Artificial Intelligence in Music Composition

William W. Saul

I. Introduction

Artificial intelligence (AI) is a pressing issue of our time, particularly for those in creative fields. In September 2023, the conclusion of a 148-day strike involving over 11,000 Hollywood screenwriters emphasized a growing demand for protective legislation from the perceived threats of AI. Many composers are also concerned about the impacts of AI but are unfamiliar with its musical applications. This essay provides an introduction to AI-generated music and its implications.

II. AI Terminology

The term artificial intelligence refers to "the capacity of computers or other machines to exhibit or simulate intelligent behaviour; the field of study concerned with this."[1] It is important to note the word "simulate", as AI systems are developed by programmers to accomplish specific tasks like problem solving or artistic expression in a human-like manner. The *Turing Test*, proposed by Alan Turing in 1950, is a methodology for judging if a simulation is distinguishable from a human. In this examination, a judge interacts with a machine and a human through text. The test is deemed successful if the judge cannot reliably identify the AI.

---

[1] Oxford English Dictionary, s.v. "artificial intelligence (n.)," December 2023, https://doi.org/10.1093/OED/7359280480.

Inspired by the human brain, AI researchers create digitized *neural networks*, which consist of layers of artificial neurons or nodes. These layers are the input layer, hidden layer (which handles computing), and output layer. The weight (i.e., strength) of connecting layers determines the output. A *large language model* (LLM) is a popular type of neural network, which collects massive amounts of text data. During 'pre-training,' the LLM utilizes *machine learning* (ML), a subset of AI that involves the unsupervised study of large datasets without a predetermined purpose. A specific type of ML, called *masked language modeling*, is used. This entails random masking of words in a sentence, relying on AI to fill in the blank spaces. The AI learns grammar and sentence structure by imitation, much like a human baby. After pre-training, the LLM is optimized for specific tasks, requiring more rigorous learning. Some examples of large language model use-cases include language translation, text analysis, and chatbots.

III. AI for Music Composition

1. How music relates to language

Since AI and neural networks specifically are involved in many text related activities, it is important to highlight the similarities between written music and language. The origins of early European music notation began with the dictation of the Gregorian Chant. "The evolution from a performance practice represented in writing, to a tradition of composing, transmission, and reading, took place over a span of centuries."[2] As notation evolved, it became more analytical, and unlike languages, shifted from being "…symbolic to iconic…" (i.e., abstract symbols

[2] Leo Treitler, "The Early History of Music Writing in the West." *Journal of the American Musicological Society* 35, no. 2 (1982), 237.

representing sound to graphical icons directly relating to musical elements).[3] Music and languages share an important similarity, however, which is that they use a finite number of signs (i.e., letters, note-heads, punctuation) to fit an infinite number of scenarios. This satisfies the necessary precondition for neural networks. AI can learn to imitate the use of pitches, for example, by analyzing large quantities of contextual data, such as compositions.

2. *History of AI for Music Composition (1958-2001)*

Ramón López de Mántaras presents a brief history of computer-generated music compositions in "Artificial Intelligence and the Arts: Toward Computational Creativity." Many early approaches utilized probabilistic processes like Markov chains to generate musical notes, testing them against music theory rules. In this instance, only the notes that adhered to the rules were retained. An example of this is Hiller and Isaacson's work on the *ILLIAC* computer in 1958. Through the use of a heuristic methodology, which attempted to mimic human-like intuitions, their work resulted in the "Illiac Suite", which is a computer-generated string quartet using classical harmony and counterpoint. Works based on Markov chains were rather limited in expressiveness, because of their reliance on randomness.

Other early researchers avoided probabilistic processes. In 1972, Moorer created complementary melodies and harmonies, by "simulating human composition processes using heuristic techniques…"[4] Levitt (1993) also argued against randomness. His work focused on creating a "descriptive language" that allowed for the expression of "musically meaningful

---

[3] Treitler, 266.

