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

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Michelle Monge

Spring 2025

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.

Advisor

Kent Wayland, Department of Engineering & Society

Introduction

Artificial Intelligence (AI) is altering the way people discover and engage with music. Major music streaming platforms, such as Spotify, now use complex AI recommendation systems that constantly analyze user listening history, preferences, and general behavior to generate personalized playlists. While this personalization offers convenience to the user, it also isolates them. In consistently delivering algorithmic recommendations tailored primarily around Spotify's commercial priorities, the platform risks constraining users within a limited, corporate-driven selection of content. This phenomenon, known as a filter bubble, occurs when personalized algorithms disproportionately favor mainstream or highly popular content, isolating users from diverse or less visible musical genres and artists. While Spotify markets personalization as elevating discovery, it paradoxically narrows individual exposure by focusing heavily on already successful content, reinforcing cultural homogeneity within music consumption. Later sections of this paper will further explore how this apparent contradiction between personalized user experiences and Spotify's commercial interests manifests in actual user behavior and content diversity outcomes.

The significance of this research lies in the exploration of the broader cultural and ethical implications of such algorithmic practices. On one hand, personalized recommendations are framed as a fun, new way to listen and find music; on the other hand, Spotify's algorithm limits opportunities for emerging and niche music. Consequently, there now exists a skew in the cultural consumption of music toward mainstream trends, which become profitable for large labels and for Spotify itself. The Spotify algorithm is not neutral in this case as it is built with company priorities in hand, and inevitably the power imbalance across all actors involved only increases. In essence, user freedom of choice diminishes as Spotify dictates whose voices are heard and ignored. In a media landscape where digital platforms increasingly shape public discourse and cultural identity, understanding the mechanisms that drive these resulting algorithmic filter bubbles becomes essential. This paper examines how passive user behavior, commercial interests, and biased algorithmic design collectively contributes to filter bubbles and the resulting narrowing of musical diversity, and what that means for both individual autonomy and societal cultural trends. In summary, this research examines: *How do the underlying mechanisms of Spotify's AI music recommendation system reinforce cultural bias and homogeneity, and what are the implications for user consumption and experience?*

Background

Spotify's AI music recommendation system must be understood within the broader context of digital media's rapid evolution. Over the past decade, the integration of machine learning algorithms into various applications has redefined user experiences across various domains - ranging from social media feeds to online shopping. In the case of music streaming, these algorithms rely on vast amounts of user data such as listening history, existing playlists, and social interactions for training. *Machine Learning* in this context, refers to a subset of AI where statistical models learn patterns from large datasets to predict or recommend content. Within this context, it is then important to acknowledge that this data driven approach is not inherently neutral. The algorithm will inevitably embed the commercial interests of the platform and developer behind it, such as increasing user retention and engagement. Algorithmic priorities can thus be optimized based on clicks or streams and inadvertently promote popular, profitable content, thus confining and controlling user's.

The sociotechnical context can also be understood through the lens of *actor-network* *theory (ANT)*, which emphasizes that human and non-human actors actively shape each other in a network. In the context of music streaming, users generally demand both ease of use and convenience, such as app navigation, quality of streaming, aesthetics, useful features, and music recommendations. Artists then seek fair representation and profit in a marketplace dominated by a handful of major labels, all the while platforms operate under business models that favor algorithms optimized for profit; which is typically measured by metrics such as time spent on the service and ad revenue. Through an ANT perspective, these human and nonhuman elements (e.g. user preferences, artist contracts, algorithms, and data analytics tools) mutually reinforce each other, creating a complex ecosystem of technological design, commercial priorities, and power dynamics. Consequently, the interplay of these actors shape not only how music is consumed, but also how it is valued in the era of digital streaming.

