Digitally Dysfunctional: How HCI Researchers Examine BIPOC Interactions with Recommendation Algorithms

Jiarun Dai[1,6], Sehba Wani[6], Zauryn Robinson[3,6], Naila Hajiyeva[4,6], and Morgan Klaus Scheuerman[5]

[1] University of Wisconsin-Madison
[2] University of Maryland, College Park
[3] Spelman College
[4] University of Texas, Austin
[5] University of Colorado Boulder
[6] iSchool Inclusion Institute

Abstract. Recommendation algorithms are often used in many mainstream media apps to provide content suggestions to users based on factors like individual preferences and virality. As a result, social media platforms circulate and popularize content trends. However, like other algorithmic approaches, recommender systems have been found to reinforce biases, particularly towards marginalized communities. This poster presents a preliminary literature review of how the Human-Computer Interaction (HCI) community has approached scholarship on both Black, Indigenous, and People of Color (BIPOC) individuals and recommendation algorithms. Our findings showcase a lack of focus at the intersection of BIPOC experiences and recommendation algorithms, with the technology and the population often being studied separately. Based on the gaps we identify in the literature, we propose future work aimed at understanding how BIPOC individuals experience recommendation algorithms, specifically on TikTok. Without proper recognition of the biases, lack of representation, and understanding of the offline impacts of the information that recommendation algorithms can create, certain populations could be at risk of encountering material that is untrue and harmful to their well-being.

Keywords: BIPOC, recommendation algorithms, human-computer interaction (HCI), bias, personalization, recommender systems, social responsibility, sociotechnical systems

1 Introduction

Recommendation algorithms, also often referred to as recommender systems, are filtering systems used to predict a user's preferences and suggest relevant content, usually to keep them engaged [1, 2]. Recommendation algorithms are a type of algorithm focused on enhancing the user experience, with typically two methods of selecting content to present personalized content to the users' system: collaborative and content-based filtering [3,4]. The two methods differ in the type of information needed to run the algorithm [4]. Such algorithms are now commonly used in a variety of digital platforms, including news websites, and social media. Applications such as TikTok are known for "time-sensitive customer feedback," such as short-video ranking or online ads, with the goal of "deliver[ing] personalized content for each individual user as a real-time response" [5].

Yet, like many other algorithmic techniques, recommendation algorithms can have concerning repercussions. Increasingly, HCI researchers have focused on how recommendation algorithms actively exclude and even target marginalized groups. For example, Andalibi et al. found that algorithmic recommendations systems perpetuate allocational harms when an automated system distributes resources (e.g., money, credit) or opportunities (e.g., jobs) unfairly to different social groups, like women, LGBTQ people, and racial minorities [6, 7, 14, 15]. Identifying biases in algorithmic systems is crucial, else we risk reinforcing social exclusion and problematic stigmas.

In an age where discussions of systematic racism, cultural appropriation, and racial representations are constantly occurring, we sought to explore how HCI research is currently engaging with how recommendation systems intersect with BIPOC (Black, Indigenous, People of Color) identities.

In this paper, we specifically examine research trends at the intersection of BIPOC identity and recommendation algorithms to understand current HCI research. We analyzed literature on both this population and technology to understand whether and how they might overlap. Specifically, we aimed to address the following questions in our literature review:

1. What have HCI scholars focused on in their study of BIPOC individuals?
2. What have HCI scholars focused on in their study of recommendation algorithms?
3. What have HCI scholars identified as important at the intersection of both BIPOC identity and recommendation algorithms?

We found many HCI scholars studying BIPOC identities in various technical domains. Meanwhile, scholarship on recommendation algorithms is primarily technical, rather than social. Generally, HCI scholars do not emphasize the intersection of recommendation algorithms and BIPOC individuals. Instead, they focus on overall algorithmic stigmatization that results in "institutional exclusions," such as notions of racism and harassment [8]. There is not as much of a focus on how recommendation algorithms amplify systemic issues directly into the personalized content of BIPOC users.