[4] Ramón López de Mántaras. "Artificial Intelligence and the Arts: Toward Computational Creativity." *The Next Step: Exponential Life,* (Madrid: BBVA, 2016), 4.

transformations of inputs, such as chord progressions and melodic lines…”[5] He achieved this through stylistic constraints, called ‘style templates.' This was embodied in a jazz inspired walking bass and ragtime piano player simulation. In 1974, Rader created a musical round generator, based on melodic and harmonic rules. What was historically significant in the implementation of these musical rules was its weighted component, allowing for the appropriateness of musical choices to be determined on the basis of contextual data. This is a similar idea to the weighted components used in neural networks.

Harmonization has always been a large area of focus for AI music compositional exploration. In 1969, Rothgeb wrote a program called *SNOBOL*. This program harmonized unfigured bass, by testing the logic of 18th century music theory. Ebcioblu created *CHORAL* in 1993, a system that can “harmonize chorales in the style of J. S. Bach.”[6] It works by the user inputting a melody, which CHORAL is then able to harmonize. Notably, this system uses logical primitives (fundamental operations in formal logic) to differentiate between musical elements. *MUSCAT*, a system utilizing neural networks to study musical harmony theory, was created by Bharucha in 1993. MUSCAT is capable of understanding the anticipatory role of dominant chords moving to the tonic. This enables a rating hierarchy of harmonic resolutions, based on satisfaction, effectively connecting theory and emotionality.

In 1998, Sabater et al. observed that purely relying on harmony rules was ineffective in generating compelling musical compositions, because “the rules do not make the music, it is the music that makes the rules.”[7] In his system, the computer first underwent a process of comparing

---

[5] López de Mántaras, 4.

[6] López de Mántaras, 6.

[7] López de Mántaras, 7.

a given harmony with a database of existing compositions. If a similar scenario was discovered, the computer would favor the existing solution over one from an academic source. If there wasn't an equivalent, the computer searched through relevant harmonic theory. Successful results were added to the database so that the AI could optimize overtime.

David Cope's *EMI* project, between 1987 and 1990, created compositions "in the styles of Cope, Mozart, Palestrina, Albinoni, Brahms, Debussy, Bach, Rachmaninoff, Chopin, Stravinsky, and Bartok."[8] The AI discovered patterns across multiple pieces by a composer, labeling them as 'signatures.' These signatures, which were tagged with contextual data from the larger composition, acted as a template for writing new work. Technical rules were also generated, though primarily from the compositions, rather than a textbook. This is essentially a statistical model of generating musical ideas like motifs and harmonic progressions, based on preexisting styles. To encourage cohesion between musical elements, EMI carefully avoids irregularities in voice leading, pitch range, etc. "The results, although not perfect, are quite consistent with the style of the composer."[9]

3. *Recent Developments in AI for Music Composition*

As AI for music composition matures into commercial viability, its development teams and target customers have shifted. Large corporations like Alphabet (Google's parent company) and OpenAI are now doing most of the groundbreaking research in AI music creation, due to their enormous resources. OpenAI has two major music composition projects, *Jukebox* and

---

[8] López de Mántaras, 7.

[9] López de Mántaras, 8.

MuseNet. Jukebox responsively generates music in various genres and instrumentations, in accordance with input conditions. It is a neural network, drawing upon large quantities of audio information. MuseNet also uses a neural network architecture, but instead learns from MIDI data to identify musical patterns and rules.[10] MuseNet creates compositions up to 4 minutes in duration, 10 instruments in size, and writes in "styles from country to Mozart to the Beatles".[11]

The most significant AI music creation advancements are associated with DeepMind, an Alphabet owned research lab. DeepMind has recently partnered with Google's YouTube platform, enabling an ongoing project called *Lyria*. YouTube provides an huge database of music, which Lyria analyzes using its neural network structure. Lyria is capable of transforming singing, beatboxing, or instrumental performances into complex renditions on digital instruments. An example published on YouTube involves Lyria turning a vocal recording by Louis Bell into a digital saxophone performance of the same melody. Lyria generates four layers of drums, bass, keys, and singing, based off Bell's overdubs. In addition to transforming beatboxing into drum loops, vocals into instruments, or MIDI pianos into choirs, Lyria can also orchestrate simple vocal performances. In a demo, melodic humming was performed and harmonized with strings and brass.