Literature

Just as all people have some inherent bias, so do machine learning algorithms and AI. People create algorithms, and thus any of their biases and priorities spew into the resulting technology. In Spotify's case, it is evident that their algorithm substantially shapes and controls user consumption and experience through its algorithmic priorities. In the public eye, because these information intermediaries automate user's core operations, they are often treated as objective and credible (Bozag, 2013). Machines, not humans, appear to make the crucial decisions, thus creating the impression that the algorithms avoid selection and description biases inherent in any human-edited media. However, bias can manifest itself within computer systems in multiple, interconnected ways. Pre-existing biases rooted in societal inequalities and stereotypes can inadvertently shape algorithm design, influencing which data points developers deem important or worthy of inclusion. For instance, cultural biases might affect decisions on what genres, artists, or demographics receive priority in data collection, thereby reinforcing existing societal hierarchies. Technical bias also arises from inherent limitations in datasets themselves; if training data disproportionately represent majority groups or mainstream content, the resulting algorithms become systematically biased against minority groups, niche genres, and emerging artists. Consequently, algorithms trained on skewed datasets amplify disparities by underrepresenting or misrepresenting less prevalent or less visible communities, further marginalizing already underserved populations within music streaming platforms.

Algorithmic bias and increasing personalization can lead to the formation of filter bubbles, where users experience reduced exposure to content that differs significantly from mainstream or popular trends. Within the current streaming industry, AI recommendation systems like Spotify's are primarily tuned to maximize user engagement, often defined by metrics such as listening duration, repeated plays, and immediate user satisfaction—rather than explicitly prioritizing diversity or discovery of niche content. While personalization features are marketed as enhancing user convenience and enjoyment, they inadvertently promote conformity by frequently recommending familiar, commercially successful content. As shown in existing research, algorithms optimized strictly for engagement tend to amplify existing preferences and popular trends, paralleling broader societal patterns seen in areas like news media and social platforms, where familiar content is favored over diverse perspectives (Interian et al. 2023). This emphasis on engagement thus contributes to a narrowing of musical exposure, reflecting deliberate design priorities rather than inherent technological limitations.

Algorithmic filter bubbles lead to a corporate-controlled narrative that offers an illusion of choice. When the algorithm confines users to isolated filter bubbles, their ability to make autonomous listening choices is compromised; this reinforces a corporate-controlled narrative rather than a more organic one. As argued by Anderson et. al, (2020) when “user’s begin to listen to more diverse music, they do so by shifting away from algorithmic consumption and increasing their organic consumption” (p.1). Spotify’s overarching business strategy consists of maximizing user engagement, but this inadvertently affects the diversity of music consumption at large and creates a controlled environment where Spotify holds a great amount of power (Li, 2022).

With the push of new AI features, musical visibility thus becomes dictated by engagement-driven categorization rather than artistic merit and complete user choice (Prey & Esteve-Del-Valle, 2024). With the inherent pressure to maximize profit and remain relevant, Spotify must constantly introduce new features that keep users entertained (Borgesius et al., 2016). Currently, AI features are everywhere within Spotify’s app. These features reinforce mainstream genres and artists while limiting discovery of niche content, especially with new features like daily AI playlists, AI DJs, and Smart Shuffle that are constantly being forced upon users (Anderson et al., 2020). Specifically, the AI Smart Shuffle feature is automatically put on every time a user wants to shuffle their own playlists, demonstrating that user’s cannot avoid AI without going out of their way to turn the feature off (Anderson et al., 2020). The underlying consequence is a system where engagement metrics take precedence over artistic diversity and choice, prioritizing engagement metrics that benefit the corporation over the user.

Spotify lacks transparency in how recommendations are generated, how data is used, or how filter bubbles may shape results, resulting in user passivity. Although Spotify offers a website called *Spotify for Developers* in which users can use the Spotify API, the site only shows a user the code on how to get their top 5 tracks, how to save those songs in a playlist, and how to listen to the songs on the website (Qutan, 2025). There is no mention of how their AI recommendation system works, or what an algorithmic filter bubble is. This opacity leads to low user awareness of the technology they are using and paying for, making it difficult for listeners to understand and actively counteract the reinforcement of filter bubbles. Without transparency, users remain passive consumers of AI generated content rather than active participants in shaping their listening experiences and the priorities of the company (Bartlett et al., 2023).

