Tempo: Development of an iOS App and Interactive Bluetooth Speaker System for Personalized Music Recommendations and Social Engagement
(Technical Project)

Bias and Diversity in AI-Based Music Recommendation Systems: Addressing Algorithmic Reinforcement and Cultural Homogeneity
(STS Project)

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Electrical Engineering

By
Michelle Monge

December, 2024

Technical Team Members:
Joey Cohen, Bella Heintges, Thomas Keathly, Naomi Solomon

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.

ADVISORS
Adam Barnes, Department of Electrical Engineering
Kent Wayland, Department of Engineering & Society

Introduction

How do AI music recommendation systems contribute to the formation of algorithmic filter bubbles, and what are the implications for cultural diversity and user autonomy?

Artificially intelligent music recommendation systems, such as those used by Spotify have changed the way users discover and engage with music. Spotify relies on algorithms to tailor playlists and suggestions based on user data, such as listening history and preferences which offer personalized results. While these features offer convenience and personalization, the algorithms risk confining users to narrow content that already aligns with their existing tastes. These isolating filter bubbles in music recommendation systems result from a combination of factors, including algorithmic design priorities such as engagement metrics, and commercial, profitable interests. This system can then be designed to maximize user interaction and retention, at the cost of promoting and prioritizing diversity. As a result, users may become less likely to encounter new or unfamiliar music, limiting their cultural exposure and reinforcing a cycle of sameness. Additionally, certain artists who post their music on Spotify may not be recommended if they do not fit into these filter bubbles. This issue contains broader implications for cultural consumption as it shapes not only individual listening habits but also societal trends in music discovery and which voices get prioritized over others.

This project aims to analyze how filter bubbles are created in AI music recommendation systems by examining algorithmic components and the underlying priorities driving its design. The project will then explore how user behavior, commercial interests, and algorithmic optimization contribute to the problem and how researchers and industry practitioners are addressing these challenges. The project will attempt to address the overarching problem through two complementary approaches. The technical project seeks to develop *Tempo*, an iOS application integrated with a Bluetooth powered speaker system where the system will incorporate user-adjustable AI music recommendations, where users can test and influence algorithmic outputs. The STS research will then complement this effort in critically examining the biases embedded in AI music recommendation systems, using Spotify as a case study. In analyzing how user behaviors and commercial priorities shape recommendation algorithms, this research seeks to contribute to a deeper understanding of the ethical and cultural implications of AI personalization in music recommendation systems.

Technical Project: Tempo; Development of an iOS App and Interactive Bluetooth Speaker System for Personalized Music Recommendations and Social Engagement

How can an iOS app and a companion Bluetooth-enabled speaker system be developed to provide personalized music recommendations while promoting social engagement and user control over algorithmic parameters?

The technical project addresses limitations in existing music recommendation systems such as Spotify, by introducing Tempo, an iOS application integrated with a Bluetooth speaker. While Spotify offers features like recommendations, collaborative playlists, and user-following systems, users often express frustration with limited functionality. Common complaints include difficulty finding friends on the platform, repetitive recommendations that overly emphasize recently played tracks, and a lack of transparency about how user statistics are calculated. Tempo seeks to address these issues by offering features that allow users to enter prompts to generate specific music suggestions, view compatibility percentages with friends alongside explanations of the calculations, and access personalized musical statistics, such as top genres and artists. The project also integrates a Bluetooth speaker that connects to the app, enabling users to control playback and interact with the system. Tempo then aims to fill a gap in the market by combining personalized recommendations with social connectivity and increased user control over algorithms.

Tempo iOS Application

The Tempo iOS app will provide personalized music recommendations while encouraging social engagement. It will use Spotify's Application Programming Interface (API), which allows external applications to retrieve user data, including playlists, liked songs, and listening history. OpenAI's Generative Pre-trained Transformer (GPT) will then process user prompts to generate recommendations tailored to specific moods or themes within the listening tab. For instance, users can input prompts such as "upbeat workout music," and the app will analyze the request to generate a list of personalized recommended tracks and then save data to Firebase, a database service powered by Google. This functionality then allows users to guide the algorithm. The app also incorporates social features to increase user interaction as within the compatibility tab, compatibility scores can be calculated by comparing shared listening data between users such as overlapping liked songs, and these scores are displayed on a community web which enables users to explore musical connections with friends. Additionally, a profile tab can be accessed to view listening data at all times, such as top 50 songs, and a top genre pie chart breakdown. Overall, transparency is a key focus, with the app allowing users to adjust parameters and view their data at all times.

