AI and Art

Humanity for millennia has associated being human with acts of creativity. We revel in the creations of our forebears who drew animals and put handprints on the walls of caves. These simple starting points in some meaningful and profound way differentiated the evolving homo sapien species from all other creatures on earth. Art and by extension the creativity involved in its creation signaled a cognitive difference that sought to establish something beyond the present, a permanence or remembrance. Art in all its forms has been long considered the most human of endeavors. The privileging of the human in the creation of art is thought to signal something in our fundamental nature that was irreplicable by computers and algorithms. Yet, as the third decade of the twenty-first century dawns technology is positioned at an inflection point. It is becoming increasingly clear that much of the art we know and love is capable of not only being reproduced by machines, but also that creativity itself might be something for which the computers of the mind and those of silicon share aptitudes.

The human mind, a biological computer capable of creation and destruction is being challenged as the sole progenitor of art. While the computational power and robustness of algorithms has improved substantially over the last 20 years, there has been concern and interest in machines that were smarter and more creative than their human counterparts dating back to 1770 when civil servant Wolfgang von Kempelen's invention of the Chess Playing automaton began to impress Empress Maria Theresa of Austria-Hungary.[1] The device known as the Turk (The eventual origin for the name of today's Mechanical Turk owned by Amazon), played chess, toured the world, and caused a stir everywhere it went. Von Kempelen even purportedly showed off the gears on the inside an effort to build the illusion. By all accounts the Turk was a wild success. The device even played chess matches against the likes of Charles Babbage (Inventor of the first mechanical computer) and Napoleon Bonaparte. While eventually the Turk was discovered to be hoax controlled by a human operator the questions it posed to the centrality of the human mind in creativity were and remain very real. Questioning the need for humans in acts of creativity and cognition have resulted in the creation of a number of tests to differentiate human from machine. Among the first of these was the Lovelace test (developed by Ada Lovelace 1843 – widely considered one of the first computer programmers) and the Turing test (developed by Alan Turing – inventor of early computers 1950) to determine when a machine is truly intelligent. Yet the groundbreaking advances of Artificial Intelligence (AI) arose in the playing of games. Most notably these advances drove public interest in AI intelligence starting with chess and the creation of IBM's deep blue which beat World Chess Champion Gary Kasparov, and subsequently through to Google's AlphaGo victory over champion Go player Lee Sedol.

---

[1] Tom Standage, "Monster in a Box," *Wired*, March 1, 2002, https://www.wired.com/2002/03/turk/.

Although games like chess and Go might seem unrelated to art, players describe the games they play as a form of creativity that leverages many of the same skills attributed to musicians and painters. There are numerous examples of computer scientists trying their hand at developing AIs that can learn from and subsequently generate art independent of the human hand. IBM and others have tried to have AIs learn from great masters of painting and subsequently generate their own art. Some AIs so closely mimic the artists on which they learn they can fool art critics, others like art.io can transform photographs into nearly any style of art imaginable. AI generated paintings have even sold at auction, taking in $425,000 at Christies in 2018.

Music, in particular is often thought of as highly mathematical and therefore susceptible AI-human labor offsets. Notes are commonly structured in measures and bars and music theorists study a variety of attributes of music that closely correlate to mathematics. Similarly, mathematicians are often fascinated with the mathematics of music and the relationship of notes in harmony and disharmony. Great musicians such as Johan Sebastian Bach were known for using music in ways that were so explicitly mathematical that by following a coded schema other composers (or mathematicians) could derive new pieces based on the structure and templates of previous ones.[2]

The streaming music start-up company AudioIn began its service small in the early 2000s. The company focused initially on independent artists and other artists with limited legal or label protections. AudioIn paid their artists reasonably for each play of a song on their service and charged users a nominal fee to listen the entire catalogue of music. Through the early 2000s the firm grew both in its music offerings and its user base and by 2010 it was one of the top five streaming services online. AudioIn began to commission artists and create their own music catalog which they marketed as the largest catalog of independent sounds on the Internet. They were able to diversify their music inventory by 2015 to include every genre of music.

In 2010, AudioIn sought to better tailor the playlists of music it offered to its users. It did this to increase user engagement and interest in the service. To do this it began to heavily recruit AI and Machine Learning (ML) programmers to find ways to provide accurate and engaging offerings to its users. The programmers followed common practices being developed elsewhere among social media companies and instituted a "like button", analytics on individual user preferences and began correlating preferences with other similar users. By 2018 their efforts were proving fruitful and AudioIn had become one of the top three streaming music providers. Their users were happy with the suggested playlists and found them extremely useful.

