Can Artificial Intelligence Play an Effective Role in the Production and Consumption of Music?

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

Michael Bliss
Spring, 2020

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

Signature  Date 5/7/2020
Michael Bliss

Approved _____ Date _____
Sharon Tsai-hsuan Ku, Department of Engineering

Fable of Tomorrow

There once existed a band unlike any other the world had ever seen. Four individuals, songwriters, and geniuses who forever changed the way music was produced and consumed. Their sweet melodies and soulful harmonies influenced an entire generation's youth and their success may never be rivaled. Every tune was produced with meticulous attention to detail, carefully thought-out lyrics, and creative thinking in a way that had never before been done. They shattered the barriers of what defined rock and roll, and their imagination knew no bounds. When they would arrive in cities on tour, crowds would swarm like bees. Thousands of people would travel miles and push and shove their way to the front of the crowd for the chance to catch a glimpse of one of their four idols. Their popularity was such that they would influence social movements, change the way people dressed, and inspire countless to pursue careers in music. Entire football stadiums would fill with fans eager to watch them perform. People would come from far and wide to see these deified celebrities perform their favorite songs.

To the dismay of millions, the Beatles broke up in 1970, after a mere six years of popularity. They had ceased their careers as touring artists four years before, in 1966. Although the four bandmembers, John Lennon, Paul McCartney, Ringo Starr, and George Harrison, each enjoyed some degree of success as individual musicians, none of their individual success came close to what they achieved as the Beatles. To make things worse, the Beatles' original founder and primary songwriter, John Lennon, was assassinated in 1970 (Unterberger, 2019, n.p.). Millions mourned in the wake of this tragedy, as candlelight vigils and memorials were held all over the world. The Beatles were a worldwide phenomenon and their influence far exceeded the confines of their home country, England. Following the disbanding of the Beatles, and following

the murder of Lennon, innumerable fans were left to experience a massive void. Given what they did in six years, one can only imagine what they could have accomplished in a longer lifetime.

But what if I told you that the Beatles had been reincarnated? What if I told you that there were finally additions being made to their decades-old list of songs and albums? There are millions of people who grew up in the 60s whose eyes would light up with excitement, but at the same time feel a glimmer of sadness and skepticism. Many would scoff and remark that John Lennon has been dead for years and that his lyrical and melodical genius will never again be matched. Contrary to what they might think, it has- by harnessing the power of artificial intelligence.

Introduction

Music is one of the most ubiquitous and popular forms of art which has existed for thousands of years. The first musical instruments known to archaeologists date back to 40,000 BC, and music of different forms has entertained populations to this day (Cassidy, 2019, n.p.). As seen in figure 1, the US music industry brought in nearly $20 billion in revenue in 2018 and this number has been on an upward trend since 2012 (Watson, 2019, n.p.).