Tempo Bluetooth Speaker System

The hardware component of the project includes a Bluetooth speaker system designed for high-quality audio and easy, interactive functionality. The speaker features two Dayton Audio DMA58-4 drivers, which provide a wide frequency range and is powered by a Sure Electronics Class-D amplifier for increased power of an audio signal to be heard through the speaker. Additionally, a Raspberry Pi microcontroller will be programmed to drive an integrated liquid-crystal display (LCD), which shows real-time playback information, including song titles, artists, and track durations. The speaker also includes physical controls, such as buttons for volume adjustment, track skipping, and muting, to provide users with a straightforward interface for managing playback. Bluetooth connectivity is implemented using the CSR8645 Bluetooth module, which supports audio streaming and playback control, enabling real-time interaction between the iOS app and the speaker.

Testing and Goals

Testing plays a critical role in validating the functionality and usability of Tempo. Component testing focuses on verifying individual features, such as Bluetooth connectivity, audio playback, and app navigation. Integration testing then evaluates the synchronization between the app and the speaker, including real-time updates on the LCD display during playback. Compliance testing is also conducted to make sure that there is adherence to Apple's App Store guidelines and Bluetooth safety standards. This structured approach to testing helps identify and resolve issues at both the component and system levels. Overall, the Tempo system aims to deliver a Bluetooth speaker with integrated playback controls and an LCD display, along with an iOS application that offers personalized, transparent, and socially engaging music recommendations.

STS Research Project: Bias and Diversity in AI Music Recommendation Systems – Understanding Algorithmic Reinforcement and Cultural Homogeneity

How do AI music recommendation systems such as Spotify reinforce cultural bias and homogeneity, and how do the Algorithms underlying mechanisms shape user experiences and cultural consumption?

Context

AI music recommendation systems have significantly reshaped how users discover music as a popular streaming application; however, Spotify's recommendation system has raised concerns and questions about the unintended consequences of algorithmic design. By optimizing for metrics like frequent user interaction and retention, Spotify often reinforces user preferences and promotes mainstream genres. Over time, this narrowing of recommendations contributes to increased filter bubbles which limit exposure to diverse music and lesser-known artists.

Filter bubbles are not merely a technical issue; they represent a broader cultural and societal concern. The bubbles shape not only individual user experiences but also patterns of cultural consumption. Artists outside mainstream genres face reduced visibility, which in turn affects their opportunities for market success and cultural representation. This creates a feedback loop where algorithms amplify what is already popular and liked, reducing the diversity of music that users encounter and thus reinforcing cultural homogeneity. Understanding these dynamics is then essential for investigating the broader ethical implications of algorithmic design in cultural industries.

Spotify operates within a complex sociotechnical system that includes multiple stakeholders, each with distinct priorities. Users expect recommendations tailored to their preferences, developers focus on optimizing engagement, artists seek greater visibility, and policymakers advocate for fairness and inclusivity. These competing priorities influence how recommendation systems are designed and deployed where investigating how these stakeholders influence and are influenced by Spotify's recommendation algorithms can provide deeper insights into the cultural and ethical consequences of algorithmic reinforcement.

Literature Review

The literature on ethics of recommendation systems highlight the mechanisms by which biases are embedded and perpetuated in algorithmic systems. Algorithms often reflect the values of the people behind the design such as developers, which can reinforce societal inequalities if left unchecked as these algorithms have never been neutral. (Baccigalupo, C., & Plaza, E. 2006), (Benkler, Y. 2006). Transparency and accountability are then necessary for understanding and addressing these biases, particularly in systems that significantly influence user behavior such as

popular streaming services (Gao and Yu 2024). With the design of these algorithms, there needs to be a sufficient level of trust and user agency if long-term engagement is to be achieved, suggesting that recommendation systems must balance personalization with user autonomy if there is any chance of success (Liu 2024).

Empirical studies provide evidence of how recommendation algorithms reinforce bias. Anderson et al. (2020) found that Spotify's algorithmic playlists often prioritize familiar content, reducing opportunities for discovery. As a result, the Spotify algorithm has created isolated information bubbles or cocoons that trap users within their own existing preferences, leading to user frustrations and a loss of cultural diversity (Borgesius et al. 2016), (O'Neil, C. (2016). Spotify recommendation systems also tend to propagate cultural stereotypes, as user preferences are often shaped by the social and demographic biases embedded within the training data (North and Hargreaves 2007). Spotify then acts as a gatekeeper, controlling which artists and genres reach large audiences and thus shaping cultural norms in the process.

Currently, Spotify offers little transparency to users about how their algorithms work. The app is currently pushing for users to try their new AI playlist feature, prompting constant ads; but, there are no links to how the system works on the app. The app is then currently pushing new features to keep engagement active compared to competitors, where Spotify must be able to provide users with music that is more satisfying to their tastes (Li, J. 2022). The lack of transparency then perpetuates inequalities as there are little user adjustable settings to balance personalization nor a widely available breakdown of how the system works (Sonboli et al. 2021).