Even if users recognize the limitations and consequences of AI driven recommendations, Spotify is unlikely to change anything about their practices unless there is collective pressure from users for Spotify to create more tools to diversify their listening habits. Without more interactive mechanisms like user input to adjust recommendations or challenge algorithmic decisions, users remain locked into algorithmically curated content loops that limit music discovery and undoubtedly reinforce the dominance of major labels and commercial interests (Espinoza-Rojas et al., 2023). Currently, there is no checkbox on the app or web player that will turn off all AI features indefinitely; there is also no easy and accessible way to leave feedback on the AI content within the app or web player without emailing Spotify (Qatun, 2025). Without such interventions, the personalization strategies used by Spotify remain largely vague and leave users in a passive state where technology shapes them, but they do not shape technology. Accordingly, in its vagueness, Spotify acts as music's cultural gatekeeper.

Although the literature has documented the narrowing effects of algorithmic curation, more information needs to be made public about Spotify's AI. Factors include: how the AI processes user data, structures its ranking systems, and why algorithmic filter bubbles occur within the code (Anderson et al., 2020). It can thus be seen that a particularly difficult aspect of AI and Machine Learning is that the system is inherently a black box–its real-time decision making processes can be often opaque, even to its own designers.

Methods & Results

This paper utilizes both case study analysis and empirical experimental data to examine how Spotify's algorithm impacts user behavior and music consumption. Since Spotify's AI features are relatively new, and the company closely protects the technical details of its recommendation systems, choosing the most informative and credible sources became essential. Specifically, the research required two distinct types of information: a detailed understanding of why Spotify shifted toward AI recommendations, and measurable evidence discussing how personalized recommendations affect user behavior.

First, a comprehensive search was performed to select a case study clearly explaining Spotify's transition to AI. Sid Yu's (2024) study provided this information, outlining Spotify's timeline in adopting machine learning algorithms. Second, quantitative evidence was gathered from a randomized field experiment performed by David Holtz (2020). This experiment tested Spotify users' reactions to personalized recommendations generated by AI. It compared two user groups: a treatment group where users' recommendations were personalized based on their music listening history, and a control group who were recommended popular content among users in their demographic group (Holtz et. al 2020).

To analyze the collected evidence, qualitative details from the case study, such as Spotify's motivations, decisions, and strategic shifts toward using AI, were integrated with the quantitative results from the field experiment data. By combining qualitative historical context with quantitative experimental results, the analysis clarified how Spotify's recommendation system shapes user listening behavior and consumption patterns. ANT emphasizes that technology (such as Spotify's recommendation algorithm) is not passive; instead, it actively interacts with users and social groups to shape behaviors and cultural practices. Listeners, artists, Spotify, and AI algorithms are all thus interconnected and inseparable parts of the larger system of music streaming. Applying ANT allowed this analysis to highlight how Spotify's algorithms not only reflect user preferences but actively influence and reshape those preferences over time.

Examination of evidence revealed several consistent themes regarding Spotify's implementation of AI. The primary themes identified include Spotify's strategic shift toward sophisticated AI-driven personalization to increase user engagement, the resulting tension between maximizing user interaction and maintaining diverse content exposure, and the significant ethical implications of algorithmically-driven recommendations on consumers. The following sections elaborate on these themes closely.

Case A: Timeline and Shift of Priorities

The selected case study described an organizational pivot in 2014 when Spotify began deploying complex machine learning models. Specifically, a multi-layered recommendation algorithm was implemented that processed both user generated data (e.g., listening history, playlist creation, likes) and musical features (e.g., tempo, genre, instrumentation) to deliver personalized playlists (Yu, 2024).

Spotify began using AI extensively after acquiring Echo Nest, a machine learning company, in 2014. The technology from Echo Nest enabled Spotify to categorize user preferences and create detailed listener profiles. Spotify's AI was still being developed in the background until 2015, when AI music recommendation was introduced. In 2015, Spotify then launched *Discover Weekly*, a personalized playlist curated weekly by Spotify based on a user's listening habits (Yu, 2024). The marketing around the feature demonstrated Spotify's dedication toward an advanced, personalized experience, solidifying its place as the number one streaming service (Yu, 2024).

