Algorithmic Bias in Algorithm-Driven User Interfaces: Recommendations for Fairness⋆

Stanley E. Abhadiomhen[1,2,*], Kiemute Oyibo[1]

[1]Department of Electrical Engineering and Computer Science, York University, 4700 Keele Street, Toronto, ON M3J 1P3, Canada
[2]Department of Computer Science, University of Nigeria, Nsukka 400241, Nigeria

Abstract

Algorithm-driven user interfaces (UI) have transformed digital experiences by enabling personalized interactions. However, these algorithms often encode biases that result in unfair, unethical, or non-inclusive user experiences. This paper examines how algorithmic personalization, including recommendation engines, dynamic pricing models, and targeted advertising, can lead to discriminatory practices, manipulative design patterns, and content exclusion. We argue that addressing these biases requires a fundamental shift in how personalization algorithms are designed and governed. To this end, we propose a framework for mitigating algorithmic bias through enhanced transparency, regular bias audits with fairness metrics, user-centric controls that allow individuals to modify algorithmic outputs, and the inclusion of diverse, representative training data.

Keywords

Algorithmic bias, UI/UX, personalization, dark patterns, fairness, ethical design

1. Introduction

Algorithm-driven user interfaces (ADUIs) use algorithms to personalize and optimize user interactions. ADUIs play a crucial role in shaping user experiences (UX), influencing everything from recruiting systems and customized news aggregation platforms to search results and personalized content recommendations [1]. For example, job platforms like LinkedIn (linkedin.com) use AI models [2] to recommend job listings based on a user's profile, experience, and preferences, prioritizing and ranking candidates or listings to enhance matching. Similarly, search engines such as Google (google.com) used algorithms to prioritize relevant websites based on factors like keywords, user intent, and search history, presenting the most useful results at the top for quicker access [3]. In addition, streaming platforms like Netflix and YouTube use recommendation models, such as collaborative filtering and deep learning, to suggest content based on viewing history, preferences, and interactions [4], continuously adapting to enhance user satisfaction and engagement. While personalization offers convenience and tailored content, it also introduces ethical concerns when underlying algorithms reinforce biases [5].

The reinforcement of bias through algorithm-driven personalization is particularly evident in areas such as job advertisement delivery, dynamic pricing, and content filtering [6]. Job ads on digital platforms, for instance, are known to disproportionately target male users and to exclude or underrepresent female users for certain roles. A notable case occurred with Facebook's (facebook.com) job advertising algorithm, which was accused of delivering job ads primarily to men and excluding women from seeing ads for such jobs [7]. Empirical evidence from the literature, such as studies by Zhang and Kuhn [8] and Galdon et al. [9], also identifies similar biases in job recommendation algorithms. Specifically, Zhang and Kuhn [8], through their study in which they audited four Chinese job boards using fictitious profiles differing only in gender, revealed that jobs recommended exclusively for male profiles often advertised higher wages and required more experience compared to those recommended to female profiles. Furthermore, the language in job ads targeted at female profiles tended to include significantly

BCSS 2025: The 13th International Workshop on Behavior Change Support Systems, May 5, 2025, Limassol, Cyprus.
*Corresponding author.
sabhadio@yorku.ca (S. E. Abhadiomhen); kiemute.oyibo@yorku.ca (K. Oyibo)
https://discover.academics.yorku.ca/Stanley.Abhadiomhen (S. E. Abhadiomhen)
0000-0002-9509-1915 (S. E. Abhadiomhen); 0000-0001-8300-3343 (K. Oyibo)

more content associated with stereotypical gender roles.

E-commerce sites are not exempt from discriminatory practices; some even adjust prices based on users' browsing behavior, resulting in economic discrimination. Illustratively, Amazon (amazon.com) has faced issues with its dynamic pricing algorithms, where users were charged different prices for the same product based on their browsing patterns [10]. In one instance [11], a "DVD case" was sold at varying prices, which Amazon attributed to random price tests. After customer outrage, Amazon refunded those who paid higher prices, acknowledging the discrepancy and resolving the issue. Further highlighting the impact of algorithmic pricing, a study by Chen et al. [12] analyzed dynamic pricing in the Amazon Marketplace, identifying over 500 sellers using such algorithms. Their findings revealed that while these sellers were more likely to win the Buy Box and achieve higher sales volumes, their prices were also more volatile, potentially leading to customer dissatisfaction.