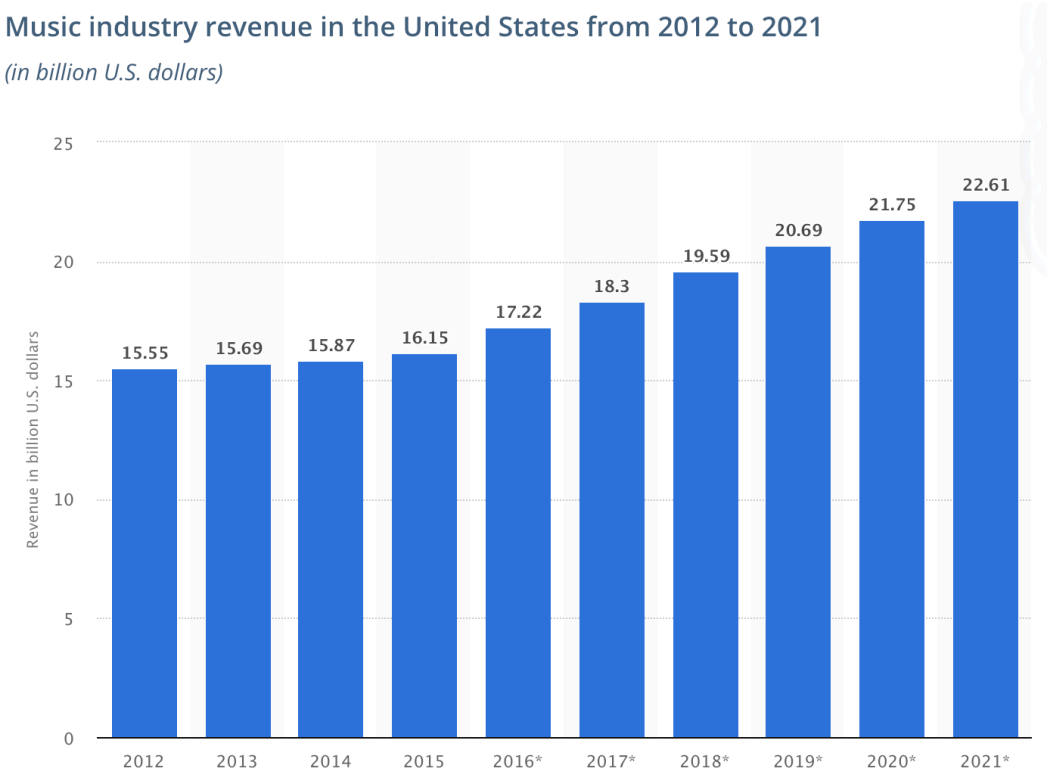

Figure 1: Music industry revenue in the United States. These figures include revenues for both live and recorded music, which are expected to exceed $20 billion in 2019. (Watson, 2009, n.p.)

A large portion of music consumption and revenue is driven by recorded music. However, in order to produce and distribute recorded music, one must go through the process of learning an instrument, procuring recording equipment, recording the song, mixing/mastering the recording, advertising the finished product, and distributing it to consumers. Because of the funding, expertise, and connections required by this process, record labels were formed and functioned to provide artists with these services (Bielas, 2013, p.3-23). Since their formation, these record labels have often been able to use their control over the industry to force artists into contracts which are non-negotiable, unfair, and indeterminate. (Eiriz & Leite, 2017, p.875-885). This has been effective in deterring unaccomplished artists from trying to get into the recorded music industry and monetize their works (Guesman, 2018, n.p.). However, within the past two decades, technological advancements have made it possible for artists to bypass these record labels and

learn, produce, distribute, and promote recorded music on their own. Technologies like YouTube are useful for learning instruments, while digital-audio workstations (DAWs) such as GarageBand are useful for mixing/mastering recorded works. Technologies such as Spotify or SoundCloud are useful for distributing music, while social media applications like Instagram or Facebook are helpful with promotion.

However, with what we now know, it seems as though there is a tradeoff between quality and accessibility. Many artists believe that the digitization of processes relating to the production and consumption of music have resulted in a depletion of artist creativity. For example, many DAWs such as GarageBand come with a preloaded set of beats and rhythms that can be reused and looped in the background of songs. By providing these beats along with the rest of the platform, software developers essentially replace the role of the music teacher in dictating how songs are created. A similar issue has surfaced in recent years with the introduction of music-related artificial intelligences. There now exist artificial intelligences that create, enhance, and teach music, and these technologies have been met with similar backlash to previous disruptive technologies in the music industry. Although people are often comfortable allowing technology to assist and facilitate other aspects of their lives, each disruptive technology in the music industry has been met with some degree of skepticism. This is likely because music encapsulates human creativity and emotion, two qualities that many people believe a machine will limit or suppress. In this paper, we will investigate the potential benefits or shortcomings of these technologies in achieving their desired goals. Additionally, we will consider external factors in the network to further discuss the likelihood of these technologies gaining acceptance and making an impact in the music industry.

