. Based on their new insights, designers may decide to refrain from using "degree" to operationalize educational diversity. They may also become inspired to question the premise that education indicates preferences for scientific documentaries. Finally, they may decide to deliberately leverage recommendations to mitigate the oppression they mapped.

Following Wong-Villacres et al. 2018 [61], there is potential in ethnographic studies of user groups to understand their needs and preferences. Similarly, participatory approaches are appropriate design methods to understand the specific experiences of privilege and oppression of affected user groups. Designers may seek access to so-called identity groups (e.g. refugees, Muslims, transgender women, Latinx) and conduct focus groups with members of these groups or expert interviews with NGO workers representing these groups to understand group members' experiences with the educational system.

Having said that, thinking of users in terms of identity groups has its limits. After all, a Black feminist intersectional lens highlights differences *within* identity groups depending on a person's position at the intersection of multiple oppressive systems. Furthermore, thinking of users as members of identity groups can have an essentializing effect [63]. Nevertheless, from a social justice point of view, identity groups are bound together by certain structural experiences. Precisely these experiences motivate the groups' political struggles for social justice [12]. It thus seems legitimate to start an inquiry with different identity groups and, from there, map diverse experiences of oppression and privilege which may transcend different groups. Yet, the process of contextualization of diversity requires further elaboration in future work.

7 RECOMMENDATIONS: TOWARDS DIVERSITY-AWARE RECOMMENDATIONS

The following recommendations and guiding questions can help designers of recommender systems avoid the compounding of previous injustice when they leverage concepts of user diversity:

1. Making a Concept of User Diversity Explicit

This may seem trivial, but it is essential to spell out what diversity dimensions we take into account, when we consider user context factors in computer models. This means not only stating a diversity category but also explaining which theoretical model we build on, and why the diversity category is relevant.

Guiding questions: What kind of diversity are we talking about? Who and what do we consider diverse? Why do we prefer this diversity concept over others?

2. Reflecting a Concept of User Diversity

Reflecting the diversity concept that lays the ground for categories of user diversity is important to understand what assumptions about human difference we buy into. Diversity concepts have been formulated in a given historical and societal context. Diversity concepts are socially constructed, often with a specific goal in mind, and thus carry baggage.

Guiding questions: What are the origins of the diversity concept? Who has developed it in which social and political context? What are the limitations of the diversity concept?

3. Contextualizing a Concept of User Diversity

Using existing diversity categories can lead to the unintentional reinforcement of injustices because the way these diversity categories are commonly understood obscures *structural* inequalities. It is therefore crucial to consider the different experiences of privilege and oppression of users in a given diversity dimension.

Guiding questions: What are privileges and oppression that users face in a certain diversity dimension? What are the cultural and historical contexts that shape a users current experience in a diversity dimension?

8 CONCLUSION

This paper discussed challenges for user-side fairness and specified a bias emerging from the conceptualization and application of user diversity categories. It highlighted concerns that categories of user diversity are often taken for granted without questioning their origins or meaning in the larger social structure. A hypothetical example involving educational diversity illustrates how recommender systems may accidentally compound historical injustice if they disregard the structural context of diversity dimensions. In order to avoid the identified pitfall, designers can contextualize user diversity by mapping structural differences (inequalities) between users, which materialize as different experiences of privilege and oppression. Preliminary suggestions for the contextualization of user diversity are given in the paper. These are only initial steps to increase fairness with regard to user-side diversity, and further research should provide solutions along a justice-oriented design process.

ACKNOWLEDGMENTS

This research has received funding from the EU Horizon 2020 project "WeNet – The Internet of Us" under grant agreement No

823783. Further funding was received in the course of the project DIGILOG@BW – DIGITALISATION IN DIALOGUE, funded by the Ministerium für Wissenschaft, Forschung und Kunst Baden-Württemberg, Germany, with resources for digitalization efforts of the state of Baden-Württemberg digital@bw. Many thanks to Astrid Franke and Daniel Gatica-Perez for their valuable advice. Further thanks to Michael Lohaus and Avi Segal for their comments on earlier drafts of the paper.