In recent years, the main machine learning algorithm used to power Spotify's music recommendation system is BaRT. This algorithm can be broken down as follows: data collection, segmentation and categorization, and targeting and recommendation (Yu, 2024). In the process, Spotify collects an extensive amount of data from song information like genre and tempo, to various user behaviors such as skipped songs, listening history, demographics of the user, and social interactions. Spotify then channels the individual data into collaborative models where users across the platform are compared to each other extensively in the backend (Yu, 2024). After collecting the data, segmentation and categorization follows. The algorithm categorizes the songs, and then the AI analyzes the sonic data from songs and identifies genres, artists, and titles (Yu, 2024).

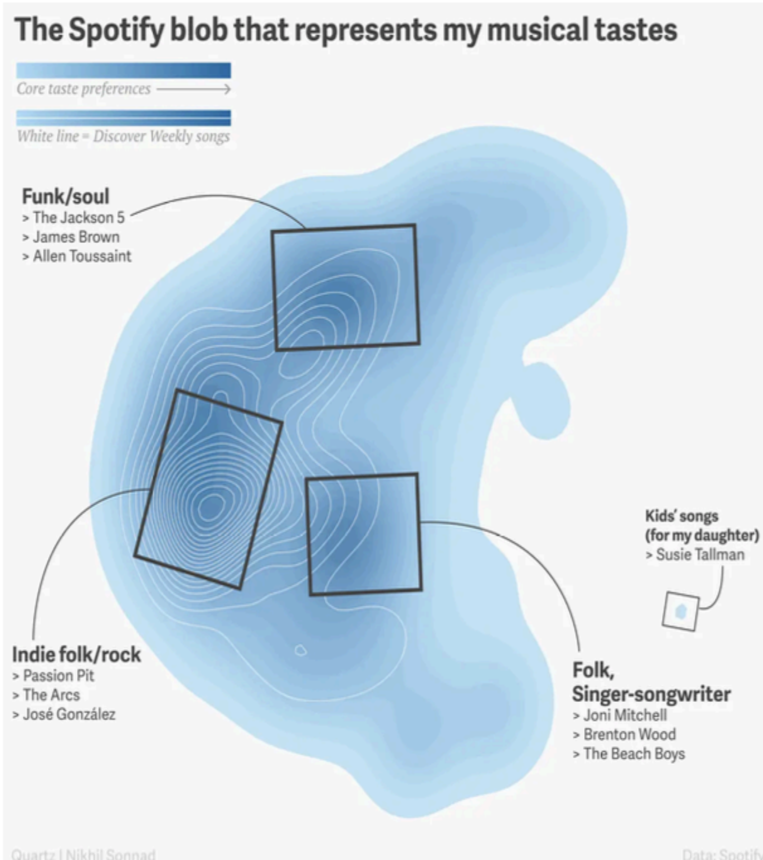

Spotify's musical map (Pasic, 2015)

Figure 1. Spotify's Musical Map (Pasick, 2015)

The model organizes tracks in a map such that its proximity to other songs relates to its similarity to those songs (Yu, 2024). For example, if given two songs that are frequently listened to in sequences are placed close to each other on the map. If multiple tracks are observed to be paired with each other in playlists, then clustering will emerge. The following behavior can be shown (Figure 1).

While Spotify often succeeds in recommending new music that users discover they enjoy, it is important to consider the ethical implications of the algorithm used and what that means for users moving forward. ANT emphasizes that technology and society mutually shape each other through a dynamic network of interactions involving human actors (such as users, artists, developers, and corporate stakeholders) and non-human actors (including algorithms, datasets, and recommendation interfaces). Within Spotify's sociotechnical network, algorithmic design choices and commercial incentives influence user listening habits, artistic visibility, and cultural trends. Ethical considerations thus revolve around the principles of transparency, autonomy, and diversity: transparency in how algorithms function and affect user choice, autonomy in preserving the listener's control over their musical preferences, and diversity in providing equitable exposure for various artists and musical styles. BaRT directly impacts which artists and genres are promoted to users. As discussed previously, any recommendation system powered by machine learning is inherently unfair. Even if the data that Spotify collects seems to be completely fair and free of bias, AI itself is encoded with bias and discriminates. There is always proportionately less data available about minorities – minority genres, artists, and listeners. Accordingly, models about minorities generally tend to be more work than those about the general population as a developer would need to keep these edge cases and special circumstances in mind when designing the algorithm. The data is thus a social mirror: if social biases against a minority group exist, they will be reflected in the training data and incorporated into machine learning algorithms.