Furthermore, dark patterns, which are deceptive and manipulative user interfaces crafted to make the user take a decision that is in the best interest of the service, are prevalent in the online environment, with research showing their effectiveness and potency [13][14]. They include hidden fees, default opt-ins, sneak into basket, trick question, and deceptive urgency messages, all of which exploit cognitive biases to drive engagement and revenue. For example, ride-sharing platforms, such as Uber (www.uber.com) and Lyft (www.lyft.com), often employ pricing algorithms that capitalize on users' fear of missing out, leading to higher prices during peak demand through practices like surge pricing [15, 16]. Additionally, while fare price increases due to factors like road repairs causing traffic may be fair in such conditions, users are typically unaware of the final charges until after the trip, which could be seen as exploiting their limited awareness of the fare. This paper argues that algorithmic bias in UI design is not an inevitable byproduct of personalization, but a consequence of prioritizing engagement and revenue over fairness and inclusivity. Therefore, we advocate for ethical design interventions to prevent discriminatory ad targeting, manipulative dark patterns, and exclusionary content curation.

2. Algorithmic Bias in UI

This section explores how algorithms can perpetuate discrimination and exclusion through various forms of biased design.

2.1. Discriminatory Advertising and Manipulative Design

Discriminatory advertising and manipulative design are becoming increasingly prevalent in digital platforms [17], disproportionately targeting specific groups, exploiting vulnerabilities, and undermining user autonomy. One significant example is personalized advertising [18], where algorithms display ads based on user data, such as demographics, browsing history, and previous interactions. As previously mentioned in a similar case involving Facebook, ads for high-paying jobs or promotions may be disproportionately targeted toward one gender, ethnicity, or socioeconomic class. Sometimes, these discriminatory practices go unnoticed, causing unfair outcomes. To detect such bias, it is crucial to examine the data distribution and performance of algorithms across different demographic groups, ensuring no unintended patterns of discrimination emerge in ad placements or content delivery.

In the case of manipulative design, a prime example is manipulative ad placement, where platforms exploit users' cognitive biases to influence their decisions [19]. This is often done by positioning urgent or limited time offers in highly visible locations, which encourages impulsive decisions. Such ads rely on tactics like time scarcity [20], social proof [21], or emotional triggers [22] to push users into making purchases or signing up for services they may not have originally intended to engage with. Additionally, some platforms use algorithms to specifically target economically vulnerable users, offering loans, high-interest financial products, or services that are not in the user's best interest. This practice not only exploits the user's financial situation but also perpetuates cycles of inequality and economic disenfranchisement [23]. Personalized pricing, where products or services are priced higher based on a user's perceived willingness to pay, is another example of such manipulative tactics.

Although discriminatory practices in advertising can involve dark patterns, they remain distinct in their focus. While discriminatory advertising targets specific groups, dark patterns manipulate users to maximize business goals, making them a form of manipulative design [24].

2.2. Dark Patterns and Exploitative UI

Dark patterns (DP) refer to user interface design choices that deceive or manipulate users into making decisions they might not otherwise make. Dark patterns (DPs) are designed to prioritize business goals, such as maximizing revenue or user retention, often at the expense of user autonomy and transparency [14], raising fairness concerns. According to Chen et al. [25], such dark patterns can exploit consumer behavior and disproportionately affect economically vulnerable users who do not realize that they are being charged more.

A typical example of DP is forced continuity subscriptions [26, 27], which occur when users sign up for a service with a free trial only to have their subscriptions automatically renewed without their explicit consent unless they take action to cancel. On the other hand, the cancelation process is often obfuscated by layers of complex steps or hidden options, taking advantage of the status quo bias, where users are more likely to accept the default setting rather than actively opting out. Services such as subscription boxes [28] or digital media platforms often use this technique to maximize user retention and revenue. Similarly, hidden opt-outs and pre-checked boxes are also common forms of DPs, where options for additional services, such as extended warranties or email newsletters, are automatically selected for users without their consent when making an online purchase. Confirm shaming, Nagging, Bait and switch, and Misdirection are also common types of dark patterns. Confirm shaming involves making users feel guilty for not opting into a decision, such as guilt-tripping them for unsubscribing or making them act differently than they normally would [24]. Nagging repeatedly prompts users to take action, often annoying them into compliance. Bait-and-switch lures users with one offer and then changes it [29], while misdirection distracts users from critical information, leading them to make uninformed choices.

2.3. Exclusionary Algorithms

Algorithmic-driven UI personalization can also result in digital exclusion. One of the most well-known forms of exclusion in this context occurs in the realm of accessibility. Take, for example, a voice recognition system or a virtual assistant trained primarily on data from native English speakers with Western accents. In this case, the system may struggle to accurately recognize users with nonnative accents or those speaking other languages. In the same vein, a facial recognition system trained predominantly on images of white individuals may not perform well on people with darker skin tones. Moreover, Buolamwini et al. [30] demonstrated that facial recognition technology tends to be less accurate for individuals with darker skin tones, often leading to biased and unfair results.

