TEMPO: A PERSONALIZED AUDIO EXPERIENCE

BEYOND THE PLAYLIST: EXAMINING THE ETHICAL DIMENSIONS OF MUSIC RECOMMENDATION ALGORITHMS IN MEDIA CONSUMPTION

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

ISABELLA HEINTGES

November 8, 2024

Technical Project Team Members

Joey Cohen, Naomi Solomon, Michelle Monge, and Thomas Keathley

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

Rider Foley, Department of Engineering and Society

Adam Barnes, Department of Electrical and Computer Engineering

INTRODUCTION

Society is currently experiencing a significant period of digitization, with many physical systems and artifacts becoming increasingly accessible online. In the realm of media, musicians have capitalized on this shift, enabling them to share their work across various platforms with their audiences. As a result, media users now have the opportunity to stream, purchase, or download artistic creations directly from the internet. This act of publicly digitizing audio for user-playback has transformed revenue generation in the United States music industry, with streaming now accounting for 84% of the total earnings in 2023, quickly growing about 20% in six years (Statista, 2024) (see Figure 1).

Figure 1

Distribution of Music Industry Revenue in the United States by Source

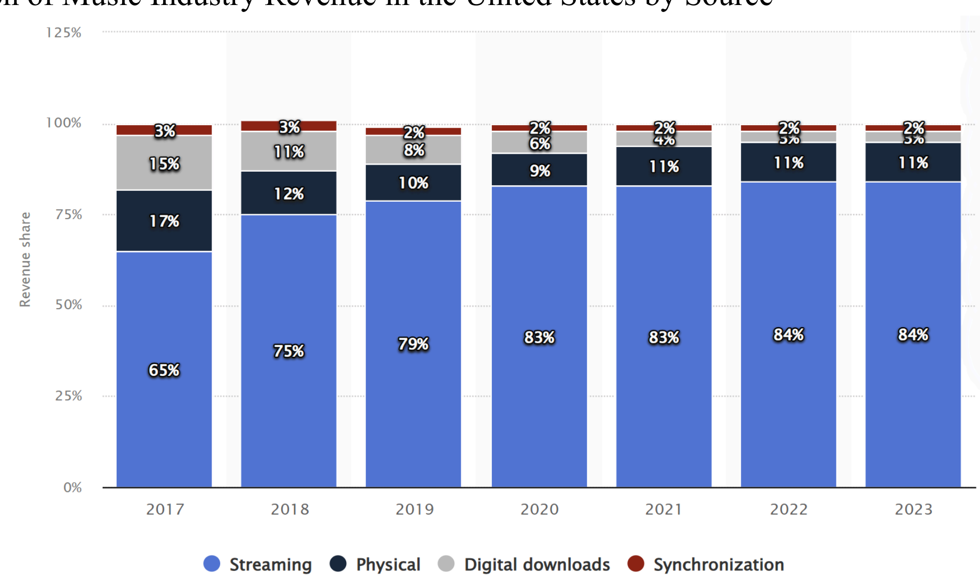

Among these streaming platforms that permit the transmission of open-source audio, Spotify dominates the market with about a 30% share, double that of its closest competitors (Turner, 2024). While the act of streaming involves the same process across all platforms, Spotify's success is directly attributable to its personalized recommendations, cultivated by its complex algorithm designed to "tailor to each listener's unique taste" (Spotify, 2024). Unlike other domains, Spotify's algorithms take into account each listener's searching, listening, skipping, or saving habits, while additionally processing personal information related to geographic location and age. While these recommendation algorithms foster a positive, individualized listening experience in the music suggestion process, the systems themselves raise significant ethical concerns in a broader context, based on their development and implementation.

At their baseline, recommendation algorithms are a subset of machine learning, now implemented in practically every online domain. In practice, these systems leverage data provided by the user, both implicitly and explicitly, to assist in the search for new content (Shetty, 2019). They utilize a filtering process specific to their application, but all function by retaining profiles of a user's long or short term activity, or stated preferences (Schafer et al., 2007). Recommendation algorithms are used in a variety of applications apart from the aforementioned personalization process offered by Spotify, with websites and social media platforms utilizing their output to "guide the user in a personalized way to interesting or useful objects in a large space of possible options" (Burke et al., 2011, pp. 361). Their processes differ from general search engines, as the results relevant to the same input differ based on who interacts with the system. Generally, by tailoring suggestions to individual tastes, these systems make users feel understood and valued. However, their continued integration online has brought to light an algorithm's "capacity to shape social and cultural formations and impact directly on individual lives" (Beer, 2009, pp. 12).