STS Framework

Actor-Network Theory (ANT) is general social theory aiming to understand technoscience and will be used to frame the analysis of the potential benefits and shortcomings of music-related artificial intelligence. ANT represents technoscience as the formation of larger networks consisting of heterogeneous actors which are both human and nonhuman. These actors are treated as equals through an ANT lense and their interactions with other actors form associations and broaden/change their networks (Sismondo, 2010, p.81-92). It will be useful to determine the pros and cons of these technologies by framing them as newly introduced nonhuman actants in different networks of music production and consumption. From there, the effectiveness of these technologies in achieving their goals can be determined by examining how the network would be affected by the replacement of the human actant with the nonhuman actant. For example, there exists an AI-assisted music teaching platform called SmartMusic (Apremaya, 2018, p.10-14). In order to determine its effectiveness as a music teacher, I plan to investigate the impacts on music-learning networks and its actants if SmartMusic were to completely replace a traditional music teacher. For example, it might have added benefit of being able to provide more detailed feedback related to the accuracy of notes played and tempo, but it might lack the personability and encouragement that a traditional piano teacher provides. This tradeoff will be discussed further in subsequent sections of this paper.

Research Method

In order to address the topic of artificial intelligence in the music industry, it is necessary to research the current technologies that exist, and the perceptions of these technologies among various stakeholders. I will analyze three digital music platforms—Sony's Flow Machines,

Google's Magenta and SmartMusic—to address the AI-facilitated music production and music education, along with their shortcomings, and formulate opinions regarding their effectiveness in achieving their goals.

1. Document analysis will be used to understand different actors' opinions, ranging from musicians, producers and music teachers who express their feelings of confidence or skepticism toward the existing technologies. In addition, I will discuss the design of the technologies to understand how they redefine future music industry in terms of content and authorship. There are currently numerous proposed technologies that have not been implemented; however, many of these technologies' potential benefits or shortcomings become visible through critical examination of the authors' plans and their optimisms.
2. I will interview Valerie Snow, who is the current Music Director at Dedham Country Day School and who has taught piano to students of all levels and abilities for the last 30 years. Because of her background in teaching, her opinions related to the existing and proposed music teaching artificial intelligences will be useful. Her career and experience in the music will offer insights specific to the industry. She also has a very strong understanding of music theory and composition, which will provide a greater understanding and analysis regarding potential music-composing artificial intelligences.
3. Secondary data drawing from news articles and existing literature will be collected to understand how users, consumers and general people perceive AI application in music production.

Data Analysis

AI Music Production

As aforementioned, there exist numerous AI technologies that aim to assist with music composition, teaching, mastering, and discovery. In order to determine the effectiveness of these existing technologies, I researched multiple platforms with goals to compose or teach music. The two music composition platforms that I analyzed are Sony's Flow Machines and Google's Magenta project. These platforms were selected because they are two of the most developed and documented AI-driven music composition platforms that exist, and because they were developed by large, reputable tech companies. Each of the two platforms utilize a machine learning algorithm called reinforcement learning to achieve their goals of producing music (Marr, 2019, n.p.). Reinforcement learning is a process through which a computer generates a model for predicting the solution to a given problem. It analyzes numerous aspects of a given input dataset, and using the model it has generated, it makes a prediction for data that might succeed its input. The quality of this prediction is measured using some reward function, and the value of the reward function is used to refine the computer's model with the goal of maximizing future predictions' rewards. This enables the computer to make better future predictions. The theory of reinforcement learning is that, when studying a large enough training dataset for a large amount of iterations, the computer will deduce a model that can accurately make predictions on its own (Budek & Osiński, 2018, n.p.). For example, if the goal is to predict text, the reinforcement learning algorithm might study some large body of text for a large amount of iterations. The algorithm would note patterns such as word frequency, word pairings, sentence structure, and grammar. If a good model has been created, then the user should be able to supply the AI with a body of input text and observe as it outputs a coherent, sensible body of possible subsequent text.