3. Design Recommendations

To mitigate algorithmic bias in UI, it is important that both designers and developers prioritize fairness-aware strategies. First, algorithmic transparency is essential to enable users to understand how their data influence UI decisions. This not only helps build trust, but also empowers users by giving them insight into the factors that shape their digital experiences. Secondly, conducting regular bias audits and evaluating algorithms for discriminatory patterns is crucial. For example, evaluating recommendation and personalization algorithms to identify and correct biases that can cause unfair or harmful results.

Furthermore, user-centric controls should be incorporated, providing users with the ability to modify algorithmic outputs, such as content filtering options, ad preference settings, and even the level of personalization they receive. Although many platforms already offer basic customization, these options often lack the depth needed to fully address biases or provide meaningful transparency. Thus, empowering users to control not only the content they see, but also the underlying algorithms that determine how content is served to them would foster a more personalized and equitable experience. Lastly, inclusive data representation is vital, where algorithms are trained on diverse datasets that represent different races, genders, ages, and cultural backgrounds. This approach will help reduce biases that may arise from narrow datasets and ensure that algorithms can better serve the needs of all users, ensuring a more equitable and inclusive user experience that minimizes the risk of bias while improving the user experience.

However, these proposed strategies come with potential trade-offs, such as the balance between fairness and personalization accuracy. For example, prioritizing fairness may sometimes lead to reduced personalization or accuracy in the recommendations or advertisements a user receives. Additionally, implementing these solutions may face challenges, including technical feasibility, especially in legacy systems, and potential resistance from stakeholders in the industry who may be reluctant to adopt new, more transparent approaches due to cost, time, or the perceived disruption to existing business models. Possible solutions include adopting incremental changes, starting with small-scale pilots to demonstrate the benefits of fairness-aware systems, and fostering collaboration between designers, developers, and industry leaders to align on long-term goals for algorithmic fairness.

4. Conclusions

Algorithm-driven personalization has the potential to enhance the user experience, but it can also introduce significant ethical risks, including creating echo chambers and serving the user with one-sided information that they are comfortable with, i.e., they want to see and hear. This paper takes the position that biased and unfair UI/UX outcomes are not inevitable consequences of personalization, but rather design choices that prioritize engagement and profitability over fairness, autonomy, and objectivity. Addressing bias in ADUIs requires a concerted effort from policy makers to designers and developers to create transparent, fair, and inclusive digital experiences. Future research should explore regulatory frameworks and technical solutions that can help mitigate algorithmic bias in interface design and online service delivery. Governments and industry bodies must establish clear guidelines for bias audits, algorithmic transparency, and data disclosure. International collaboration on regulatory standards will ensure fairness across global platforms. In addition, incentivizing companies to prioritize fairness and inclusivity in their algorithms will help foster a more ethical and user-centered approach.

5. Acknowledgments

This work was undertaken thanks in part to funding from the Connected Minds Program, supported by Canada First Research Excellence Fund, Grant No. CFREF-2022-00010.

References

[1] D. Shin, B. Zhong, F. A. Biocca, Beyond user experience: What constitutes algorithmic experiences?, International Journal of Information Management 52 (2020) 102061.

[2] X. Yan, J. Yang, M. Obukhov, L. Zhu, J. Bai, S. Wu, Q. He, Social skill validation at linkedin, in: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2019, pp. 2943–2951.

[3] R. K. Chaudhary, How seo impacts on websites rank?, Multi-Disciplinary Explorations: The Kasthamandap College Journal 2 (2024) 161–174.

[4] P. Suman, M. Yang-Sae, Netflix, amazon prime, and youtube: comparative study of streaming infrastructure and strategy, Journal of Information Processing Systems 18 (2022) 729–740.

[5] H. Ashman, T. Brailsford, A. I. Cristea, Q. Z. Sheng, C. Stewart, E. G. Toms, V. Wade, The ethical and social implications of personalization technologies for e-learning, Information & Management 51 (2014) 819–832.

