Whiplash
(Technical Project)

Music Recommendation Software's Impact on New Artists
(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 Computer Engineering

By
Max McCullough

December 8, 2021

Leonardo Anselmo, Uriel Gomez Ibarra, John Lilly, Davis Lydon

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

MC Forelle, Department of Engineering and Society

Harry Powell, Charles L. Brown Department of Electrical and Computer Engineering

I. Introduction

In 2021 over 988 billion songs were streamed in the United States alone (Billboard, 2021, p. 3). Around two-thirds of that music was played on the top three streaming platforms: Spotify, Apple Music, and Amazon Music (MIDiA, 2021). These services have provided music listeners with a great way to have affordable access to seemingly unlimited music. However, these listeners may not be entirely aware of how the types of music they are listening to are being affected. Music recommendation software is used by nearly all popular music streaming services to sort music and provide recommendations to users. One of the potential issues with this is the ability of these services to unknowingly transfer bias on the recommendations based on how the software is designed. This means that a substantial portion of who is determining popular music is not only these companies and their patrons but also the software developers who design the music streaming platforms.

Historically most of the power of deciding who would go on to become popular artists was given to recording companies because they had all the resources required to distribute the music. This included the studios, production equipment for vinyl and CDs, and they even held the advertising departments (Lopes, 1992, p. 70). As time went on, we saw the introduction of new technologies that have slowly taken away some of the power of big record companies. Changes in how radio stations were distributed nationally caused record companies to shift their focus from recording their own musicians' music to seeking out artists and focusing on the production and distribution side. They relied more on searching for new artists on the rise to promote instead of continuing to promote their existing artists (Lopes, 1992, p. 56). The next major shift was the introduction of the internet in the 1980s which then paved the way for music streaming services as we know them today. This began in 1999 with Napster and continued to evolve throughout the 2000s (Dowling, 2019).

In addition to music recommendation software, popular streaming services have made it easier than ever for amateur artists to share their music with the world. Additionally, recording equipment has become increasingly accessible to musicians due to decreases in prices over the last few decades (Recording Connection, 2022). Unfortunately, that does not mean these artists can succeed right out of the gate. Since recommendation software heavily relies on user data, they are unable to properly recommend songs until it has been listened to by a variety of different users. If these systems continue to work in this same manner, new artists will never truly have equal footing when compared to artists backed by huge production companies. Machine learning based systems that use digital signal processing have the potential to create a change in how this works. Systems that use this technique can process the audio of songs and find similarities in various characteristics of the music so that they can suggest songs before anyone has even listened to them (Nikki & Craig, 2020, p. 106). My technical project also uses digital signal processing to break down music. However, instead of using the output of this digital signal processing being used to recommend music, it plays a drum along with the music.

II. Technical Topic

Currently, several working models of robotic drum-playing machines exist in a variety of different physical forms, ranging from purely functional to humanoid and even somewhere in between. One of the most prominent robotic-humanoid drummers is named “Stickboy” (Barnes, 2007). Stickboy is a six-armed robot that has been on tour since 2007. It has no AI but can play a full drum set. Therefore, while Stickboy may be much more capable of playing than any drummer, it is also incapable of creating its own music. Stickboy must also be programmed to play as it has no internal algorithm to play autonomously to a specific song. On the other end of the spectrum, the Polyend Perc (Polyend, 2015) is a less aesthetically focused device capable of playing an entire drum set as well. The device has aluminum enclosures that contain small cylinders to “beat” the drum, rather than strike it. Similarly, this device has no algorithm to automatically track music and play along, as it must be entirely programmed to the user’s specifications.