Google's Magenta project is an open source music composition AI that utilizes TensorFlow's deep learning library. The platform consists of four different applications, each with a different purpose. These different applications allow the user to generate beats, interpolate beats between two provided bars, continue a beat based off of an input, or groove with an input. The AI was trained using the Yamaha e-Piano Composition dataset, which consists of roughly 1400 MIDI recordings that capture the performances of skilled musicians. One very important feature of Google's Magenta project is the AI's ability to learn expressive timing and dynamics. Expressive timing and dynamics describe the length for which notes are held and the volume at which they are played, and these qualities are essential for conveying emotion in music. In order to capture these features, the AI observes four qualities for each note played in the input dataset. These qualities are pitch, the time when the note is played, the time when the note stops playing, and the velocity at which the note was struck. Since the AI studies actual human performances as opposed to score sheets, these qualities are easily observable and allow the AI to generate beats where the expressive timing and dynamics are easily observable. However, despite these positive qualities of the output, Google's Magenta project only creates small snippets of a beat and these beats do not have much long-term structure. The platform also does not allow for much customizability of an output. It is a black-box platform and the only quality that can be changed by the user (besides the input beat) is the "temperature" of the output. The temperature is a value between 0 and 2 and it dictates the randomness of the model used for prediction (Oore & Simon, 2017, n.p.). Therefore, the user has limited ability to control important aspects of a beat such as emotion, tempo, and style.

Sony's Flow Machines is another music composition AI with a slightly different goal. Instead of composing a few bars of a beat/melody for one instrument, it has the goal of creating

complete melodies and harmonies either autonomously or collaboratively with humans. Another goal of Flow Machines is to model an artist's style as a computational object, allowing the user to customize the style of the output accordingly (MIDI Association, 2020, n.p.). For example, Flow Machines released its first song, "Daddy's Car," in the style of the Beatles. The melody for "Daddy's Car" was composed completely by Flow Machines, but the lyrics were written by Flow Machines composer Benoît Carré. The machine's model is derived from studying a 13,000-lead sheet dataset consisting of pop, jazz, and Brazilian music. According to the project lead François Pachet, these genres were chosen as they focus heavily on melody and harmony, whereas genres such as heavy metal or rap commonly do not. In order to create its model, the machine and its creators dissect each song into four stages. The first of which is the lead sheet itself, which just contains symbolic information about which notes should be played and when and for how long. The second stage is the orchestration of the arrangement, meaning which instruments will be used to produce the desired sound. The third stage is the production of the piece and the special effects with which each note is played. Lastly, the fourth stage is the complete performance of the piece. Through this analysis, the machine is able to observe and understand complex patterns that are essential to the production of a good song, such as emotion. While the AI is able to operate completely autonomously, Pachet notes that, "A good song is rare […] A song tells a story: a beginning, a middle and an end. Today AI techniques are not very good to produce structure." (Ferreira, 2016, n.p.). Because of this, the Flow Machines' AI allows the user to customize its output. Once the AI has composed a work, the user can tell the machine what parts it likes, which it doesn't, and the machine will learn from this feedback and recompose.

AI-Facilitated Music Education

SmartMusic is a platform that provides a suite of web-based music education tools that assist music teachers with their lessons. Each student has their own account where they can practice their respective instruments and receive real-time feedback on the correctness of their performance, based on pitch and rhythm. The platform is able to identify areas of weakness in students' performances and it possesses a tool that loops these tricky sections for the student to gain extra practice. SmartMusic saves audio recordings of these performances and teachers are able to view the correctness analysis, hear these recordings, and provide feedback. Teachers are able to post assignments tailored to each student, and the platform contains a tool that can generate sight reading exercises based on a specified range, rhythm, time and key (SmartMusic, 2019, n.p.). Cornell University Student Lavana Apremaya has proposed an AI-driven music teacher based off of SmartMusic's model called the Autonomous Intelligent Music Teacher (AIMT). AIMT is an attempt to enhance the platform provided by SmartMusic and minimize the involvement of the human teacher. Apremaya proposes AIMT as a four-stage platform. The first stage is music generation. Using composition algorithms similar to those previously discussed, the AIMT would compose scores for the students to practice based on specifications such as key and form. The student would then perform according to this composed lead sheet and AIMT would evaluate aspects of the student's performance, similar to the way SmartMusic does. However, the AI will note more complex patterns in the student's performance and to factor in to a more detailed feedback grade. AIMT aims to be able to evaluate emotion in a student's performance by harnessing the power of artificial intelligence. For example, if a piece indicates that it should be played with the style of agitado, the AIMT would evaluate whether or not the performance is played in an agitated tone, with sharp accents to deliver the intended effect. The

final stage is the adaptation to student performance. Based on student strengths and weaknesses, the AIMT will identify areas that need improvement and compose new music for the student to practice accordingly. Apremaya encapsulates this four-stage platform in a flowchart, seen in figure 2 (Apremaya, 2018, p.10-14).