With the rise of streaming and the growing popularity of recommendation systems—particularly Spotify's successful personalization features—my technical project focuses on developing an iOS application that tailors music suggestions based on a user's query and streaming preferences. This app will connect to an interactive Bluetooth speaker designed for high-quality audio transmission, delivering a personalized listening experience that enhances enjoyment of recommended music and displays relevant song information on the device. My project integrates advanced recommendation algorithms that adapt music suggestions based on each user's preferences and habits. Given the powerful role of recommendation systems in shaping media consumption, my goal is to design algorithms that prioritize both efficiency and fairness. I will also explore how recommendation algorithms can offer relevant content while remaining transparent and ethically responsible, addressing potential concerns about privacy, manipulation, and user well-being in today's increasingly algorithm-driven digital landscape.

TEMPO: A PERSONALIZED AUDIO EXPERIENCE

My technical project can essentially be divided into two main components: **(1)** a Bluetooth-enabled speaker with a capacitive-touch LCD display, and **(2)** a social audio-recommendation iOS application that integrates the Spotify API. The speaker is designed for high-quality audio output, incorporating a low-noise audio amplifier, two 8-ohm full-range drivers, and a Bluetooth extension board equipped with high-resolution digital-to-analog converters. Additionally, a microcontroller will manage the functionality of the LCD display, which provides users with touch control for playback directly on the device. Complementing this hardware, the iOS application will generate music recommendations based on each user's unique preferences, including keyword inputs and past listening habits. Beyond recommending music, the app will feature a social media aspect, allowing users to connect, follow others, and view shared music history and preferences. The following sections will delve into the technical specifications of both the Bluetooth-enabled speaker and the social audio-recommendation iOS application, outlining each component in greater detail.

(1) Bluetooth Speaker with a Capacitive-Touch LCD Screen

This Bluetooth-enabled speaker is designed to provide high-quality audio and user-friendly functionality, with integrated controls for power, volume, and playback (see Figure 2). At its core, the system is powered by the STM32F769I-DISCO microcontroller, chosen for its efficient processing abilities, operating at up to 216 MHz. This microcontroller handles Bluetooth connectivity and data signals, manages user information through a Bluetooth module, and operates a capacitive-touch LCD screen. The screen displays real-time song information, such as the title and duration, while allowing users to control playback directly on the speaker.

To maintain a smooth audio experience, the system synchronizes data transfer with the Bluetooth module, reducing any potential delay. This module sends audio signals to a low-noise, Class-D amplifier that drives two 2-inch full-range speakers, providing a balanced, high-fidelity sound output (Iversen et al., 2015). The amplifier is tuned to minimize distortion, keeping Total Harmonic Distortion[1] (THD) below 10%, and reproduce a wide audio range, from deep bass to clear highs, without needing additional woofers or tweeters. Each audio channel offers a 50W output at 8 Ohms, delivering ample volume and clarity for various listening environments.

For reliable power, the speaker is equipped with an AC/DC adapter that supports standard outlets and converts 110–220 VAC to 12V, providing up to 150W of power. Users can control the speaker's operation using an onboard power switch, which allows easy toggling between active and standby modes. LED indicators provide visual feedback on the device’s status, enhancing the user experience with intuitive, at-a-glance interaction.

---

[1] A measurement used in audio and electronic systems to quantify the amount of distortion added by a device, typically an amplifier or speaker, to the original signal. Distortion occurs when the system introduces unwanted harmonics (multiples of the original frequency) into the output sound (Williams, 2017).