Ultimately what is missing in the literature, is a detailed investigation into how Spotify algorithms operate such as what is in the code, and if the algorithms are specifically trained to favor certain genres and artists over others? Scholars have identified the problem of filter bubbles and biases, but there is less focus on how these issues emerge from the interplay between algorithmic design, user behavior, and commercial priorities. By examining how Spotify's algorithms are shaped by these factors, this research will contribute to a deeper understanding of the cultural and ethical stakes of recommendation systems.

Methodology

This research uses the theoretical frameworks of mutual shaping and value-laden design to investigate how Spotify's recommendation algorithms reinforce cultural biases and influence user behavior. With this system, there exists a dynamic feedback loop between users and technology, where user interactions inform algorithmic adjustments, and these adjustments, in turn, shape user experiences. This framework will guide the analysis of how user engagement metrics—such as repeat interactions with familiar content—drive Spotify's algorithms and contribute to filter bubbles.

Value-laden design highlights how algorithms are not neutral but reflect the values and priorities of their developers and the platforms they operate on. By focusing on Spotify's prioritization of engagement and retention, this research will critique the embedded values that shape algorithmic outputs, such as favoring mainstream, already liked content over diversity. This framework will then provide insights into how Spotify's design choices contribute to cultural homogeneity and limit artistic visibility.

To study these dynamics, the research will employ a document-based analysis of Spotify's publicly available technical reports, privacy policies, developer guidelines, and patent filings. These documents will be examined to identify the platform's design priorities, data-handling practices, and algorithmic mechanisms. Academic studies, such as peer-reviewed articles and case studies about Spotify and similar platforms, will also provide insights into the technical and cultural impacts of recommendation systems.

Additionally, the research will include a study of recommendation systems by analyzing the metadata of playlists generated through Spotify's API. By simulating user behaviors and analyzing playlist outputs, the study will explore how algorithmic outputs differ based on user input. For example, variations in prompts, listening habits, or levels of user engagement will be used to evaluate the algorithm's tendency to reinforce existing preferences versus promoting diverse content. Finally, the research will synthesize findings from these methods to develop a comprehensive critique of Spotify's recommendation systems. The focus will be on understanding the mechanisms by which these systems operate and their broader cultural implications.

Conclusion

The methodology focuses on analyzing existing documentation, academic studies, and computational outputs. The project will then critically examine and test the mechanisms driving Spotify's recommendation algorithms and their implications for cultural diversity, contributing to the broader discourse on AI ethics in music recommendation systems.

References

Anderson, A., Maystre, L., Mehrotra, R., et al. (2020). Inferring similarity between music objects with applications to recommendation. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. https://doi.org/10.1145/1101826.1101840

Baccigalupo, C., & Plaza, E. (2006). Case-based sequential ordering of songs for playlist recommendation. In *Advances in Case-Based Reasoning* (pp. 114-127). https://doi.org/10.1007/11805816_22

Benkler, Y. (2006). *The wealth of networks: How social production transforms markets and freedom.* Yale University Press. *Social Science Computer Review, 26*(2), 259-261. https://doi.org/10.1177/1084713807301373

Borgesius, F. J. Z., Trilling, D., Möller, J., et al. (2016). Should we worry about filter bubbles? *Internet Policy Review, 5*(1). https://doi.org/10.14763/2016.1.401

Gao, D., & Yu, D. (2024). Challenges and cracks: Ethical issues in the development of artificial intelligence. *Science, Technology and Society, 29*(3), 415-434. https://doi.org/10.1177/09717218241246372

Li, J. (2022). Analysis of the trend of Spotify. *BCP Business & Management, 34*, 919–926. https://doi.org/10.54691/bcpbm.v34i.3112

Liu, W. (2024). How do we move towards true artificial intelligence? *Proceedings of the IEEE International Conference on High Performance Computing and Communications* (HPCC-DSS-SmartCity-DependSys53884). https://doi.org/10.1109/HPCC-DSS-SmartCity-DependSys53884.2021.00321

O'Neil, C. (2016). Weapons of math destruction: How big data increases inequality and threatens democracy. Crown Publishing. https://doi.org/10.5860/crl.78.3.403

North, A. C., & Hargreaves, D. J. (2007). Lifestyle correlates of musical preference: Relationships, living arrangements, beliefs, and crime. *Psychology of Music, 35*(1), 58–87. https://doi.org/10.1177/0305735607068898

Sonboli, N., et al. (2021). Fairness and transparency in recommendation: The users' perspective. *arXiv preprint*. arXiv:2103.08786