Lyria's interface is very dynamic, allowing musical styles, instruments, and other edits to be made on the track level. It also appears to use descriptive tagging, like those on YouTube videos. Lyria is currently being developed in collaboration with musicians, who provide insights

---

[10] MIDI (Musical Instrument Digital Interface) is a protocol for communicating information like pitch or velocity between musical instruments and computers.

[11] Christine Payne, "MuseNet," April 25, 2019, https://openai.com/research/musenet.

and recording samples. One of DeepMind's stated objectives is to "set the standard for the responsible development and deployment of music generation tools."[12]

IV. Creativity and Implications of Musical AI

1. AI can be musically creative; practically speaking

In "The Intersection Of AI And Human Creativity: Can Machines Really Be Creative?", Bernard Marr cites the Oxford dictionary, saying creativity is "The use of imagination or original ideas to create something."[13] He proceeds to declare that "imagination" is "the genesis of creativity", but mentions that imagination is "largely still a mystery", and that the significance of creativity is the thought process behind making things.[14] I differ slightly with this assessment. For one, many dictionaries like Cambridge and Merriam-Webster define creativity without mention of imagination.[15] [16] Since the *genesis* of creativity and creativity itself are two distinct categories, one has to question if it is appropriate to presuppose imagination, given creation. For the sake of argument, I will grant that imagination is necessary. If a human and an AI subject both produce music compositions that pass a Turing test, based on rules from a contrapuntal or

---

[12] "Transforming the Future of Music Creation," Google DeepMind, November 16, 2023, https://deepmind.google/discover/blog/transforming-the-future-of-music-creation/.

[13] Bernard Marr, "The Intersection of AI and Human Creativity: Can Machines Really Be Creative?," *Forbes*, March 27, 2023, https://www.forbes.com/sites/bernardmarr/2023/03/27/the-intersection-of-ai-and-human-creativity-can-machines-really-be-creative/.

[14] Marr.

[15] Cambridge Dictionary, s.v. 'Creativity,' December 20, 2023, https://dictionary.cambridge.org/dictionary/english/creativity. Reference for alternative definition of creativity.

[16] Merriam-Webster Dictionary, s.v. "creativity," accessed December 24, 2023, https://www.merriam-webster.com/dictionary/creativity. Reference for alternative definition of creativity.

tonal harmony textbook, why would one piece be considered a creation, but not the other? As Marr later acknowledges, this becomes a semantic reductio ad absurdum.[17]

While Generative AI is not creative (if biological processes are presupposed), I would argue a more nuanced point; that AI is *practically* creative. AI is capable of producing intricate, novel, and valuable data, reliably. Ergo, for all intents and purposes, AI should be considered as such. Critics often argue that AI only arranges preexisting materials, therefore the output data is not truly novel, and therefore not creative. This standard becomes a 'no true Scotsman,' however, when applied to most music.[18] Almost every musical innovation builds on previous knowledge and materials, which are descendant of humans.

In the 1920s, serialism became a novel way of composing with pitches of the 12-tone chromatic scale. By organizing notes into 'tone rows' (cycles of all 12 pitches without repetition), atonality was achieved. Though the pitches and instruments preceded this methodology, composers like Arnold Schoenberg arranged them in a new and interesting ways. Post-World War II, serialism became more influential and composers like Pierre Boulez expanded the concept of rows to control rhythm and dynamics. Though Boulez applied a preexisting concept to preexisting parameters and instruments, these explorations were very creative. Thousands of composers have used Schoenberg and Boulez's systems to make original rows and pieces. Despite the lack of novelty in pitch, rhythmic, dynamic, or instrument usage, most composers would still consider this creating music.