Figure 2
Design and Arrangement of the Speaker Components

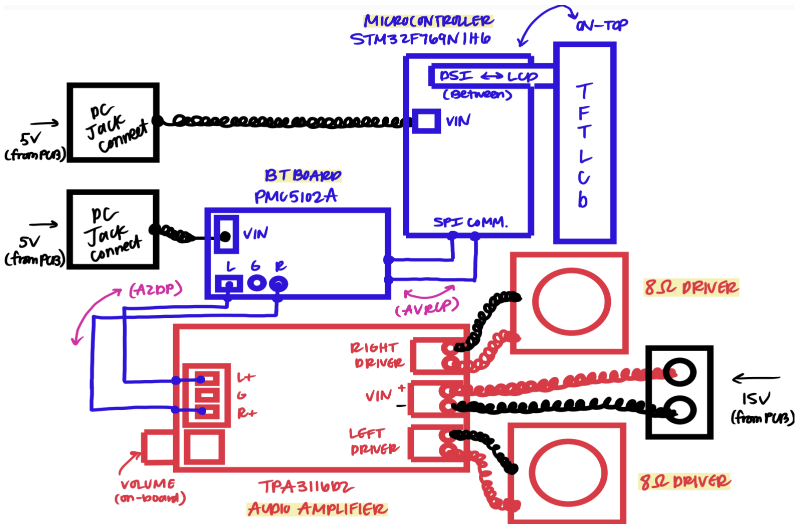

The enclosure will be constructed from composite wood, specifically Medium-Density Fiberboard (MDF), due to its excellent vibration absorption properties and density, which helps prevent unwanted resonance (DIY Audio & Video, 2024) (see Figure 3). MDF is a preferred material for its durability and acoustic qualities, contributing to overall sound quality. However, careful consideration will be given to ensure that any improvements in audio performance do not compromise the enclosure's durability. A passive radiator vibration membrane will be mounted in the center of the enclosure to enhance sound clarity.

Figure 3
Three-Dimensional Model of the Speaker's Enclosure Design

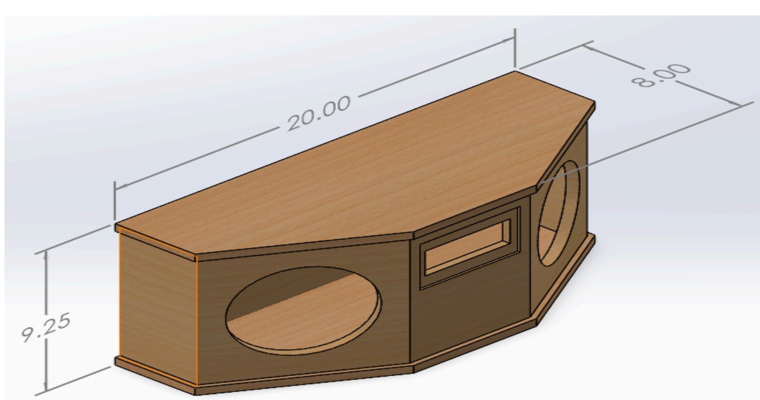

(2) Audio Recommendation iOS Application Integrated with Spotify API

Tempo, the iOS application, aims to aid in music discovery and social connectivity by delivering a highly personalized and engaging user experience (see Figure 4). Central to its functionality is a music recommendation engine that uses machine learning to analyze user listening habits, preferences, and social interactions. This system ensures that users receive music suggestions specific to their tastes, and dynamically adapt based on their evolving musical interests.

Figure 3

Wireframe Design of the Tempo Application

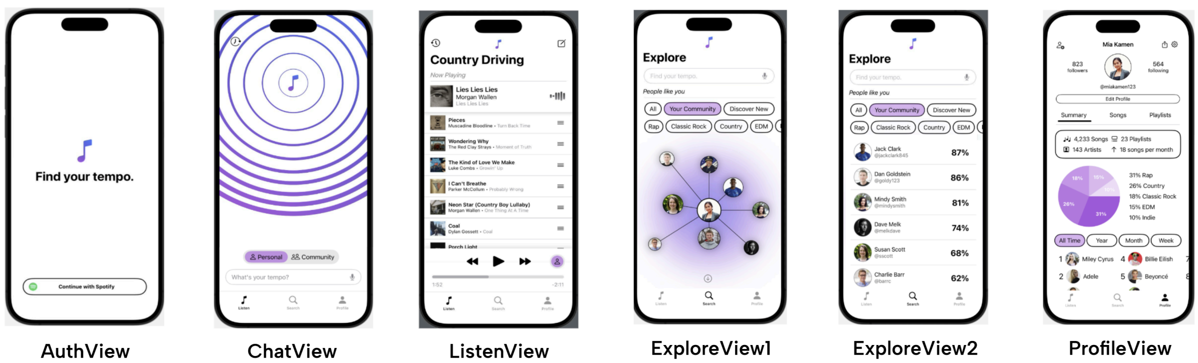

The music recommendation algorithm will implement a pipeline that integrates OpenAI's GPT with Spotify's Application Programming Interface (API)[2], to fulfill musical user requests based on prompts. This API allows us to leverage statistical data automatically collected by Spotify, such that our application can interact directly with the service and permits the retrieval of user metadata, management of playlists, and control of playback (Spotify). The system will then use a Generative Pre-Trained Transformer (GPT)[3] for natural language processing and sentiment analysis to parse the prompt and extract structured data including the mood, artist,

[2] A set of rules and protocols that allows different software applications to communicate with each other. It defines how requests and responses should be structured, what data is needed, and what functions or services are available from one software system to another (Spotify).

