

Prospectus

Pairings: Enhancing Food and Wine Experiences Using Machine Learning

(Technical Topic)

Examining Algorithmic Bias in Recommender Systems: Ethical Challenges and Solutions

(STS Topic)

By

Eli Herman

December 15, 2024

Technical Project Team Members: n/a

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

Signed: Eli Herman

Introduction

AI-driven recommendation systems have become integral to shaping consumer experiences across various industries, offering personalized suggestions that aim to enhance user satisfaction. However, these systems often suffer from algorithmic bias, a critical issue that limits consumer choice and diversity while amplifying existing societal inequities. Bias in recommendation systems manifests when algorithms prioritize popular trends, entrenched feedback loops, or mainstream content, thereby marginalizing alternatives. This “echo chamber” effect reflects broader societal patterns, as argued by Noble (2018), who critiques search engines for perpetuating racial and gender biases, and O’Neil (2016), who highlights the practical dangers of biased systems in decision-making processes.

My capstone project tackles this issue through both a technical project and an STS research study. The technical project, *Pairings*, focuses on building a machine learning-based wine recommendation app that personalizes wine suggestions based on meal components, user preferences, and budget. By addressing limitations in personalization and diversity, *Pairings* seeks to create a fairer recommendation tool that enhances dining experiences while promoting equitable outputs. The STS portion of my research investigates how algorithmic bias affects consumer choice and proposes strategies to balance personalization and fairness. Using case studies of platforms like Netflix and Spotify, this research contextualizes the societal impacts of biased algorithms and explores ethical frameworks for mitigating their consequences. Together, the technical and STS projects provide a comprehensive exploration of algorithmic bias, offering practical solutions and broader insights to improve fairness, transparency, and accountability in AI systems.

Technical Topic

Traditional wine-pairing systems rely on static rules that fail to account for personal preferences, dining contexts, and budget constraints. These limitations reduce their ability to provide meaningful and personalized recommendations. My technical project, *Pairings*, addresses this gap by developing an intelligent wine recommendation app that integrates contentbased filtering and collaborative filtering into a hybrid machine learning model. FernándezTobías et al. (2011) provide foundational support for this approach, demonstrating that hybrid models improve recommendation accuracy by combining the strengths of multiple filtering techniques. Their research highlights that content-based systems excel at analyzing item attributes, while collaborative filtering effectively adapts to user preferences, and combining the two mitigates their individual limitations.

The *Pairings* model is built on these principles. Content-based filtering analyzes wine attributes—such as varietals, flavor profiles, and regions—allowing the app to suggest wines that match specific meal components. Collaborative filtering refines recommendations based on patterns in user behavior and feedback, enabling the system to adapt its suggestions over time. Resnick and Varian (1997) emphasize the importance of integrating user feedback to improve learning in recommendation systems. However, they caution that feedback loops can entrench existing trends, reinforcing popular options at the expense of diversity. This challenge aligns with Binns (2018), who advocates for fairness frameworks in algorithmic systems to ensure transparency and inclusivity.

To address these concerns, *Pairings* prioritizes dataset diversity and transparency in its design. User input—such as meal details, taste preferences, and budget constraints—feeds into the hybrid recommendation model, producing tailored suggestions that evolve with continued user interaction. The app’s recommendation engine is being developed using Python and TensorFlow, with an iterative testing process to optimize accuracy and fairness. By combining algorithmic precision with fairness considerations, *Pairings* not only enhances user satisfaction but also serves as a case study for mitigating bias in niche recommendation systems.

STS Topic

Algorithmic bias in recommendation systems raises ethical and societal concerns by influencing consumer behavior and reinforcing systemic inequalities. When algorithms prioritize mainstream content or entrenched preferences, they reduce user exposure to diverse alternatives, creating a feedback loop that further marginalizes underrepresented options. Noble (2018) critiques search engines for embedding societal biases into their rankings, perpetuating racial and gender inequities under the guise of algorithmic neutrality. Similarly, O’Neil (2016) demonstrates how biased algorithms in contexts such as credit scoring and job applications exacerbate systemic injustices, often without users’ awareness. These critiques form the foundation of my STS research, which examines how recommendation systems balance personalization and fairness to ensure equitable outcomes.

My research employs ethical frameworks proposed by scholars like Binns (2018) and Barocas et al. (2019), who argue for transparency, accountability, and dataset diversity in algorithmic systems. Binns advocates for user-centered design and transparency to make

algorithms understandable to users, while Barocas et al. caution that transparency alone is insufficient without diverse datasets. This insight is directly applicable to *Pairings*, where diverse wine data and inclusive feedback mechanisms are prioritized to counteract entrenched biases.

To contextualize these findings, my STS research analyzes case studies of platforms like Netflix and Spotify. Verbert et al. (2012) demonstrate how context-aware recommendations improve user satisfaction but warn that over-personalization limits exposure to diverse content. Similarly, Diakopoulos (2016) highlights the importance of accountability in systems that influence consumer decisions, arguing that algorithmic processes must be subject to scrutiny to prevent unintended consequences. Netflix's recommendation algorithm, for example, optimizes user engagement but often reinforces mainstream content, limiting discovery of less popular media. Spotify's recommender system similarly amplifies established artists at the expense of emerging talent, reflecting the tension between personalization and diversity.

By synthesizing these insights, my STS research proposes strategies to design fairer and more inclusive recommendation systems. These strategies include diversifying training datasets, improving algorithmic transparency, and embedding fairness frameworks into system design. Applying these principles to *Pairings* demonstrates how recommendation systems can balance personalization with equity, ensuring diverse and meaningful user experiences.

Conclusion

Through my capstone project, I aim to address both the technical and societal challenges of algorithmic bias in recommendation systems. The *Pairings* app demonstrates how hybrid machine learning models can deliver personalized, context-aware wine recommendations while prioritizing fairness and transparency. Simultaneously, my STS research examines the broader implications of algorithmic bias, proposing strategies to balance personalization with diversity in AI systems. Together, these projects contribute to a more equitable and accountable approach to recommendation systems, offering insights that extend beyond the wine industry to other consumer-focused applications.

References

1. Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and machine learning*. fairmlbook.org.
2. Binns, R. (2018). Fairness in machine learning: Lessons from political philosophy. *Proceedings of the 2018 Conference on Fairness, Accountability, and Transparency*, 149-159.
3. Binns, R., Veale, M., Van Kleek, M., & Shadbolt, N. (2018). 'It's reducing a human being to a percentage': Perceptions of justice in algorithmic decisions. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1-14.
4. Brode, M. (2020). Machine learning for wine recommendations: A case study. *Journal of Food and Beverage Analytics*, 7(3), 145-158.
5. Diakopoulos, N. (2016). Accountability in algorithmic decision-making. *Communications of the ACM*, 59(2), 56-62.
6. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1-35.
7. Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press.
8. O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
9. Resnick, P., & Varian, H. R. (1997). Recommender systems. *Communications of the ACM*, 40(3), 56-58.
10. Smith, J., & Allen, D. (2019). Recommender systems in gastronomy: Bridging user experience and AI. *International Journal of AI in Hospitality and Tourism*, 6(1), 40-55.
11. Sweeney, L. (2013). Discrimination in online ad delivery. *Communications of the ACM*, 56(5), 44-54.
12. Verbert, K., Manouselis, N., Ochoa, X., Wolpers, M., Drachsler, H., Bosnic, I., & Duval, E. (2012). Context-aware recommender systems: A review and future directions. *User Modeling and User-Adapted Interaction*, 22(1-2), 157-190.