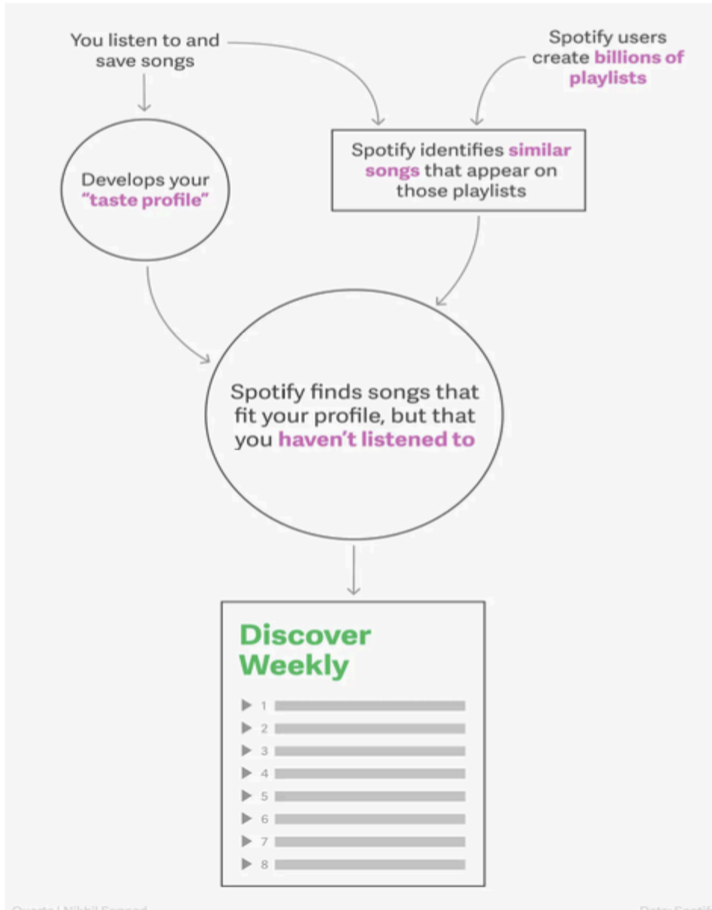

Diagram of music recommendation (Pasick, 2015)

Figure 2. Diagram of Music Recommendation (Pasick, 2015)

Looking at diversity within the musical map and collaborative filtering approach, there are inherently more data points representing the majority group; therefore, the algorithm tends to promote music from the majority in efforts to increase the likelihood that the user is satisfied with the recommendation (Yu, 2024) (Figure 2). Ultimately, Spotify faces the dilemma of utilizing BaRT to boost user engagement and perpetuate the homogeneity of music, or reducing the reliance on BaRT to promote diversity at the expense of user engagement.