---

[2] Marcus Du Sautoy, *Creativity Code: Art and Innovation in the Age of AI* (Cambridge, MA: Harvard University Press, 2019).

AudioIn had increased their userbase substantially over the previous eight years and had seen their profits rise and their shareholders, employees, and customers (users) were all extremely happy. However, by 2019 their userbase growth had begun to stagnate and profits had begun to flatten. While the company still earned healthy profits, its growth potential was increasingly constrained. To address flattening growth in users and profits the company initially increased the fees it charged its users, but only minimally to prevent declines in its userbase. Although there were some grumbles at the increased fees, the quality of tailored playlists kept most users on the platform from leaving.

By 2020, AudioIn had by all accounts become one of the market leaders in streaming music, it earned strong profits, but its prospects for growth were still stagnating. AudioIn was not unique in its market challenges, similar market stagnation was prevalent across all major competitors as the market became fully saturated. AudioIn Corporate leadership was at a crossroads and was seeking solutions to cut costs and increase its profit margins. It couldn't lower its payments to artists for fear they would withdraw their catalogs from its platform, and it couldn't charge its users more money for fear they would defect to other platforms. So again, in early 2020 AudioIn began hiring AI professionals.

The first task of these professionals was to identify the costliest subsectors of the AudioIn catalog, i.e., which artists and genres were paying out the most money. Some of these subsections were comprised of non-vocalist artists who used instruments to create soothing sounds for workouts, night-time sleep routines, yoga studios, mediations, elevator or corporate clients, classical works, jazz artists and more. At this point the AI team proposed an idea to AudioIn's corporate officers. What if it was possible to have artificial intelligences create new content of equivalent quality that users would like, but that would dilute the catalog and reduce aggregate payouts to human artists within the catalog. They took as their inspiration the work of David Cope, EMI (Experiments in Musical Intelligence) and others who pioneered AI generated music. They noted that Cope and other AI experts had been able to leverage computers to write and perform music that even many music critics could not distinguish from great composers. Moreover, they noted that AudioIn had access to a vast trove of human generated music that could begin the process of training AIs to develop music independent of human collaboration.

AudioIn had a number of advantages working in its favor. First, it had an enormous catalogue of music on which to train its AIs. Second, it had nearly 20 years user preference data and behavioral data to use to hone its efforts to those portions of its catalogue where it could reap the most financial reward. The question then became, could AudioIn generate enough high-quality content that its users would like to lower its payouts to human artists? Would its users know or care that they were listening to AI generated music rather than human created music? Should AudioIn let their users know they were listening to music generated by a machine rather than a human.

AudioIn slowly rolled out its efforts to dilute its catalogue with AI generated tunes. An initial decision was made to roll out new songs under pseudonyms given to each AI and respective category. As new music was generated, it was given similar treatment given conventional human released music. A song name was assigned, and each song was assigned to a fictional musician or group. It was believed by both the AI team and the marketing teams that by giving the songs names and assigning them to fake artists and albums they could mitigate human biases against machine generated music. AudioIn even hired an artistic design team to develop album art that would display when the songs played in any given user's playlist.

The team started in the classical genre, where computer science experts had previously had a great deal of success in approximating the styles of classical composers. Starting with piano concertos and other single instrument pieces they gradually began to increase to quartets and eventually to entire symphonies. The AI team and corporate teams were all extremely excited by the quality and progress of the AIs and gave approval for the content teams to begin the process of including music into the main music catalogues on a trial basis.

The initial forays into diluting the music catalogue began by feeding in one AI generated song for every nine human made songs. This testing process extended for several months. Rather than user consternation, AI generated songs began seeing direct requests under their constructed names and artists. Eventually within the classical section of the catalogue the AI team was able to shift more than 50% of all songs in its automated playlists to AI generated music. AudioIn corporate leadership was extremely impressed. The dilution of the playlists with AI generated music was resulting in real cost savings within that section of its catalogue and they authorized the slow and deliberate expansion into other genres. By mid 2022 thousands of songs had been introduced into multiple sections of the non-vocalist portions of its catalogue and the savings were starting to accumulate and the profit margin for AudioIn was starting to grow again. It seemed everyone was happy. Customers were happy and unknowingly requesting the AI generated music. AudioIn's shareholders were happy with the improved profit margins and Arlington, the city where AudioIn was based was happy with the influx of new high-paying jobs.