While both Stickboy and Polyend Perc fill their niche of being able to showcase drumming capability, neither can play without being programmed to a song’s specifications. Our project, Whiplash, utilizes a digital signal processing (DSP) element to autonomously track the beat of a song and play a physical drum to that beat (Anselmo et al., 2022). Rather than hard-code each drum hit as the current technology does, we can focus on making our project generalizable to any musical audio. The chipset we have chosen is a Texas Instruments (TI) Mixed Signal Processing 432 (MSP432) chip. We are also using a TI CC3220SF launchpad (Texas Instruments, 2018) that enables us to easily access each input/output of the chip. The user plays audio from their preferred device (phone, computer, etc.) and the signal is passed to the microcontroller. The microcontroller then performs the DSP to parse out the song's beat. The beat signals are then transformed internally into pulse-width modulated (PWM) signals that can be read by servos, specifically HS-805BB servos (HiTEC, n.d.). The generated PWM signals are sent to a servo driver printed circuit board (PCB) created to protect the MSP432 and servos from overcurrent. The servo driver amplifies the signals and sends them off to the servos, enabling them to spin the drumsticks and hit the drum.

The device will have several other features. The strength of the drumstick hit will depend on how loud the song is playing at the moment of that hit. The device will also have an output to a speaker, so the user is able to hear the music live. The output speaker will have a delay when playing audio so the device can hit the drum at the exact moment it is expected to. The physical body of the device was designed using FreeCAD (FreeCAD, 2002) and printed using a 3D printer. The body has mounts for the servos, while the drumsticks will be screwed into a mounting plate on the servos themselves.

Whiplash helps in solving the previously stated problem by providing a non-human musical companion to practice with. The user will get the benefits of physical feedback while being able to practice as long as they want. The main benefit of the device is the ease of use. A user can simply play any song they would like directly into the device and then continue to play along. Whiplash will handle the rest of the work, calculating the beat and playing autonomously to any song without the user having to think about hard coding in drum hits. Research and development of the device has revealed technologies that can be applied to other fields of research. A common side-effect of research is that it typically leads to revelations that are later applied to wildly separate fields (Cetinic & She, 2022). For example, the beat detection algorithm enables humans to dive deep into the digital bits of a song, allowing us to further understand why we enjoy music on a technical level.

While the device provides an extra musician where there may not be one, it also provides context to the human and social dimensions of technology in general. Whiplash is a bridge technology that may eventually expand to devices with onboard AI, capable of generating their own algorithmic compositions. Currently, the device is nowhere near capable of fully replacing a human musician; however, if the project were to be expanded and researched further, it may very well be. The device can be reformatted to parse out individual drumbeats in a song or can be "taught" by machine learning algorithms to generate its own drumbeats. This would allow the device to connect to a full-scale drum set and create and play music better than a human could. This issue then expands into technology's effect on musical creativity and general creativity.

III. STS Topic

To understand how music recommendation software's effect new artists in the music industry it is important to define the general technologies that are encompassed within music recommendation software. The first variety is the main method used today. Since Spotify is the most popular music streaming service in the United States (MIDiA, 2021), it will be used as an example, but many other large streaming services use these methods for their music recommendation. These systems use the history of music that an individual user listens to and compares it to other users with similar preferences to find who would enjoy that music. In an article published by Spotify, they show that this is still able to recommend music from all levels of popularity (Anderson et al., 2020). Although this system is effective for music that has a small number of listens, it still leaves room for improvement for music that has never been heard or only a couple of listens. This is where using machine learning that uses digital signal processing to recommend music comes in. By using the acoustic quality of the music instead of the current user listening history it can help the underrepresented musicians that have almost no listens. This technique uses things like tempo, volume, and frequency analysis to compare music and record music (Purwins, et al., 2019). Despite the vast potential of this music recommendation technique is seldom used in modern systems due to its greater use of computational resources and difficulty to implement (Schedl, 2019, p. 1). -

To help break down the impact that music recommendation system has had on new artists in the music I plan on using Actor-Network Theory, particularly paying attention to Bruno Latour's work. Actor-Network Theory combines both human and nonhuman actors to create an Actor-Network to help break down how they affect each other. When Actor-Network Theory is applied it assigned notable actors to a node within the network and bridges them together with how they affect each other. This a powerful tool when writing about how a specific technology works because you can use a visual representation of a network to support the claims being made about the way a particular technology has been shaped by society. (Latour, 1992). Since music recommendation systems are entirely focused on predicting human behavior in the form of which songs they will like, Actor-Network is a perfect fit.