Our findings reveal a significant lack of research on the experiences and effects of recommender systems on BIPOC experiences. Given the gaps we identified, we propose future work aimed at providing a fundamental understanding of BIPOC user experiences with recommendation algorithms. Given we did not see any work focused on the impact of recommendation algorithms on the overall well-being of BIPOC individuals, we propose a future study for attending to the BIPOC experience of recommendation algorithms. Such future work can fill current gaps in understanding the effects of potential assumptions, and stereotypes about BIPOC identities embedded in recommender systems. Through conducting this work, we aim to contribute new insights to the ongoing discourse surrounding the societal and ethical impacts of algorithms on BIPOC communities. Our proposed future research aims to illuminate and rectify the embedded biases within these systems, ensuring a more equitable technological landscape for all.

2 Method

To conduct our literature review, we reviewed papers from the Association of Computing Machinery (ACM) Digital Library (DL). We chose to review papers in the ACM DL because many premiere HCI and algorithmic fairness conferences are sponsored by the ACM. Within the ACM DL, we examined peer-reviewed articles, eliminating opinion articles. We used the following search terms to identify relevant literature: “BIPOC,” “recommendation algorithms,” and “BIPOC recommendation algorithms.” We identified 20 papers for “BIPOC recommendation algorithm,” 134 papers for “BIPOC,” and 50 papers for “recommendation algorithm.” Our final corpus included 203 papers. About eighty percent of our corpus were North American contributors.

We categorized the 203 papers to understand the topics of the papers. We identified five categories: diversity, equity, and inclusion (DEI)/representation; education; social; entertainment; technical. Each category represents a domain or concept that technology intersects with in the papers reviewed.

To finalize which papers we would closely analyze, we examined each paper to see if it included the terms “BIPOC” or “recommendation algorithm.” Our final analysis included 118 papers.

Within our final analysis, we centered our attention on methodologies and conclusions. For data collection, we looked at what tools and approaches were utilized. For conclusions we noted whether the authors mentioned any shortcomings of their results towards fairness and bias.

3 Results of the Literature Review

3.1 BIPOC Identity in HCI

We found limited depth in the current HCI research focusing on BIPOC identities. The dominant themes within the ACM DL centered on inclusion and representation within educational settings. In particular, numerous works were focused on increasing representation in computing-related courses [9, 31]. These works often offered insights into the predominant theories and frameworks employed [27, 28, 29, 30].
For example, Xie et al. examined how the increasing number of students in computing courses lead BIPOC students to feel discouraged and disadvantaged [9]. These research findings reflect real-world scenarios, like the Tech Leavers Study which reported about 40% of underrepresented populations leave the tech industry due discrimination [10]. Researchers focused on how a lack of diversity and inclusion impacts a student’s identity in computing spaces [9, 29, 30].

We also identified literature which focused on the intersection of BIPOC identity and technical representations. For example, Bennett et al. showcased the complexity of providing textual descriptions for visual content, particularly images of people, in a way that is both accessible and sensitive to issues of representation. The paper presents findings from interviews with screen reader users who are Black, Indigenous, People of Color (BIPOC), non-binary, and/or transgender, focusing on their practices and preferences for image descriptions. While textual descriptions are crucial for accessibility, guidelines often lack specifics on how to describe people's appearances without misrepresenting them. The authors argue for a nuanced approach that considers the salience of appearance in different contexts and the preferences of the individuals depicted [25].

Similarly, Abebe et al. examined racism in technology design approaches, specifically in HCI. Their study aimed "to interrogate present methods and suggest alternative approaches to HCI, grounded in critical inquiry, seeking abolitionist opportunities for the mutual and ongoing entanglement of people and machinery" [28]. With the focus of presenting "anti-racism" within technology, the authors proposed four prototypes to address anti-racist HCI beyond traditional approaches such as DEI initiatives.

We did see HCI scholarship focused on issues of fairness in machine learning, more broadly. Simpson and Semaan, who conducted a study on the creative work on TikTok, discussed how the "infrastructure" of the platform resulted in recommending creators' work to a low variety and number of users [12]. In these cases, BIPOC users experience harms related to representation and resource allocation, discrepancies between user' goals and the systems objectives, and reduced influence over the system [11, 24]. Without re-evaluating algorithmic design approaches the "primary causes of algorithmic hate are the unintended results of decisions made while designing these systems" such as suppression of certain content, especially content made by BIPOC creators [11, 29].