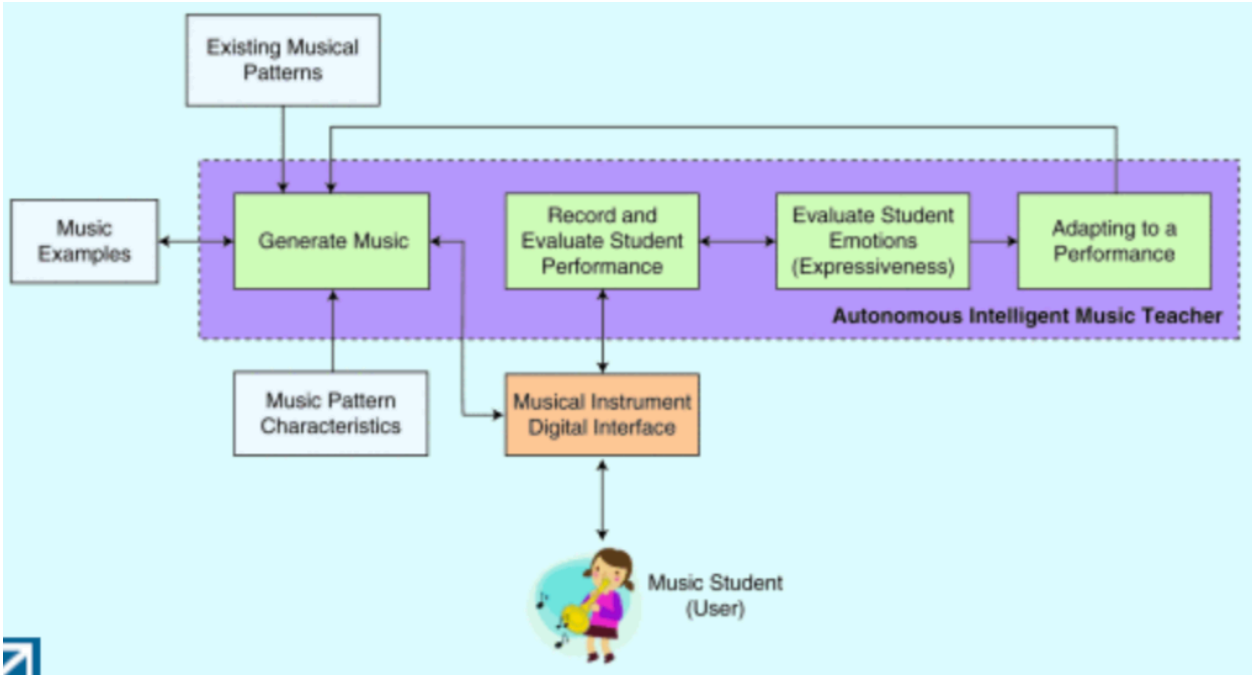

Figure 2: Flowchart of proposed Autonomous Intelligent Music Teacher, as Described by Lavanya Apremaya (Apremaya, 2018, p.10-14).