Another, notable feature of Actor-Network Theory is that it can be applied over time. Since I am looking at how new artists' ability to gain popularity has changed, I can compare how the Actor-Networks have changed. I can establish the key Actor-Networks including the most important categories of listeners, large artists, new artists, streaming companies, and music producers. The reason why it is essential to use Actor-Network Theory to analyze how music recommendation systems have affected new artists' ability to gain popularity is that there are many factors contributing to why new artists succeed. For example, one focus of smaller artists is to perform at local shows and sell physical copies of their music to get their name out and gain a closer connection with their fans (Hoare et al., 2014). This will need to be accounted for when forming the Actor-Network that my argument will be based on because although this affects how an artist is growing in popularity that cannot be explained by music recommendation systems. If I ignore instances like this in my Actor-Network, I will not get the entire picture and may conclude about the effectiveness of music recommendation systems that were caused by external factors.

Next, I can analyze the listener and the software itself. The primary reason of adding music recommendation system to a music streaming platform is to impact the choices that the user makes. Steven Woolgar's "Configuring the user: the case of usability trials" provides a in depth analysis of the relationship of the user to technology that can be translated to his same situation (Woolgar, 1991). I would particularly draw upon the idea of "Configuring the User", which shows how when the developers behind a technology, in this case largely the programmers, create software they are defining who exactly the people that use their software are. In this case, they are even going so far as to influence their music preference.

IV. Research Question & Methods

This leads to the research question: How has the introduction of music recommendation software affected new artists' ability to gain popularity? To address this topic, I will analyze the differences in the rise of new artists before and after the introduction of music recommendation systems. This will be done with the use of technical papers on the past and present state of music recommendation systems.

To better represent how these systems interact with their users, research will be conducted to develop a deep understanding of what makes modern music recommendation systems work the way that they do. Specifically understanding why they might recommend music with lower monthly listeners. Spotify R&D will be used heavily as they have chosen to be very transparent about how their software works, such as publishing a paper on the effects of their software (Anderson et al., 2020).

I will use the various sources that I have mentioned to form my Actor-Network both before and after the introduction of music recommendation systems. For before their introduction I will use articles including “Innovation and Diversity in the Popular Music Industry” (Lopes, 1992) because it provides general information about how new artists became popular that will allow me to figure out some of the key actors and their relationships. To form the Actor-Network for after the introduction of music recommendation systems I will use sources including ” Recommendation Systems as Technologies of the Self: Algorithmic Control and the Formation of Music Taste” (Karakayali, Kostem, & Galip, 2018) as it goes over some of the most important actors from recent time. Once I have both of these I can dive deep into what specifically has changed between the two and more importantly what has stayed the same to isolate how this changing technology has included new artists.

Failing to understand the effects that music recommendation software has on new artists could lead to undetected ramifications in the music industry. These systems have been around for about two decades and we have seen tons of changes in who and what becomes popular for better or for worse. As stated in this paper from Bilkent University regarding current works, “These pioneering works provide invaluable insights about how recommender systems can exert control over users. But we still know little about how users utilize these systems in practice.” (Karakayali, Kostem, & Galip, 2018, p. 4). Although taste in music may not have a direct link to any life or death problems, understanding how people's preferences can be shaped by technology is very important to be aware of in today's world of digital influences.