REFERENCES

[1] Patrícia Alves, Pedro Saraiva, João Carneiro, Pedro Campos, Helena Martins, Paulo Novais, and Goreti Marreiros. 2020. Modeling Tourists' Personality in Recommender Systems. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 4–13. https://doi.org/10.1145/3340631.3394843

[2] Sharon K. Anderson and Valerie A. Middleton (Eds.). 2018. *Explorations in Diversity: Examining the Complexities of Privilege, Discrimination, and Oppression* (third edition ed.). Oxford University Press, New York.

[3] Alisha Banskota and Yiu-Kai Ng. 2020. Recommending Video Games to Adults with Autism Spectrum Disorder for Social-Skill Enhancement. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 14–22. https://doi.org/10.1145/3340631.3394867

[4] Michael Barz, Kristin Altmeyer, Sarah Malone, Luisa Lauer, and Daniel Sonntag. 2020. Digital Pen Features Predict Task Difficulty and User Performance of Cognitive Tests. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 23–32. https://doi.org/10.1145/3340631.3394839

[5] Ruha Benjamin. 2019. *Race after Technology : Abolitionist Tools for the New Jim Code.* Polity Press, Cambridge and Medford MA.

[6] Jakim Berndsen, Barry Smyth, and Aonghus Lawlor. 2019. Pace My Race: Recommendations for Marathon Running. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 246–250. https://doi.org/10.1145/3298689.3346991

[7] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. *Proceedings of Machine Learning Research* 81 (2018), 1–15. http://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf

[8] Robin Burke. 2017. Multisided Fairness for Recommendation. https://arxiv.org/pdf/1707.00093

[9] Pablo Castells, Neil J. Hurley, and Saul Vargas. 2015. Novelty and Diversity in Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer US, Boston, MA, 881–918. https://doi.org/10.1007/978-1-4899-7637-6_26

[10] Ian Cheney and Sharon Shattuck. 2020. Picture a Scientist.

[11] Patricia Hill Collins. 2000. *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment* (2nd ed. ed.). Routledge, New York and London.

[12] Combahee River Collective. 1996. A Black Feminist Statement. In *Words of Fire*, Beverly Guy-Sheftall (Ed.). The New Press, New York, 232–240.

[13] Argyris Constantinides, Anna Maria Pietron, Marios Belk, Christos Fidas, Ting Han, and Andreas Pitsillides. 2020. A Cross-cultural Perspective for Personalizing Picture Passwords. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 43–52. https://doi.org/10.1145/3340631.3394859

[14] Sasha Constanza-Chock. 2020. *Design Justice: Community-led Practices to Build the Worlds We Need.* The MIT Press, Cambridge.

[15] Brittney C. Cooper. 2016. Intersectionality. In *The Oxford Handbook of Feminist Theory*, Lisa Jane Disch and M. E. Hawkesworth (Eds.). Oxford University Press, Oxford and New York, 385–406.

[16] Felipe Soares da Costa and Peter Dolog. 2019. Collective Embedding for Neural Context-aware Recommender Systems. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 201–209. https://doi.org/10.1145/3298689.3347028

[17] Kimberlé Crenshaw. 1988. Race, Reform, and Retrenchment: Transformation and Legitimation in Antidiscrimination Law. *Harvard Law Review* 101, 7 (1988), 1331–1387.

[18] Kimberlé Crenshaw. 1995. Mapping the Margins: Intersectionality, Identity Politics, and Violence Against Women of Color. In *Critical Race Theory*, Kimberlé Crenshaw, Neil Gotanda, Gary Peller, and Kendall Thomas (Eds.). New Press and Distributed by W.W. Norton & Co, New York, 357–383.

[19] Kimberlé Crenshaw, Luke Charles Harris, Daniel HoSang Martinez, and George Lipsitz (Eds.). 2019. *Seeing Race Again: Countering Colorblindness Across the Disciplines.* University of California Press, Oakland, California.

[20] Matthew F. Delmont. 2016. *Why Busing Failed: Race, Media, and the National Resistance to School Desegregation.* American Crossroads, Vol. 42. University of California Press, Oakland, California.