The lack of diverse perspectives in algorithmic design and implementation is causing tangible harm to BIPOC groups [22]. Measures of how BIPOC identity is accounted for in machine learning are crucial. As Ashktorab et al. identified, how identity is grouped when measuring fairness outcomes is a difficult socially-driven problem that can result in diverse outcomes for machine learning models [13].

We identified two major themes of research specific to BIPOC identity in HCI: issues of representation of BIPOC people in the computing field and issues of fairness in machine learning systems, broadly, not specific to BIPOC populations. Our literature review highlights significant gaps in the topics which include BIPOC individuals and the centrality of BIPOC identity to research, emphasizing a need for expanded HCI research on BIPOC populations and issues.

3.2 Recommendation Algorithms in HCI

Recommender systems often use historical user data to predict what content or information the user would like to be engaged with in the future, allowing for convenience and overall better experience [11, 16]. The examination of recommendation algorithms within HCI research presented a shift away from users towards model techniques and model evaluation. Evaluating recommendation algorithms is a priority within the field of recommender systems. Roy and Dutta conducted a meticulous examination of diverse research papers focused on recommender systems in the entertainment domain, published between 2011 and 2021 [19]. They analyzed different application fields, techniques, simulation tools, performance metrics, and the challenges encountered within different recommender systems. Roy and Dutta's work underscores a pervasive use of Java and Python for developing recommender systems, attributed to the extensive standard libraries in these languages. Roy and Dutta also noted a rising trend in the adoption of hybrid and optimization techniques, enhancing the performance efficiency of recommender systems. Despite the extensive research in the domain, they highlight a significant gap: the absence of a standard measure for evaluating the performance of recommender systems.

Researchers have in the past proposed a test suite that can be used to validate the correctness of a recommendation algorithm to correct the performance and behavior of the algorithms [17, 31]. Although many technical researchers are aware of the need for assessing real-world impact, there remains "the relentless focus on 'accurate' recommender systems ignores important factors that contribute to the trustworthiness of recommender systems, such as transparency, fairness, and robustness to noise and attack" [17].

We identified several papers focused on algorithmic bias. Divisions in how content is delivered to certain group has led to issues of polarization, isolation, and division [18]. Being introduced to only certain media, both harmful and non-harmful, changes one's experience with the applications that they use, leading to highly specific "algorithmic imaginaries," mental models of how algorithms must work, and, also, "algorithmic irritation," negative perceptions of how algorithms work based on personal mental models [11, 32, 33].

Overall, we found that the literature on recommender systems is broad but predominantly focused on the technical aspects of implementation and methodology. We saw a significant lack of research on ethical considerations and the diverse impacts on various communities, including BIPOC individuals.

3.3 BIPOC Identity and Recommendation Algorithms

When examining literature focused on *both* BIPOC identities and recommendation algorithms, we found that there are few researchers studying recommendation algorithms from a human-centered perspective. McConvey et al. noted that AI software "includes many examples of algorithmic harm in higher education including facial-recognition software for exam proctoring providing 'allegedly discriminatory experiences for BIPOC students'" [20]. Cheng et al. noted that the impact of "socially indifferent algorithms" mirror society, with the most affected groups being BIPOC and women [21].

Das et al. examines the Quora Recommendation Algorithm (QRA) on the Bengali Quora platform (BnQuora), specifically expressions of communities including persons within the BIPOC community [26]. They explore sociotechnical mechanisms of governance within the BnQuora, discovering that a large proportion of user experience mediated by BnQuora leads to the privileging of certain identities especially concerning religious background, and linguistic practices. The study emphasizes how the QRA initially seeks to learn about user experience yet still is amplifying distinctions amongst users and recommending controversial conversations that uplift colonial structures and reinforce harmful ideologies [26].

Abebe et al. advocated for understanding the impact of these algorithms, highlighting concerns related to mental health, societal assumptions, and the perpetuation of stereotypes [22]. Echoing these insights, they illuminated the ongoing and adaptable nature of anti-racist actions within HCI.