[3] A type of artificial intelligence language model developed by OpenAI. It is designed to understand and generate human-like text based on a given prompt (Amazon).

genre, and user-preferences (Narducci, 2020). To effectively map these moods to musical features, a predefined dictionary will translate common themes into genres and features. These themes are further converted to parameters for the Spotify API, consisting of seed genres and target audio features.

With these parameters identified, the system will authenticate through Spotify API using O-Auth[4] tokens to retrieve the relevant music data. This involves converting the artist names to Spotify artist IDs, and then aligning the genre codes with predefined genres. The system will then formulate specific API queries based on user intent: playing songs by specific artists, finding similar travels, and/or generating playlists that match the users goal. Further, the system will implement caching and batch API requests to improve the applications efficiency and reduce the server load (Kulkarni et al., 2023).

Overall, in integrating OpenAI's GPT and Spotify's API, the recommendation algorithm pipeline creates a system that will interpret simple or complex user prompts, map prompts to attributes, and will generate a queue that is tailored to the user's preferences, allowing for an interactive experience alongside the physical speaker.

Relation to the Broader Context

Spotify's recommendation process has received criticism from listeners, with frequent complaints about repetitive suggestions that prioritize recently played tracks and a lack of transparency in how user statistics are calculated. As former Spotify subscriber Chayka (2024) noted, "The platform interface has gradually made it harder to find the music I want to listen to. With the latest app updates, I'd had enough." Our technical project aims to improve user engagement by offering intuitive controls and a personalized music experience that encourages a

[4] A type of security token used in the OAuth protocol, which is an open standard for authorization. OAuth allows a user to grant third-party applications limited access to their resources without sharing their passwords (OAuth).

more meaningful connection to users' musical tastes. Addressing the limitations of Spotify's recommendation system, our algorithm will introduce a sophisticated scoring method to deliver more contextually relevant suggestions. Additionally, the application will adhere to data privacy regulations to protect user information and clearly outline how user statistics are generated. As recommendation systems grow more common in media consumption, society is becoming more wary of their influence.

STS ANALYSIS OF THE RECOMMENDATION PROCESS

Recommendation algorithms play an active role in shaping how society consumes media and continually influence users' thoughts, actions, and perspectives. Widely used in platforms like Spotify, these systems not only enhance the user experience but also exert significant influence over user behavior, often in ways that are neither transparent nor fully understood. As these algorithms become more integrated into media consumption habits, it becomes increasingly important to consider the ethical implications of their design, and how their integration has changed existing user-domain dynamics. The balance between the efficiency of content delivery and fairness in representation is a key challenge in addressing these concerns and ensuring that these systems serve users ethically (Zarsky, 2016).

Audience manipulation is a significant ethical concern arising from the use of recommendation systems, particularly in the context of misinformation, polarization, and bias. These algorithms are designed to predict and suggest content based on users' past behaviors, essentially reinforcing one's preferences to create a personalized experience. This customization, while sought after, often results in the creation of *filter bubbles*, "whereby users are only recommended content narrowly aligned with their historical interests" (Gao et al., 2022, pp. 423). As a result, user-suggested content becomes skewed, contributing to the spread of misinformation, while entrenching users in e*cho chambers*: spaces "that only echo(es) their own views and beliefs," (Bojic, 2024, pp. 104). Ethically, this raises critical questions about user autonomy: Are users making informed, independent choices, or are they being steered into consuming content that serves the interests of the platform? The challenge lies in striking a balance between the efficiency of recommendation algorithms and the fairness of content delivery.

Recommendation algorithms rely heavily on the collection of vast amounts of user data, including listening habits, search histories, and demographic information, to create accurate predictions of user preferences and enhance recommendations. While this data-driven personalization improves the user experience, it also brings forth significant privacy concerns. The non-removability of networked data, the monetization of personal attention, and the constant surveillance of human behavior raise questions about the ethics of data collection (Gao & Yu, 2024). Users often have little awareness of the extent to which their information is being gathered, shared, and utilized. Moreover, the inability to guarantee anonymity in these systems reinforces information asymmetries and power imbalances, posing a challenge to traditional notions of privacy (Mccarthy, 2017). As users' data is commodified and networked, they may lose control over their personal information, undermining trust in the platform. The ethical dilemma here lies in balancing the efficiency of data collection, which improves recommendation accuracy, with the responsibility to safeguard user privacy. Developers must prioritize data protection by providing transparency about data usage and offering users meaningful choices regarding what they share (Martin, 2019). This ensures users can make informed decisions about their privacy while still enjoying the benefits of personalized content.