[21] Bernd Dudzik, Hayley Hung, Mark Neerincx, and Joost Broekens. 2020. Investigating the Influence of Personal Memories on Video-Induced Emotions. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 53–61. https://doi.org/10.1145/3340631.3394842

[22] Michael D. Ekstrand, Robin Burke, and Fernando Diaz. 2019. Fairness and Discrimination in Retrieval and Recommendation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, Benjamin Piwowarski, Max Chevalier, Eric Gaussier, Yoelle Maarek, Jian-Yun Nie, and Falk Scholer (Eds.). ACM, New York, NY, USA, 1403–1404. https://doi.org/10.1145/3331184.3331380

[23] Andrew Emerson, Michael Geden, Andy Smith, Eric Wiebe, Bradford Mott, Kristy Elizabeth Boyer, and James Lester. 2020. Predictive Student Modeling in Block-Based Programming Environments with Bayesian Hierarchical Models. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 62–70. https://doi.org/10.1145/3340631.3394853

[24] Farzad Eskandanian and Bamshad Mobasher. 2020. Using Stable Matching to Optimize the Balance between Accuracy and Diversity in Recommendation. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 71–79. https://doi.org/10.1145/3340631.3394858

[25] Farzad Eskandanian, Bamshad Mobasher, and Robin Burke. 2017. A Clustering Approach for Personalizing Diversity in Collaborative Recommender Systems. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, Maria Bielikova, Eelco Herder, Federica Cena, and Michel Desmarais (Eds.). ACM, New York, NY, USA, 280–284. https://doi.org/10.1145/3079628.3079699

[26] Wenqi Fan, Yao Ma, Dawei Yin, Jianping Wang, Jiliang Tang, and Qing Li. 2019. Deep Social Collaborative Filtering. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 305–313. https://doi.org/10.1145/3298689.3347011

[27] Nancy Fraser. 1996. Social Justice in the Age of Identity Politics: Redistribution, Recognition, and Participation: Stanford University. *Tanner Lectures on Human Values* (1996). https://tannerlectures.utah.edu/_documents/a-to-z/f/Fraser98.pdf

[28] Evgeny Frolov and Ivan Oseledets. 2019. HybridSVD: When Collaborative Information Is Not Enough. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 331–339. https://doi.org/10.1145/3298689.3347055

[29] 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. http://arxiv.org/pdf/2006.02046v2

[30] Justin Anthony Haegele and Samuel Hodge. 2016. Disability Discourse: Overview and Critiques of the Medical and Social Models. *Quest* 68, 2 (2016), 193–206. https://doi.org/10.1080/00336297.2016.1143849

[31] Alex Hanna, Emily Denton, Andrew Smart, and Jamila Smith-Loud. 2020. Towards a Critical Race Methodology in Algorithmic Fairness. *FAT Conference Proceedings, Barcelona, January 2020* (2020).

[32] Sally Anne Haslanger. 2012. *Resisting Reality: Social Construction and Social Critique.* Oxford University Press, New York.

[33] Deborah Hellman. 2020. Measuring Algorithmic Fairness. *Virginia Law Review* 106, 4 (2020), 811–866.

[34] Nico Herbig, Tim Düwel, Mossad Helali, Lea Eckhart, Patrick Schuck, Subhabrata Choudhury, and Antonio Krüger. 2020. Investigating Multi-Modal Measures for Cognitive Load Detection in E-Learning. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 88–97. https://doi.org/10.1145/3340631.3394861

[35] Daniel Herzog and Wolfgang Wörndl. 2019. User-centered Evaluation of Strategies for Recommending Sequences of Points of Interest to Groups. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 96–100. https://doi.org/10.1145/3298689.3346988

[36] Murium Iqbal, Kamelia Aryafar, and Timothy Anderton. 2019. Style Conditioned Recommendations. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 128–136. https://doi.org/10.1145/3298689.3347007