An emphasis on the dynamic nature of anti-racist efforts highlights the pressing need for continuous engagement with how recommendation systems impact BIPOC individuals. Such an approach is crucial for ensuring that the interactions of BIPOC communities with recommendation algorithms are characterized by fairness. Despite the substantial focus on technical and empirical aspects, the ethical dimensions, especially concerning BIPOC communities, remain largely unexplored.

4 Future Work

We found that many HCI researchers are concerned about BIPOC user experiences; however, scholarship focused on recommendation algorithms are largely only about technical implementation. For the few researchers who have focused acutely on BIPOC identity in recommendation algorithms, we found researchers advocating for applying a human-centered lens to develop less harmful algorithms.

We see vast opportunities to explore specific research that centers BIPOC individuals and their experiences with recommendation algorithms. There should be an emphasis that recommendation algorithms are more than just algorithms but rather sociotechnical systems that humans interact with [23]. We did not identify HCI focused on understanding how BIPOC individuals experience recommendation systems in their daily lives, including the potential benefits and harms of those systems. Understanding the benefits and harms caused by recommendation algorithms is a crucial initial step towards developing fairer and just sociotechnical futures [6].

In our future work, we plan to dig deeper into how recommendation algorithms might perpetuate stereotypes, and misinformation that actively impact BIPOC communities. TikTok, a highly popular social media platform with a robust recommendation algorithm, represents a timely platform to focus on given its mainstream popularity and accessibility to BIPOC users. Thus, our next steps involve conducting empirical research on how BIPOC individuals interact with TikTok. Our exploration will augment collective wisdom about the real-world impact of recommendation systems on BIPOC communities, offering insights to enhance their well-being. Aiming to eradicate algorithmic biases, our future work aspires to promote equality and inclusivity in technological advancements for a well-rounded societal upliftment.

Acknowledgments

We would like to sincerely recognize and thank the iSchool Inclusion Institute, which is generously funded by The Andrew W. Mellon Foundation, for their feedback on this work. Without their aid and insight, we would not have been able to come this far and be empowered to move forward.

References

1. Liang, C.A., Munson, S.A., Kientz, J.A.: Embracing Four Tensions in Human-Computer Interaction Research with Marginalized People. ACM Trans. Comput.-Hum. Interact. 28, 14:1-14:47 (2021). https://doi.org/10.1145/3443686, accessed last on 2024/16/02
2. Lee Y A., Mieczkowski H, Ellison B N., and Hancock T. J.: The Algorithmic Crystal: Conceptualizing the Self through Algorithmic Personalization on TikTok | Proceedings of the ACM on Human-Computer Interaction, https://dl.acm.org/doi/10.1145/3555601, last accessed 2023/10/02.
3. V. Thannimalai and L. Zhang. A Content Based and Collaborative Filtering Recommender System in IEEE Conference Publication in IEEE Xplore, https://ieeexplore.ieee.org/document/9737238, last accessed 2023/10/02.
4. Applied Sciences | Free Full-Text | An Improved Collaborative Filtering Recommendation Algorithm Based on Retroactive Inhibition Theory, https://www.mdpi.-com/2076-3417/11/2/843, last accessed 2023/10/02.
5. Liu, Z., Zou, L., Zou, X., Wang, C., Zhang, B., Tang, D., Zhu, B., Zhu, Y., Wu, P., Wang, K., Cheng, Y.: Monolith: Real Time Recommendation System With Collisionless EmbeddingTable,http://arxiv.org/abs/2209.07663,(2022). https://doi.org/10.48550/arXiv.2209.07663.
6. Andalibi, N., Pyle, C., Barta, K., Xian, L., Jacobs, A.Z., Ackerman, M.S.: Conceptualizing Algorithmic Stigmatization. In: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. pp. 1–18. Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/10.1145/3544548.3580970.
7. Amazon Scraps Secret AI Recruiting Engine that Showed Biases Against Women - Machine Learning - CMU - Carnegie Mellon University, https://www.ml.c-mu.edu/news/news-archive/2016-2020/2018/october/amazon-scraps-secret-artificial-intelligence-recruiting-engine-that-showed-biases-against-women.html, last accessed 2023/10/02.
8. Shelby, R., Rismani, S., Henne, K., Moon, Aj., Rostamzadeh, N., Nicholas, P., Yilla, N., Gallegos, J., Smart, A., Garcia, E., Virk, G.: Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction, http://arxiv.org/abs/2210.05791, (2023). https://doi.org/10.48550/arXiv.2210.05791.last accessed 2024/16/02.
9. Xie B., Oleson A., Everson J., and Ko J. A.: Surfacing Equity Issues in Large Computing Courses with Peer-Ranked, Demographically-Labeled Student Feedback | Proceedings of the ACM on Human-Computer Interaction, https://dl.acm.org/doi/10.1145/3512912, last accessed 2023/10/02.
10. ScottA.,KleinKaporF.,.pdf,https://www.kaporcenter.org/wpcontent/uploads/2017/08/TechLeavers2017.pdf.
11. Recommender Systems and Algorithmic Hate | Proceedings of the 16th ACM Conference on Recommender Systems, https://dl.acm.org/doi/10.1145/3523227.3551480, last accessed 2023/10/02.
12. Simpson, E., and Semaan B. Rethinking Creative Labor: A Sociotechnical Examination of Creativity & Creative Work on TikTok | Proceedings of the 2023 CHI