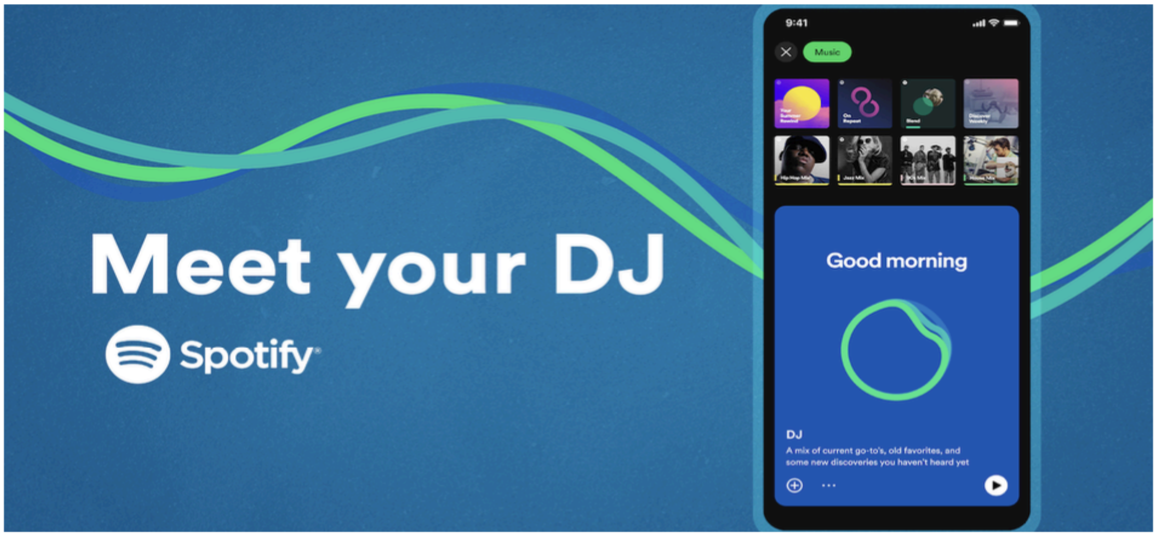

(Spotify, 2023)

Figure 3. Spotify DJ Ad. Advertisement from Spotify for the then-new AI Spotify DJ feature (Spotify, 2023).

Spotify will continue to develop its music recommendation experience with AI, focusing on strengthening the relationship between the consumer and the brand. Spotify aims to achieve this in humanizing the platform with the introduction of Spotify DJ in 2023 (Figure 3). With meeting *your* personalized DJ, the feature extends the BaRT recommendation system interacting with the user in a human-like manner, providing insights to music recommendations, and implementing reinforcement learning by interacting with the user directly (Yu, 2024). Likely, the platform will continue on this route of humanizing AI to increase general acceptance.

Case B: Field Experiment

A review of the field experiment highlights that algorithmic personalization has measurable impacts on user behavior. By analyzing outcome measures such as total streams and skip rates, the experiment reveals a pronounced spike in engagement. The data presentation here is methodical, with results indicating a cause-and-effect relationship between personalization and user activity. The experiment also demonstrates strong internal validity by incorporating a control group in which no AI recommendations were given to the user.

Table 1: Impact of Spotify's AI-Driven Personalization on User Engagement and Diversity

Metric	Effect of personalized recommendations
Average Increase in Total Podcast Streams	+28.90%
Individual-level Content Diversity (Shannon entropy)	-11.51%
Aggregate (Platform-wide) Content Diversity	+5.96%
Home-Screen Recommended Content Diversity	-17.70%
Non-Home ("Organic") Content Diversity	-3.31%

In the experiment, both control and treatment users were given recommendations, with the sole aim of increasing consumption. Treatment users recommendations were personalized with AI based on their listening history, whereas control users were recommended popular content among users in their demographic group. The experiment demonstrates a 28.90% increase in user engagement measured by total streams, indicating that personalized AI recommendations effectively boost user interaction with the platform. However, these recommendations also reduce individual-level diversity by 11.51%, meaning each user's content consumption becomes narrower. Meanwhile, platform-wide content diversity sees a modest increase of 5.96%, indicating users as a whole engage with a broader variety of content, yet individually, they consume less varied content. Additionally, personalized recommendations presented on Spotify’s home screen further reduced diversity by 17.70%, suggesting strong algorithmic influence in shaping user habits. Even outside direct recommendations, users' organic listening choices show 3.31% lower diversity, highlighting the pervasive effect of personalization algorithms on user behavior (Holtz et al., 2020). The results provide evidence of an engagement-diversity trade-off when recommendations are optimized solely to drive consumption.

These findings highlight the need for academics and practitioners to continue investing in personalization methods that explicitly take into account the diversity of recommended content, especially for music streaming. The results of this experiment clearly show how Spotify’s recommendation system significantly increases user engagement at the expense of diversity of consumed content (eg. Podcasts in the Field Experiment. Combining the case study and experimental data highlights how Spotify's AI implementation reinforces existing cultural preferences, amplifying popular or mainstream content. The data suggests that Spotify's focus on maximizing short-term user engagement through AI-driven personalization comes at the potential cost of limiting long-term cultural diversity and user autonomy.