Despite of the potential benefit AIMT could bring to music education, questions about the essence of music and music education were raised by the Director of Music at Dedham Country Day School (in Dedham, MA), Valerie Snow. Snow has been teaching piano to students of all ages and abilities for the last 30 years, many of whom are neuro-atypical. She pointed out the biggest challenges that she faces in her regular piano class is students' frustration due to a slow learning process. "Note reading is hard, and rhythm is hard for students to understand. Often students want to learn songs that they hear on the radio or on YouTube and they are simply not ready" (Snow, 2020). As a traditional music teacher who believes in face-to-face learning, Snow had not used many technologies and was reluctant to use AI in teaching

music. The technologies that she had used, however, were more for practicing already-learned skills than for developing new skills. She said the platforms would evaluate student performance and provide feedback, but they would not actually provide any instruction on how to actually play the piano. Her attitude toward Autonomous Intelligent Music Teacher was skeptical. She noted that she could see the benefit for somebody who wanted to quickly learn an instrument; however, many important qualities of a human teacher simply cannot be provided by the machine:

> Students often require attention and approval for inspiration. Many of my students will constantly look at me for approval while playing their pieces. Also, there is a lot of psychology that goes into understanding students. Students learn at different paces and handle failures differently. I don't think my students would enjoy practicing piano if I were to only pick pieces that highlight their weaknesses (Snow, 2020).

As a classically trained teacher/pianist, Snow pointed out the importance of developing a strong understanding of music theory for being an aspiring musician, an aspect that AI might be able to help:

> Unless you have perfect pitch (which is very rare), it is essential to understand music theory if you wish to compose anything of quality. There are many patterns explained by music theory that are important to achieve a good sound… it is almost all based on mathematics, so an AI would be able to deduce theory by studying sufficient compositions (Snow, 2020).

However, regarding the capacity of AI's potential involvement in music composition, she responded with similar skepticism to many others, saying that "It makes me cringe to think that someone would think that something artificial could compose or teach music. At the same time, I don't know enough about it to be convinced otherwise" (Snow, 2020).

Discussion

Each of the technologies described previously present numerous advantages and disadvantages compared to their human counterparts in the process of achieving the goals that they aim to achieve. However, in order to determine the potential successes of these platforms, it must also be considered whether or not they would be accepted into their respective networks. As discussed in the introduction, each disruptive technology in the music industry has been met with some degree of skepticism and most have introduced some level of tradeoff between accessibility and ease of use vs. quality and complexity. However, each of these previous technologies still necessitated some form of human operator for their use. A digital-audio workstation does not produce music without a musician operating the computer and YouTube cannot provide music lessons if no teacher records and uploads themselves giving lessons. There have never existed platforms that so explicitly make the decisions that humans traditionally make such as these AI platforms, so they are sure to be met with a great deal of doubt. Based on the previous descriptions of Sony's Flow Machines, Google's Magenta project, SmartMusic, and the Autonomous Intelligent Music Teacher, it is fairly obvious that none of them currently possess

the capability to replace their human counterparts in their respective networks. This raises the question: could they be an effective addition to their networks, and would they be accepted?

In order to determine whether or not these technologies will be accepted, one must consider the musicians who would use them and why they would be skeptical. Although people are generally comfortable allowing machines to facilitate more trivial aspects of their lives, many are uncomfortable with the idea of a machine performing tasks that mimic human creativity. This raises the question: what is creativity? Is Sony's Flow Machine creative for composing "Daddy's Car" in the style of the Beatles? Or does it lack creativity because it achieved such by studying the score sheets of every Beatles song that exists? One could argue that it is creative and that every musician has received some level of influence from hearing and studying other musicians' performances, similar to how the AI is heavily influenced by its training dataset. One could also argue that it is not creative and maybe even an infringement of copyright laws because the AI's model is explicitly and directly trained by other artists' works. But what if an artist trains an AI to mimic their own style to assist with composition? Would that be considered creative? It is important to consider the artists' and their listeners' perspectives to determine if these technologies will be successful.