Conference on Human Factors in Computing Systems, https://dl.acm.org/doi/ 10.1145/3544548.3580649, last accessed 2023/10/02.
13. Ashktorab, Z., Hoover, B., Agarwal, M., Dugan, C., Geyer, W., Yang, H.B., Yurochkin, M.: Fairness Evaluation in Text Classification: Machine Learning Practitioner Perspectives of Individual and Group Fairness. In: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. pp. 1–20. ACM, Hamburg Germany (2023). https://doi.org/10.1145/3544548.3581227 accessed last on 2023/10/02
14. Overview‹GenderShades—MITMediaLab, https://www.media.mit.edu/projects/ gender-shades/overview/, last accessed 2023/10/02.
15. NIST.https://www.nist.gov/news-events/news/2019/12/nist-study-evaluates-effects-race-age-sex-face-recognition-software-GoogleSearch, https://www.google.-com/search?q=22.24.+https%3A%2F%2Fwww.nist.gov%2Fnews-events %2Fnews%2F2019%2F12%2Fnist-study-evaluates-effects-race-age-sex-face-recognition-software&rlz=1C1CHBF_enUS1060US1060&oq=22.24.%09https %3A%2F%2Fwww.nist.gov%2Fnews-events%2Fnews%2F2019%2F12%2Fnist-study-evaluates-effects-race-age-sex-face-recognition-software&gs_lcrp=EgZ-jaHJvbWUyBggAEEUYOTIGCAEQRRg60gEHNDczajBqNKgCALACAA&sourceid=chrome&ie=UTF-8&bshm=rime/1, last accessed 2023/10/02.
16. Xu, Z., Jiang, S.: Study on Personalized Recommendation Algorithm of Online Educational Resources Based on Knowledge Association. Computational Intelligence and Neuroscience. 2022, e2192459 (2022). https://doi.org/ 10.1155/2022/2192459. last accessed on 2024/02/16
17. Michiels L., Verachtert R., Ferraro A.,Falk K., and Goethals B. A Framework and Toolkit for Testing the Correctness of Recommendation Algorithms | ACM Transactions on Recommender Systems, https://dl.acm.org/doi/10.1145/3591109, last accessed 2023/10/02.
18. Link recommendation algorithms and dynamics of polarization in online social networks | PNAS, https://www.pnas.org/doi/full/10.1073/pnas.2102141118, last accessed 2023/10/02.
19. A systematic review and research perspective on recommender systems | Journal of Big Data | Full Text, https://journalofbigdata.springeropen.com/articles/ 10.1186/s40537-022-00592-5, last accessed 2023/10/02.
20. McConvey, K., Guha, S., Kuzminykh, A.: A Human-Centered Review of Algorithms in Decision-Making in Higher Education. In: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. pp. 1–15. Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/ 10.1145/3544548.3580658.
21. Cheng, L., Varshney, K.R., Liu, H.: Socially Responsible AI Algorithms: Issues, Purposes, and Challenges. J. Artif. Int. Res. 71, 1137–1181 (2021). https:// doi.org/10.1613/jair.1.12814.
22. Abebe V., Amaryan G., Beshai M., Ilene, Gurgen Ekin A., Ho W., Hylton R N., Kim D., Lee C., Lewandowski C., Miller T K., Moore A L., Sylwester R., Thai E., Tucker N F., Webb T., Zhao D., Zhao Charles H., and Vertesi J.: Anti-Racist HCI: notes on an emerging critical technical practice. In Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems, https:// dl.acm.org/doi/10.1145/3491101.3516382, last accessed 2023/10/02.