Given these ethical considerations, I will draw upon "The Ethnography of Infrastructure" as defined by Susan Leigh Star to conduct my research (Star, 1999). Infrastructure, as a concept, is both relational and ecological. It refers to the systems, structures, and technologies that serve as foundational elements within a particular environment, but its meaning can vary depending on the context and the groups involved. In many cases, infrastructure is embedded within a given setting, making it often invisible and taken for granted by those who interact with it. When studying information systems through an ethnographic lens, infrastructure is not just about physical or technical structures but also the social and cultural practices intertwined with them. This framework is particularly well-suited for further exploring recommendation systems in terms of the changing dynamics of media consumption.

RESEARCH QUESTION AND METHODS

The research question that I plan to explore further is: *How have recommendation algorithms impacted the dynamics of social media consumption, and how might these systems be better designed to ensure fairness and efficiency?* This question is crucial in the context of recommender systems, given that the guiding algorithms have become central to content delivery on all platforms, influencing the media that users see and how they interact with the content. Understanding their effects on user behavior, privacy, and the potential for manipulation is essential to developing ethical systems with greater transparency.

I plan to conduct my analysis using a mixed-methods approach, similar to the processes outlined by Anandhan et al. (2018), to explore the evolving dynamics of media consumption and examine potential strategies for addressing ethical considerations. To assess the effects of recommendation systems on audience manipulation and data privacy, I will explore the platforms of TikTok and Instagram.

The first step will be data collection, where I will gather user-generated content such as posts, hashtags, and interaction metrics (likes, comments, shares, and views) through APIs provided by these platforms, alongside implementing available third-party tools. The next phase will involve content analysis, where I will evaluate the posts and hashtags on both platforms. This will include categorizing types of content and identifying trends in how algorithms promote specific topics, influencers, or content genres. I will also assess how engagement metrics, such as the number of likes, comments, and shares, correlate with the content recommended by the algorithms, exploring whether certain content (e.g., viral trends or influencer posts) is disproportionately promoted. I hope to gather insights to understand how the recommendation systems shape user behavior and attention, possibly leading to user manipulation.

To provide historical context to my research, I will conduct a literature review examining the evolution of recommendation algorithms from their early implementation to their current role in shaping media consumption. This review will highlight changes in algorithmic complexity, transparency, and ethical considerations over time.

Analyzing this framework using Star's framework of infrastructure, will offer a unique approach in understanding how the integration of recommendation algorithms have transformed media consumption and the ethics surrounding their increasing presence in daily life. Star argues that infrastructure operates as "a relational property, not as a thing stripped of use," highlighting the importance of exploring this technology not in isolation but in relation to the people, institutions, and other systems it interacts with (Star, p. 113). This perspective will allow for a deeper analysis of recommendation systems, focusing on how their design influences fairness, privacy, and user autonomy.

CONCLUSION

Recommendation algorithms significantly shape media consumption by personalizing content, but their widespread use also raises critical ethical concerns, including audience manipulation, data privacy, and bias reinforcement. The project’s technical deliverable involves developing an iOS application integrated with an interactive Bluetooth speaker for personalized music recommendations. Generally, the research in this paper will explore the effects of recommendation systems on social media dynamics and propose ways to enhance fairness and efficiency. Using a mixed-methods approach, the study will examine the evolving dynamics of media consumption by analyzing quantitative data, such as interaction metrics and engagement levels, alongside qualitative insights from content analysis, user surveys, and a literature review on the historical development of recommendation algorithms. Additionally, Susan Leigh Star's infrastructure framework will be applied to examine the relational dynamics between algorithms, users, and broader systems.

REFERENCES

Anandhan, A., Shuib, L., Ismail, M. A., & Mujtaba, G. (2018). Social Media Recommender Systems: Review and Open Research Issues. *IEEE Access*, *6*, 15608–15628. https://doi.org/10.1109/access.2018.2810062

Beer, D. (2009). Power through the algorithm? Participatory web cultures and the technological unconscious. *New Media & Society*, *11*(6), 985–1002.