Leon Smart is an electronic music producer from the UK. During an interview, when asked about "Daddy's Car," Smart remarked, "That's cool, but I'd really want AI to further what we do, not just… do what we do" (Cills, 2017, n.p.). Both Sony's Flow Machines and Google's Magenta project have the capabilities to do such. Both are able to produce beats or melodies based on a provided input. However, the level of human creativity required to go from the machine's output to a finished product varies between the platforms. The output that Flow Machines provides is much closer to a finished product than what is produced by Magenta. Flow

Machines' training dataset is much larger and its algorithms are more thorough and computationally-intensive. At the same time, the dataset contains copyrighted material and its output can often mimic such copyrighted material. Although Sony states the goal of Flow Machines as to not replace the human artist but to augment their creativity, the platform comes much closer to replacement of the human artist than Google's platform does. Conversely, while Google's Magenta project produces much shorter and less complex beats and melodies, the argument that it can assist human creativity is stronger. It takes much more human contribution to go from Magenta's output to a finished song than from Flow Machines's output. It takes much more musical talent to create a complete song based off of four bars of drum beats than from a complete, multi-instrument melody.

In determining whether or not the music-teaching artificial intelligences would be accepted, the creative aspect is less of a factor. The most important factor of a good teaching technology is its ability to effectively communicate with students and to effectively teach good practices. Both SmartMusic and the Autonomous Intelligent Music Teacher can be useful in evaluating student performance for correctness. However, both of these platforms assume that the student already possesses some level of ability to play their instrument. Both platforms are able to determine if the user has played a C note, but neither instructs the user *how* to play a C note on their instrument. They also both lack human-like feedback and encouragement, a quality that is necessary to many developing learners. When asked whether she thinks AIMT could be effective for teaching her students, Snow said, "I think that it could be a useful practice tool, but I think that many of my students would grow frustrated in time. I think that many would become frustrated if they were consistently tasked with playing pieces that highlight their weaknesses and their only feedback was a number" (Snow, 2020).

Conclusion

In conclusion, none of the existing AI technologies in question are ready to fully replace their human counterparts in their respective networks. Even if they were, it is doubtful that they would be accepted by the musicians, the listeners, the students, or the learning institutions. However, in a way similar to other disruptive music-related technologies of the past, I think that some of these AI-driven technologies have the opportunity of providing useful additions to their respective networks. It is clear that neither Sony's Flow Machines nor Google's Magenta Project are ready or able to create complete, high-quality songs; but either could provide great utility in driving artist inspiration. Artists can use pieces of the machine output as building blocks for more complex melodies. One could argue against this technique on the basis of creativity or lack thereof. However, the technique of using computer-generated bars as inspiration for a larger musical work is not something new. Artists have been using this technique in the form of digital-audio workstations for decades. Does Rihanna lack creativity for using a GarageBand loop in her hit song *Umbrella*? Or is she creative for starting with a short beat and building it up to ultimately create one of the most popular songs of the time? I would argue that she is creative in this regard and I think that the real argument regarding creativity will arise when these music-composition AIs are ready to fully replace their human counterparts in their respective networks. However, when these technologies do exist, the ethical and legal issues that will arise will be so large in quantity that I'm not sure they will be accepted.

As the music composition AIs are not ready to fully replace the artist, music teaching AIs are not ready to fully replace the teacher. A platform like SmartMusic or AIMT could be a useful addition to the classroom for teachers, but for most learners it will not be sufficient. Learning music can be frustrating and if one doesn't possess the proper resources to learn effectively, it is

likely that they will grow discouraged or develop bad habits. The AI-driven platforms that exist/have been proposed do not actually teach the student how to play an instrument or how to read music, they just evaluate the student's ability to do so. They also do not provide any of the important feedback or encouragement that a human teacher is able to provide. I think that for a music-education AI to encapsulate more of the teaching process in hopes of eventually replacing the human music teacher, some type of psychology-related aspect would need to be included. Music students are often children who thrive off of the encouragement from their teachers. Also, every student is different in terms of ability and in terms of personality. Adapting to the learning styles of different students is an essential capability of the human teacher and it is a capability that the current proposed technologies do not possess.