[37] Marius Kaminskas and Derek Bridge. 2017. Diversity, Serendipity, Novelty, and Coverage: A Survey and Empirical Analysis of Beyond-Accuracy Objectives in Recommender Systems. *ACM Transactions on Interactive Intelligent Systems* 7, 1 (2017), 1–42. https://doi.org/10.1145/2926720

[38] Thomas Kundinger and Andreas Riener. 2020. The Potential of Wrist-Worn Wearables for Driver Drowsiness Detection: A Feasibility Analysis. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 117–125. https://doi.org/10.1145/3340631.3394852

[39] Monica Landoni, Maria Soledad Pera, Emiliana Murgia, and Theo Huibers. 2020. Inside Out: Exploring the Emotional Side of Search Engines in the Classroom. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 136–144. https://doi.org/10.1145/3340631.3394847

[40] Jurek Leonhardt, Avishek Anand, and Megha Khosla. 2018. User Fairness in Recommender Systems. In *Companion Proceedings of the The Web Conference 2018 (ACM Other conferences)*, Pierre-Antoine Champin, Fabien Gandon, Mounia Lalmas, and Panagiotis G. Ipeirotis (Eds.). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, 101–102. https://doi.org/10.1145/3184558.3186949

[41] Noemi Mauro, Liliana Ardissono, and Federica Cena. 2020. Personalized Recommendation of PoIs to People with Autism. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 163–172. https://doi.org/10.1145/3340631.3394845

[42] Peggy McIntosh. 1988. White Privilege and Male Privilege: A Personal Account of Coming to See Correspondences through Work in Women's Studies: Working Paper 189. https://www.nationalseedproject.org/images/documents/White_Privilege_and_Male_Privilege_Personal_Account-Peggy_McIntosh.pdf

[43] Sean M. McNee, John Riedl, and Joseph A. Konstan. 2006. Being Accurate Is Not Enough. In *CHI '06 Extended Abstracts on Human Factors in Computing Systems - CHI EA '06*, Gary Olson and Robin Jeffries (Eds.). ACM Press, New York, New York, USA, 1097. https://doi.org/10.1145/1125451.1125659

[44] Martijn Millecamp, Nyi Nyi Htun, Cristina Conati, and Katrien Verbert. 2020. What's in a User? Towards Personalising Transparency for Music Recommender Interfaces. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 173–182. https://doi.org/10.1145/3340631.3394844

[45] Kathryn Pauly Morgan. 2019. Describing the Emperor's New Clothes: Three Myths of Educational (In-)Equity. In *The Gender Question in Education*. Routledge, Taylor et Francis Group, New York.

[46] James Neve and Ivan Palomares. 2019. Latent Factor Models and Aggregation Operators for Collaborative Filtering in Reciprocal Recommender Systems. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 219–227. https://doi.org/10.1145/3298689.3347026

[47] Safiya Umoja Noble. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism.* New York University Press, New York.

[48] Alessio Palmero Aprosio, Stefano Menini, and Sara Tonelli. 2020. Adaptive Complex Word Identification through False Friend Detection. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 192–200. https://doi.org/10.1145/3340631.3394857

[49] Zacharoula Papamitsiou, Kshitij Sharma, and Michail N. Giannakos. 2020. On the Dependence Structure Between Learners' Response-time and Knowledge Mastery: If Not Linear, Then What?. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 201–210. https://doi.org/10.1145/3340631.3394865

[50] Changhua Pei, Yi Zhang, Yongfeng Zhang, Fei Sun, Xiao Lin, Hanxiao Sun, Jian Wu, Peng Jiang, Junfeng Ge, Wenwu Ou, and Dan Pei. 2019. Personalized Re-ranking for Recommendation. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 3–11. https://doi.org/10.1145/3298689.3347000

[51] John Rawls. 1971. *A Theory of Justice: Original Edition.* Harvard University Press, Cambridge, MA. https://doi.org/10.2307/j.ctvjf9z6v

[52] Soyoung Rhim, Uichin Lee, and Kyungsik Han. 2020. Tracking and Modeling Subjective Well-Being Using Smartphone-Based Digital Phenotype. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 211–220. https://doi.org/10.1145/3340631.3394855