[6] T. P. Liang, H. J. Lai, Y. C. Ku, Personalized content recommendation and user satisfaction: Theoretical synthesis and empirical findings, Journal of Management Information Systems 23 (2006) 45–70.
[7] K. Hao, Facebook's ad algorithms are still excluding women from seeing jobs, MIT Technology Review (2021). Retrieved January 21, 2022.
[8] S. Zhang, P. J. Kuhn, Measuring Bias in Job Recommender Systems: Auditing the Algorithms, Technical Report, National Bureau of Economic Research, 2024.
[9] G. Galdon Clavell, M. Martín Zamorano, C. Castillo, O. Smith, A. Matic, Auditing algorithms: On lessons learned and the risks of data minimization, in: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, 2020, pp. 265–271.
[10] K. Lippert-Rasmussen, L. A. Munch, Price discrimination in the Digital Age, 2021.
[11] A. Gautier, A. Ittoo, P. Van Cleynenbreugel, Ai algorithms, price discrimination and collusion: a technological, economic and legal perspective, European Journal of Law and Economics 50 (2020) 405–435.
[12] L. Chen, A. Mislove, C. Wilson, An empirical analysis of algorithmic pricing on amazon marketplace, in: Proceedings of the 25th international conference on World Wide Web, 2016, pp. 1339–1349.
[13] K. Oyibo, The influence of user knowledge and usage behaviour on decision-making and perceived reputation of streaming sites that use dark patterns, Behaviour & Information Technology (2025) 1–20.
[14] A. Mathur, G. Acar, M. J. Friedman, E. Lucherini, J. Mayer, M. Chetty, A. Narayanan, Dark patterns at scale: Findings from a crawl of 11k shopping websites, Proceedings of the ACM on human-computer interaction 3 (2019) 1–32.
[15] J. D. Martini, International regulatory entrepreneurship: Uber's battle with regulators in france, San Diego Int'l LJ 19 (2017) 127.
[16] H. H. Perritt Jr, Don't burn the looms: Regulation of uber and other gig labor markets, SMU Sci. & Tech. L. Rev. 22 (2019) 51.
[17] J. K. Bahangulu, L. Owusu-Berko, Algorithmic bias, data ethics, and governance: Ensuring fairness, transparency and compliance in ai-powered business analytics applications (2025).
[18] F. De Keyzer, N. Dens, P. De Pelsmacker, Is this for me? how consumers respond to personalized advertising on social network sites, Journal of Interactive Advertising 15 (2015) 124–134.
[19] M. Amirpur, The role of cognitive biases for users' decision-making in is usage contexts (2017).
[20] M. A. G. P. Pattinaja, M. Mangantar, M. Pandowo, The impact of user interface and time scarcity on purchase intention through e-commerce shopee among young adults in manado, Jurnal EMBA: Jurnal Riset Ekonomi, Manajemen, Bisnis dan Akuntansi 11 (2023) 149–160.
[21] K. K. Kim, W. G. Kim, M. Lee, Impact of dark patterns on consumers' perceived fairness and attitude: Moderating effects of types of dark patterns, social proof, and moral identity, Tourism Management 98 (2023) 104763.
[22] J. A. Galindo, S. Dupuy-Chessa, N. Mandran, E. Céret, Using user emotions to trigger ui adaptation, in: 2018 12th International Conference on Research Challenges in Information Science (RCIS), 2018, pp. 1–11.
[23] L. SANCHEZ CHAMORRO, Disentangling vulnerability to manipulative designs: An experiential perspective to rethink resistance strategies (2024).
[24] M. Potel-Saville, M. Da Rocha, From dark patterns to fair patterns? usable taxonomy to contribute solving the issue with countermeasures, in: Annual Privacy Forum, Springer, 2023, pp. 145–165.
[25] J. Chen, J. Sun, S. Feng, Z. Xing, Q. Lu, X. Xu, C. Chen, Unveiling the tricks: Automated detection of dark patterns in mobile applications, in: Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, 2023, pp. 1–20.
[26] A. Mathur, M. Kshirsagar, J. Mayer, What makes a dark pattern... dark? design attributes, normative considerations, and measurement methods, in: Proceedings of the 2021 CHI conference on human factors in computing systems, 2021, pp. 1–18.
[27] F. Nygren, P. Tran, Streaming in the dark: Analysing video streaming services for dark patterns:

A user interface study (2024).
[28] N. Umashankar, K. H. Kim, T. Reutterer, Understanding customer participation dynamics: the case of the subscription box, Journal of Marketing 87 (2023) 719–735.
[29] R. Riaz, A. Vasconcelos, P. Pinto, An overview of user psychological manipulation techniques in ui/ux web design, in: 2024 Cyber Awareness and Research Symposium (CARS), IEEE, 2024, pp. 1–6.
[30] J. Buolamwini, T. Gebru, Gender shades: Intersectional accuracy disparities in commercial gender classification, in: Conference on fairness, accountability and transparency, 2018, pp. 77–91.