Although these technologies are not ready to replace their human counterparts yet, they will likely continue to develop and perhaps one day they will possess the technological ability to replace their human counterparts in their networks. This raises the question of whether or not these technologies will be accepted into their intended networks. I think that the music education AIs have a much better chance of being accepted than their compositional counterparts. People tend to value the end result of learning more than they value the method through which they were taught. As discussed previously by Snow, music teachers and scholars will have plenty of issues with an AI music teacher. But for your typical learner, if the AI teacher is able to provide everything that a human teacher can at a fraction of the cost and on demand, why wouldn't you save the money and experience the added leisure? There would still exist many more ethical debates regarding if the AI can actually effectively perform all the roles that a human teacher could, but it is more likely to be accepted than a music-composition AI.

Regarding the music-composition AI, I do not think that enough of the involved parties would be in favor of the machines for them to gain much popularity. First of all, the artists would take issue with their replacement. Secondly, I think many people look for meaning/emotion in the music they listen to and that they'd be hesitant or uncomfortable listening to computer generated and composed music. Thirdly, there are so many legal issues that would arise surrounding copyright laws that it might make the whole phenomenon of AI-composed music problematic. How does the copyright law work when a computer has generated its compositional model from thousands of other artists' work? Do Google and Sony own the copyright for creating the software? Or do the artists whose works formed the model deserve ownership? One can argue that the human artist learns from countless other artists in a similar way to the AI, but with the AI it is explicit and is documented in the AI's code. For these reasons, I think we are a long way away from having music fully composed by AI. Even if the technology is to exist, the societal implications are such that it may not be accepted for a long time.

References

Apremaya, L. (2018). An autonomous intelligent music teacher. *IEEE Potentials, 37*(1), 10-14. Retrieved from https://ieeexplore-ieee-org.proxy01.its.virginia.edu/document/8253748

Bielas, I. (2013). The rise and fall of record labels. (Undergraduate senior thesis). Retrieved from https://scholarship.claremont.edu/cmc_theses/703/.

Budek, K., & Osiński, B. (2018). What is reinforcement learning? The complete guide. Retrieved from https://deepsense.ai/what-is-reinforcement-learning-the-complete-guide/

Cassidy, S. (2019). The complex genetics of musical talent. [Blog post]. Retrieved from https://www.jax.org/news-and-insights/jax-blog/2018/may/complex-genetics-of-musical-talent.

Cills, H. (2017). Can AI make musicians more creative? Retrieved from http://www.mtv.com/news/2983208/ai-artificial-intelligence-music/

Eiriz, V. & Leite, F. P. (2017). The digital distribution of music and its impact on the business models of independent musicians. *The Service Industries Journal, 37*(13-14), 875-895. https://doi.org/10.1080/02642069.2017.1361935.

Ferreira, E. (2016). How Sony Flow Machines' AI takes itself for the Beatles. Retrieved from http://www.makery.info/en/2016/10/14/comment-lia-de-sony-flow-machines-se-prend-pour-les-beatles/

Guesman, B. (2018, February 25). Rehashing the Recording Process: A look at the democratization of audio production. Retrieved September 30, 2019, from https://medium.com/steam-stories/rehashing-the-recording-process-a-look-at-the-democratization-of-audio-production-ec00f4afd9db.

Marr, B. (2019). The amazing ways artificial intelligence is transforming the music industry. Retrieved from https://www.forbes.com/sites/bernardmarr/2019/07/05/the-amazing-ways-artificial-intelligence-is-transforming-the-music-industry/#42f4e3885072

MIDI Association. (2020). MIDI and music AI. Retrieved from https://www.midi.org/articles-old/midi-and-music-ai

Oore, S., & Simon, I. (2017). Performance RNN: generating music with expressive timing and dynamics. Retrieved from https://magenta.tensorflow.org/performance-rnn

Sismondo, S. (2010). An introduction to science and technology studies. Chichester, West Sussex, U.K.: Wiley-Blackwell.

SmartMusic. (2019). What is SmartMusic? Retrieved from https://www.youtube.com/watch?v=H_lRZBSo_zU

Snow, V. (2020, Feb 18). Personal Interview.

Unterberger, R. (2019) The Beatles. Retrieved from https://www.allmusic.com/artist/the-beatles-mn0000754032/biography

Watson, A. (2019, August 27). U.S. music - statistics and facts. Retrieved from https://www.statista.com/topics/1639/music/.