Despite all the excitement over the improved profits, happy customers and shareholders, another group began asking questions, the musicians. Since late 2020 non-vocal musicians began witnessing a steady decline in the number of times their songs were played on the AudioIn platform and with the drop in the number of plays came a precipitous decline in artist revenues. Social media platforms were soon abuzz with rumors as to why artists were losing revenue on the AudioIn platform. Some speculated that AudioIn's analytics were faulty, yet soon it became apparent that when using automated playlists their songs simply did not come into the mix as frequently and new "unknown" artists were used on a more regular basis. When the human artists tried to find these new artists in the real world, they could not find them. Strangely most of the newer artists being cycled into playlists had no social media presence, no websites, and perhaps most strangely of all did not perform publicly.

All of the facts pointed to these artists not being "real." This led to a great deal of confusion among AudioIn users as the fervor of musicians intensified. AudioIn users similarly could not find out anything about the artists they were increasingly listening to. Quickly attention turned to an analysis of AudioIn's hiring practices over the last several years and both users and artists noticed the substantial increase in the number of new AI and ML specialists hired. As the pieces of the puzzle started fitting together AudioIn began to face a number of tough questions and was eventually forced to admit that it had been using AIs to develop music to dilute portions of its catalogue.

Many of AudioIn's users and human artists felt betrayed and began to depart the platform. Users felt as though they had been manipulated into listening to what they were increasingly referring to as "fake art". The artists felt manipulated and betrayed as they lost their revenue streams and their music had been used to "train" AIs to create new music inspired by their own hard work. By early 2023 two class action lawsuits were filed against AudioIn. The first lawsuit was comprised of users and the second of artists.

Discussion Question #1
It is hard to define art. Its definition seems to change over time and by culture, yet what is rarely debated is the creator of art. Although in the last decade there have been challenges to the definition of art creator from within the animal world, a more profound definitional change is arising out of computer science, in particular the subfield of AI. Is art fundamentally a product of the human mind derived from imagination, creativity, emotion, beauty, or the ideas of the artist transferred through technical proficiency and skill to a canvas, stone, sheet of music or other medium? What makes art – art? Is it a human endeavor or can computers also be creative?

Discussion Question 2: *AudioIn sought to maximize profits and minimize costs through the use of AI in the creation of non-copyright protected music which saved the company money. In the process they satisfied their users and shareholders by withholding information as to the origin of large portions of their evolving catalogue. What responsibilities, if any, did AudioIn have to be transparent about the origins of the new content generated for their platform? Was the creation of fake artist names, groups, and album art unethical? Did their attempt to hide the providence of the music being generated undermine their legitimacy and credibility? Should it matter whether the artist was human or machine?*

Discussion Question 3: *AudioIn maximized in part by training content generating AIs on the creative works of human artists. Without these prior works AudioIn would have been unable to generate the content it did. Did AudioIn engage in unethical and immoral practices? They did not explicitly violate any copyright laws and all works were new, but they leveraged human creative*

labor to train machines to offset or replace human artists, what, if any, responsibilities did AudioIn have to tell the artists about what it was doing?

Discussion Question 4: *AudioIn was able to maximize profits and tailor its AI music generation to meet user preferences by utilizing data collected on users. AudioIn requested consent for the sharing of this data, however users thought it would be used to tailor music content from human artist. Is the fact that Audio in used it for non-human generated music fundamentally different? The data provided detailed insights into what types of music AudioIn users liked, absent this data, AudioIn would not have been able to create AI generated music to lower costs. What, if any, responsibilities did AudioIn have to its users to tell them that their data was being used not only to create tailored playlists of human artists, but also tailored playlists of AI generated music?*

Discussion Question 5: *AudioIn was successful in creating AI generated art largely because of its access to a vast trove of music on which to train its AIs. Is training an AI different than training an art student by taking them to a museum and showing them different styles of art? Is it different than giving a student books on music theory or having them practice piano or violin pieces by famous composers? What makes a computer learning on thousands of pieces of music different from a human learning on the same music? Is there a difference? Is the difference one of biological vs non-biological entities? Is the difference one of scale (a human learns far more slowly)?*

Reflecting on AI and Art

The story of AudioIn is not hypothetical and is modeled on news stories plaguing streaming giant Spotify in 2017.[3] Scholars working on advances in artificial intelligence are increasingly leveraging algorithms in ways once thought impossible and art is at the frontier of AI learning. There are attempts underway at universities and corporations around the world to use AI algorithms to generate different forms of art. Substantial efforts have made headway in creating art in the style of multiple classical and modern painters using Oil, Acrylics, Watercolor, and other mediums. AI Art has both mimicked prior artists and adapted on those artists by ingesting large volumes of data and generating new alternative outcomes. Some AIs focus on portraiture, while others emphasize landscapes or surrealist art. Some AIs have learned solely by viewing prior works of art, others by watching artists paint, and yet still others by interpreting the responses of critics to art.