Bojic, L. (2024). AI alignment: Assessing the global impact of recommender systems. *Futures*, *160*, 103383. ScienceDirect. https://doi.org/10.1016/j.futures.2024.103383

Burke, R., Felfernig, A., & Göker, M. H. (2011). Recommender Systems: An Overview. *AI Magazine*, *32*(3), 13. https://doi.org/10.1609/aimag.v32i3.2361

Chayka, K. (2024, July 31). Why I Finally Quit Spotify. Retrieved from The New Yorker website: https://www.newyorker.com/culture/infinite-scroll/why-i-finally-quit-spotify

DIY Audio Speaker Box Building Guide. (2024). Retrieved from DIY Audio & Video website: https://www.diyaudioandvideo.com/Guide/BuildSpeakerBox/

Gao, D., & Yu, D. (2024). Challenges and Cracks: Ethical Issues in the Development of Artificial Intelligence. *Science, Technology & Society*, *29*(3). https://doi.org/10.1177/09717218241246372

Gao, Z., Shen, T., Mai, Z., Bouadjenek, M. R., Waller, I., Anderson, A., … Sanner, S. (2022). Mitigating the Filter Bubble While Maintaining Relevance: Targeted Diversification with VAE-based Recommender Systems. *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2524–2531. https://doi.org/10.1145/3477495.3531890

Iversen, N. E., Knott, A., & Michael. (2015). Low Impedance Voice Coils for Improved Loudspeaker Efficiency. *Welcome to DTU Research Database*, 9389. Retrieved from https://orbit.dtu.dk/en/publications/low-impedance-voice-coils-for-improved-loudspeaker-efficiency

Kulkarni, A., Mahajan, P., Nimbokar, G., & Tupe, U. (2023, June 23). AI Based Song Recommendations System. Retrieved November 8, 2024, from ResearchGate website: https://www.researchgate.net/publication/371804043_AI_Based_Song_Recommendations_System?__cf_chl_tk=LzMYFkmmP_9XfWg.O.2JHz_Eo73X2FYYhvV0e6Flt0U-1731093181-1.0.1.1-NUID7QfpaGxQAj.DY.oLK8qy0ieKOTCooe2l6Zqfrp0

Martin, K. (2019, June 1). Designing Ethical Algorithms. Retrieved from papers.ssrn.com website: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3056692

Mccarthy, M. T. (2017). The Semantic Web and Its Entanglements. *Science, Technology and Society*, *22*(1), 21–37. https://doi.org/10.1177/0971721816682796

Narducci, F. (2020). Evaluating ChatGPT as a Recommender System: A Rigorous Approach. Retrieved November 8, 2024, from Polytechnic University of Bari website: https://arxiv.org/html/2309.03613v2

Schafer, J. B., Frankowski, D., Herlocker, J., & Sen, S. (2007). Collaborative Filtering Recommender Systems. *The Adaptive Web*, *4321*, 291–324. https://doi.org/10.1007/978-3-540-72079-9_9

Shetty, B. (2019). An In-Depth Guide to How Recommender Systems Work. Retrieved from Built In website: https://builtin.com/data-science/recommender-systems

Spotify. Web API | Spotify for Developers. Retrieved from developer.spotify.com website: https://developer.spotify.com/documentation/web-api

Star, S. L. (1999). The Ethnography of Infrastructure. *American Behavioral Scientist*, *43*(3), 377–391. https://doi.org/doi/10.1177/00027649921955326

Statista. (2024, June 7). U.S. music industry - revenue distribution 2017-2023, by source. https://www.statista.com/statistics/186304/revenue-distribution-in-the-us-music-industry/

Turner, A. (2024, April 8). Spotify users: How many people have Spotify? (2024). BankMyCell. https://www.bankmycell.com/blog/number-of-spotify-users/

Understanding recommendations on Spotify. (2024). Retrieved November 6, 2024, from Spotify website: https://www.spotify.com/us/safetyandprivacy/understanding-recommendations

What is an Access Token - OAuth 2.0. Retrieved from oauth.net website: https://oauth.net/2/access-tokens/

What is GPT AI? - Generative Pre-Trained Transformers Explained - AWS. Retrieved from Amazon Web Services, Inc. website: https://aws.amazon.com/what-is/gpt/

Williams, D. (2017, February 20). Understanding, Calculating, and Measuring Total Harmonic Distortion (THD). Retrieved from Allaboutcircuits.com website: https://www.allaboutcircuits.com/technical-articles/the-importance-of-total-harmonic-distortion/

Zarsky, T. (2016). The Trouble with Algorithmic Decisions. *Science, Technology, & Human Values*, *41*(1), 118–132. https://doi.org/10.1177/0162243915605575