23. Jannach, D., P. Pu, F. Ricci, and M. Zanker.: Recommender systems: Trends and frontiers. AI Magazine 43: 145–50. https://doi.org/10.1002/aaai.12050, last accessed on 2024/16/02.
24. Karizat, Dan Delmonaco, Motahhare Eslami, and Nazanin Andalibi. 2021. Algorithmic folk theories and identity: How TikTok users co-produce Knowledge of identity and engage in algorithmic resistance. Proceedings of the ACM on Human-Computer Interaction 5, CSCW2 (2021), 1–44, last accessed 2024/16/02.
25. Bennett L. C., Gleason C, Scheuerman Klaus M., Bigham P. J., Guo A., and Tao A. 2021. “It’s Complicated”: Negotiating Accessibility and (Mis)Representation in Image Descriptions of Race, Gender, and Disability. https://doi.org/10.1145/3411764.3445498, last accessed 2024/16/02
26. Das D., Østerlund C., and Semaan B.: “Jol” or “Pani”?: How Does Governance Shape a Platform’s Identity?. Proc. ACM Hum.-Comput. Interact. 5, CSCW2, Article 473 (October 2021), 25 pages. https://doi.org/10.1145/3479860, last accessed on 2024/16/02
27. Guzdial M. :Talking about race in CS education. Commun. ACM 64, 1 (January 2021), 10–11. https://doi.org/10.1145/3433921, last accessed on 2024/16/02
28. Reeves D. C. and Richard T. G.: This is So Cool: Exploring Identity Development and Roles for Middle School Girls of Color Within a STE(A)M-focused Summer Camp. In Proceedings of the 22nd Annual ACM Interaction Design and Children Conference (IDC '23). Association for Computing Machinery, New York, NY, USA, 578–582. https://doi.org/10.1145/3585088.3593861, accessed last on 2024/16/02.
29. Lee L., Latulipe C., and Frevert T.: Using Discrimination Response Ideation to Uncover Student Attitudes about Diversity and Inclusion in Computer Science. ACM Trans. Comput. Educ. 22, 4, Article 52 (December 2022), 32 pages. https://dl.acm.org/doi/pdf/10.1145/355048730, last accessed on 2024/16/02.
30. Dawson Head Lynn E., Anderson E., Moghadam Hejazi S., Dorough E., Sessoms l., and Hovey Lynnly C.: Improving Diversity, Equity, and Inclusion in Doctoral Computing Education. In Proceedings of the 52nd ACM Technical Symposium on Computer Science Education (SIGCSE '21). Association for Computing Machinery, New York, NY, USA, 764–765. https://doi.org/10.1145/3408877.3432577
31. Long J., Chen T., Nguyen Viet Hung Q., and Yin H.: Decentralized Collaborative Learning Framework for Next POI Recommendation. ACM Trans. Inf. Syst. 41, 3, Article 66 (July 2023), 25 pages. https://doi.org/10.1145/3555374, accessed last on 2024/16/02.
32. Buchner, T.: The algorithmic imaginary: exploring the ordinary affects of Facebook algorithms. Information, communication & society 20, 1 (2017), 30–44, accessed last on 2024/16/02.
33. Ytre-Arne B. and Moe H.: Folk theories of algorithms: Understanding digital irritation. Media, Culture & Society 43, 5 (2021), 807–824. accessed last on 2024/16/02.