---

[3] Andrew Flanagan, "Spotify Is Accused Of Creating Fake Artists — But What Is A Fake Artist?," *NPR*, July 12, 2017, https://www.npr.org/sections/therecord/2017/07/12/536670493/spotify-is-accused-of-creating-fake-artists-but-what-is-a-fake-artist.

Nowhere has AI made more advances than in computer assisted art. In particular in the development of animation for movies and video games. What previously took teams of artists months or years to create can now be augmented by AIs and done in weeks. AI assisted artists can now create dynamic worlds that change and adapt to user interactions in two-dimensional and three-dimensional environments.

Musical composition has also been of great interest to scholars of AI. A number for research projects have sought to mimic, interpret, and build upon the works of artists from nearly every genre as described above. Entire concerts of AI generated music have been held often without the audience knowing that the composition was generated by an AI. The use of AI in music generation for movies, television, sound effects and more is of increasing interest to the entertainment industry. Not only is AI potentially able to reduce the cost of music composition, it can also be done more quickly.

The story of AudioIn is one that is becoming increasingly common. AIs are trained on the work of humans and then humans are subsequently replaced or augmented by the AI. It is a story of legitimacy, transparency, inequality, accountability, and equity.

Legitimacy: One of the initial questions posed and one that needs to be addressed at the outset is - what is art? Is it a purely human endeavor or is it something more? If it is a human endeavor, then does the substitution of AI generated content lack legitimacy? Can it be judged by critics, patrons, users in the same way as human generated art? If art is a reflection of something deeper in the mind or soul of a human is a machine capable of generating something of equivalence? Does viewing art or listening to a composition made by an AI undermine its value? Would you as an individual who just viewed a painting that moved you deeply and emotionally, feel violated if you were to later learn that the painting was created by an AI? Does a song that moves you mean less if you know it was that it was created by an AI rather than a human? These are fundamental questions residing at the heart of the debate in AI. These are questions that extend beyond the Turing and Lovelace tests and get at the core of what it means to be human and what it means for something to be art.

Transparency and Inequality: In the digital economy corporations often hold disproportionate power over independent contractors and users. This power asymmetry is particularly acute on platforms be they social media or streaming services. Understanding how your data as a contractor (musicians providing music to AudioIn) or as a user (generating data through interaction behaviors) is used can impact the decisions you make. A lack of transparency on platforms fosters systemic inequality between the platform operators and those who use the platform. If musicians had known that their songs would be used as training materials to eventually put them out of work would they still use the AudioIn platform? Would users who sought out music by artists still use AudioIn if they knew the platform was serving them music

generated by an AI? AudioIn did not give its users or contractors an option, their lack of transparency fed the roots of a power imbalance that undermined the economic livelihoods of musicians and the confidence of their users.

Accountability: When a firm violates the trust of its users and contractors does it have a responsibility to those whose trust it violated? How did AudioIn attempt to avoid accountability? Was its behavior ethical? How can firms such as AudioIn be held accountable for violations of trust? Corporate accountability extends across all forms of businesses and their actions. Honesty, is an important concept in building a brand. When firms supplant the expectations of their customers through deception whether using algorithms or the alteration of materials in a product the result is a decline in the perceived trust associated with a firm. When the trust of clients and even customers is violated there are often attempts by injured parties to seek legal redress.

Equity: Were AudioIn's practices fair? Is it fair to generate new content by training an AI from the work of others? Is it different to train a machine as opposed to a person? Does training a machine make it unethical? Was it fair for AudioIn to dilute its music catalogue to reduce its payments to artists? Fairness is a complex question in the context of AI-labor substitution. Balancing the interests of all actors involved in complex market arrangements can be complicated. Yet, undermining the livelihoods of groups through the co-optation of their work without their consent due to a structural imbalance in power undermines the equity of a transparent contractual relationship.

Flanagan, Andrew. "Spotify Is Accused Of Creating Fake Artists — But What Is A Fake Artist?" *NPR*, July 12, 2017. https://www.npr.org/sections/therecord/2017/07/12/536670493/spotify-is-accused-of-creating-fake-artists-but-what-is-a-fake-artist.

Sautoy, Marcus Du. *Creativity Code: Art and Innovation in the Age of AI*. Cambridge, MA: Harvard University Press, 2019.

Standage, Tom. "Monster in a Box." *Wired*, March 1, 2002. https://www.wired.com/2002/03/turk/.