Conclusion

Spotify’s increased reliance on AI-driven personalization significantly enhances user interaction, but it simultaneously reduces individual exposure to a diverse range of musical content, thereby creating isolated user filter bubbles. This research identifies clear measurable impacts: while personalized algorithms boost short-term engagement and user satisfaction, they concurrently limit long-term musical diversity, shaping user preferences towards mainstream and commercially successful content. Such practices conflict directly with Spotify’s stated mission to “unlock the potential of human creativity-by giving a million creative artists the opportunity to live off their art and billions of fans the opportunity to enjoy and be inspired by it.” (Spotify, 2015). This mission statement was created ten years ago, before AI. Spotify’s initial mission statement no longer holds value as Spotify’s current practices prioritize reinforcing existing tastes and increasing corporate revenue. These practices create filter bubbles, closing off the user to only encounter familiar and predictable content. Such algorithmically-driven curation restricts opportunities for emerging, niche, or diverse artists, directly opposing Spotify’s stated mission to promote human creativity by connecting artists and fans on a broad scale. While Spotify’s original mission emerged five years prior to its extensive AI integration, the rapid adoption of advanced machine learning technologies necessitates reconsideration of how algorithms align with, or challenge this vision.

Moving forward, research can look further into the role of transparent user controls and whether or not they would be generally accepted. In a landscape where AI only continues to gain widespread adoption by both companies and consumers, it is critical to examine how these technologies will impact the long-term experiences of users and what voices are deemed worth listening to.

References

About Spotify. Spotify. (2025, February 7). https://newsroom.spotify.com/company-info/#:~:text=Our%20mission%20is%20to%20unlock,and%20be%20inspired%20by%20it

Anderson, A., Maystre, L., Mehrotra, R., Anderson, I., & Lalmas, M. (2020). Algorithmic effects on the diversity of consumption on Spotify. Proceedings of The Web Conference 2020 (WWW '20), 2155–2165. https://doi.org/10.1145/3366423.3380281

Bartlett, M., Morreale, F., & Prabhakar, G. (2023). Analysing privacy policies and terms of use to understand algorithmic recommendations: The case studies of Tinder and Spotify. *Journal of the Royal Society of New Zealand, 53*(1), 119–132. https://doi.org/10.1080/03036758.2022.2064517

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

Bozdag, E. (2013). Bias in algorithmic filtering and personalization. *Ethics and Information Technology*, *15*(3), 209–227. https://doi.org/10.1007/s10676-013-9321-6

Espinoza-Rojas, J., Siles, I., & Castelain, T. (2023). How using various platforms shapes awareness of algorithms. *Behaviour & Information Technology, 42*(9), 1422–1433. https://doi.org/10.1080/0144929X.2022.2078224

Holtz, D., Carterette, B., Chandar, P., Nazari, Z., Cramer, H., & Aral, S. (2020). The engagement-diversity connection: Evidence from a field experiment on Spotify. In

Proceedings of the 21st ACM Conference on Economics and Computation (pp. 75–76). ACM. https://doi.org/10.1145/3391403.3399532

Interian, R., Marzo, R. G., Mendoza, I., & Ribeiro, C. C. (2023). Network polarization, filter bubbles, and echo chambers: An annotated review of measures and reduction methods. *International Transactions in Operational Research, 30*(6), 3122–3158. https://doi.org/10.1111/itor.13224

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

Pasick, A. (2015, December 21). *The magic that makes Spotify's discover weekly playlists so damn good*. Quartz. https://qz.com/571007/the-magic-that-makes-spotifys-discover-weekly-playlists-so-damn-good

Prey, R., & Esteve-Del-Valle, M. (2023). The algorithmic network imaginary: How music artists understand and experience their algorithmically constructed networks. *The Information Society, 40*(1), 18–31. https://doi.org/10.1080/01972243.2023.2274070

Qatun, S., Sandhya, K., & Sree, K. N. (2025). Spotify API: Fetching artist details like popularity, followers, top tracks, and its details. *Recent Trends in Information Technology and its Application, 8*(2), 14–22. https://doi.org/10.5281/zenodo.14998645

Yu, S. (2024.). *The engagement-diversity connection: Evidence from a field experiment on Spotify* [Case study]. Stanford Students in Entertainment.

https://www.stanfordssie.org/s/Sid-Yu-SSIE-Winter-Research.pdf