[53] Dimitris Sacharidis. 2020. Building User Trust in Recommendations via Fairness and Explanations. In *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 313–314. https://doi.org/10.1145/3386392.3399995

[54] Dimitris Sacharidis, Carine Pierrette Mukamakuza, and Hannes Werthner. 2020. Fairness and Diversity in Social-Based Recommender Systems. In *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 83–88. https://doi.org/10.1145/3386392.3397603

[55] Mete Sertkan, Julia Neidhardt, and Hannes Werthner. 2020. Eliciting Touristic Profiles: A User Study on Picture Collections. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 230–238. https://doi.org/10.1145/3340631.3394868

[56] Nasim Sonboli, Robin Burke, Zijun Liu, and Masoud Mansoury. 2020. Fairness-aware Recommendation with librec-auto. In *Fourteenth ACM Conference on Recommender Systems.* ACM, New York, NY, USA, 594–596. https://doi.org/10.1145/3383313.3411525

[57] Nasim Sonboli, Farzad Eskandanian, Robin Burke, Weiwen Liu, and Bamshad Mobasher. 2020. Opportunistic Multi-aspect Fairness through Personalized Re-ranking. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 239–247. https://doi.org/10.1145/3340631.3394846

[58] Yuan Tian, Ke Zhou, Mounia Lalmas, Yiqun Liu, and Dan Pelleg. 2020. Cohort Modeling Based App Category Usage Prediction. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 248–256. https://doi.org/10.1145/3340631.3394849

[59] Ningxia Wang, Li Chen, and Yonghua Yang. 2020. The Impacts of Item Features and User Characteristics on Users' Perceived Serendipity of Recommendations. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 266–274. https://doi.org/10.1145/3340631.3394863

[60] Amy A. Winecoff, Florin Brasoveanu, Bryce Casavant, Pearce Washabaugh, and Matthew Graham. 2019. Users in the Loop: A Psychologically-informed Approach to Similar Item Retrieval. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 52–59. https://doi.org/10.1145/3298689.3347047

[61] Marisol Wong-Villacres, Arkadeep Kumar, Aditya Vishwanath, Naveena Karusala, Betsy DiSalvo, and Neha Kumar. 2018. Designing for Intersections. In *Proceedings of the 2018 on Designing Interactive Systems Conference 2018 - DIS '18*, Ilpo Koskinen, Youn-kyung Lim, Teresa Cerratto-Pargman, Kenny Chow, and William Odom (Eds.). ACM Press, New York, New York, USA, 45–58. https://doi.org/10.1145/3196709.3196794

[62] Meng Wu, Ying Zhu, Qilian Yu, Bhargav Rajendra, Yunqi Zhao, Navid Aghdaie, and Kazi A. Zaman. 2019. A Recommender System for Heterogeneous and Time Sensitive Environment. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 210–218. https://doi.org/10.1145/3298689.3347039

[63] Iris Marion Young. 1990. *Justice and the Politics of Difference.* Princeton University Press, Princeton, N.J.

[64] Nira Yuval-Davis. 2011. *Gender & Nation* (reprinted. ed.). Sage, London.

[65] Sebastian Zepf, Neska El Haouij, Jinmo Lee, Asma Ghandeharioun, Javier Hernandez, and Rosalind W. Picard. 2020. Studying Personalized Just-in-time Auditory Breathing Guides and Potential Safety Implications during Simulated Driving. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Tsvi Kuflik, Ilaria Torre, Robin Burke, and Cristina Gena (Eds.). ACM, New York, NY, USA, 275–283. https://doi.org/10.1145/3340631.3394854

[66] Zhe Zhao, Lichan Hong, Li Wei, Jilin Chen, Aniruddh Nath, Shawn Andrews, Aditee Kumthekar, Maheswaran Sathiamoorthy, Xinyang Yi, and Ed Chi. 2019. Recommending What Video to Watch Next: A Multitask Ranking System. In *13th ACM Conference on Recommender Systems*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). The Association for Computing Machinery, New York, New York, 43–51. https://doi.org/10.1145/3298689.3346997