V. Conclusion

The anticipated final project of the technical report behind this paper will be a drum system that can take an audio input from a phone or computer's audio port, use a microcontroller to process the signal and find the beat of the song, and finally play that beat in real time via a servo-controlled drumstick. Although this does not exist in exactly the same space as my STS topic, where they overlap is important. The developers of our technical project and the designers of music recommendation software are faced with the same challenge of limiting biases towards certain types of music. The research I will do in this area will help give developers a better understanding of what they can do differently to make this software equal for all types of musicians and their patrons. This builds off of how the STS research portion of this paper's goal is to break down the impacts of music recommendation software on new artists' ability to succeed in the music industry. The central idea of this is to use computer technology to analyze music and provide something of use to their users. Relative to other spaces, music is one of the most difficult forms of media for computers to understand, so the goal of the combination of these two projects is to bring attention to what is possible in this space (Purwins, et al., 2019, p. 207).

VI. References

Anderson, A., Maystre, L., Mehrotra, R., Anderson, I., & Lalmas, M. (2020, December 3). *Algorithmic Effects on the Diversity of Consumption on Spotify*. Retrieved from Spotify R&D | Research: https://research.atspotify.com/2020/12/algorithmic-effects-on-the-diversity-of-consumption-on-spotify/

Anselmo, L., McCullough, M., Lydon, D., Lilly, J., & Gomez Ibarra, U. (2022). *Whiplash [Unpublished engineering capstone project]*. University of Virginia.

Barnes, F. (2007). *The Drummer*. Retrieved from Stickboy - Robocross Machines: http://robocross-machines.com/robots/stickboy/

Billboard. (2021). *U.S. Music Year-End Report.* Luminate Data.

Cetinic, E., & She, J. (2022). Understanding and Creating Art with AI: Review and Outlook. *ACM Transactons on Multimedia Computing, Communications, and Applications*, 1-22.

Dowling, S. (2019, May 31). *Napster turns 20: How it changed the music industry*. Retrieved from BBC: https://www.bbc.com/culture/article/20190531-napster-turns-20-how-it-changed-the-music-industry

FreeCAD. (2002).

Hoare, M., Benford, S., Jones, R., & Milic-Frayling, N. (2014). Coming in from the Margins. *ACM Confrence on Human Factors in Computing Systems*, 1295-1304.

(n.d.).*HS-805BB Mega Giant Scale Servo.* HiTEC .

Karakayali, N., Kostem, B., & Galip, I. (2018). Recommendation Systems as Technologies of the Self: Algorithmic Contorl and the Formation of Music Taste. *Theory Culture & Society*, 2-24.

Latour, B. (1992). Where are the missing masses? The sociology of a few mundane artifacts. In J. Law, & W. E. Bijker, *Shaping Technology / Building Society* (pp. 225-258). Cambridge, MA: MIT Press.

Lopes, D. P. (1992). Innovation and Diversity in the Popular Music Industry, 1969 to 1990. *American Sociological Review*, 56-71.

MIDiA. (2021). *MIDiA Research Music Subscriber Market Share Model.*

Nikki, P., & Craig, M. G. (2020). Neural Network Music Genre Classification. *Canadian Journal of Electrical and Computer Engineering*, 170-173.

Polyend. (2015). Polyend Perc.

Purwins, H., Li, B., Virtanen, T., Schluter, J., Chang, S.-Y., & Sainath, T. (2019). Deep Learning for Audio Signal Processing. *IEEE Journal of Selected Topics in Signal Processing*, 206-219.

Recording Connection. (2022). *How Much Does a Studio Cost to Build*. Retrieved from Recording Connection: https://recordingconnection.com/reference-library/how-much-does-a-recording-studio-cost-to-build/#:~:text=The%20entry%2Dlevel%20home%20recording,much%20easier%20to%20record%20music.

Schedl, M. (2019). Deep Learing in Music Recommendation Systems. *frontiers in Applied Mathematic and Statistics*, 1-9.

Texas Instruments. (2018). SimpleLink Wi-Fi CC3220SF Wireless Microcontroller Launchpad ,Development Kit.

Woolgar, S. (1991). Configuring the user: the case of usability trials. (J. Law, Ed.) *A sociology of Monsters: Essays on Power, Technology and Domination*, 58-99. this was a collected volume.