---

[17] Marr. Bernard Marr describes arguing for or against the position that AI is creative as "at this point, really just a matter of semantics."

[18] A no true Scotsman is a logical fallacy, in which one excludes something from a group based on arbitrary reasons, often to maintain consistency in an argument.

Musical AI is currently unable to create ideas as potent as serialism, at least not without guidance from a human user. It is, however, able to recognize and compose in specific musical styles. The capabilities of AI appear to be between that of a relatively generic composer, and someone who expands on existing concepts in a unique way. One only needs to examine the homogeneity of many popular music genres to grasp the contradiction of considering a composer of generic music as creative, while dismissing the potentiality of AI.

2. Implications of AI Music Composition

Musical AI is increasingly capable of composing classical, and as of more recently, contemporary pop genres. Associating the use of pitches, rhythms, and instrumentations with musical genres, composers' styles, and emotional implications, AI can generate music that fits a wide range of needs. Currently the most complete work is being done with neural networks, using similar logic as LLMs. Therefore, it can be assumed that the most common musical ideas will be the most optimized, because they provide the most data. As AI musical technology is deployed, the intention may be one of responsibility; however, supply and demand economics will ultimately determine the monetary value of musical skills.

One noteworthy area to track is library music, also called production music. In this field, composers upload pieces of music to digital libraries, where customers can license them for media work. A 'sad ballad' could be purchased by a company for use in advertising, or a hip hop beat could be licensed for someone to rap over. Since the length and style of a piece are not customized for the media project, this is not ideal in many cases, but it does offer predictability and affordability. For example, many film students use library music in their academic work,

instead of hiring media composers. Libraries provide customers with a variety of generic music for their projects, quickly, reliably, and with low costs.

Below is a simple template I made for creating a sad ballad, which could be used as a starting point for creating library music.

Tempo/meter:	**Harmony:**	**Melody:**	**Instrumentation:**
Approx. 60 bpm, duple or triple meters (3/4, 4/4, etc.).	Minor key/aeolian, slow rate of harmonic change.	Mostly diatonic.	Mainly strings, and or piano.

If a sophisticated AI tool was given a prompt to create a sad ballad, it would likely make similar calculations, based on a large database of media scores. AI is already capable of sorting emotions like sadness by application, so the user could also specify if the use-case was for film/ television, or something else. A user could demo the material quickly, and then decide to make any revisions. As an example, maybe the user decides the music needs to be more dramatic. The AI could update the template accordingly:

Tempo/meter:	**Harmony:**	**Melody:**	**Instrumentation:**
*Approx. 80 bpm, duple or triple meters (3/4, 4/4, etc.).	*Minor key w/modal modulation, slow rate of harmonic change.	*Follow chord changes.	Mainly strings, and or piano.

If an AI tool could make these determinations, it could offer some unique advantages over traditional library music. In its infancy, the AI would likely produce highly generic compositions, but it would offer ample opportunities for revision. For example, a customer could specify parameters like duration or dynamics. The more interaction a user has with the AI, the more customized the outputs would become. This method would also be at a lower operating cost, possibly by subscription, and generate large quantities of music faster than a human

composer. Much like many Top 10 product websites are now run with AI, this type of platform could be useful for mass generating media.

Another area that may be impacted is orchestration or arrangement. My chosen software for music notation, *Dorico*, has recently introduced a feature called "Generate Notes from Chord Symbols."[19] After inputting chord symbols, the program can distribute notes from these chords across a selection of instruments, using common doublings. As discussed, Lyria is capable of orchestrating melodies from audio files, but its full potential is still unknown. While orchestrating AI is unlikely to categorically impact orchestrators, these technologies allow artists to add simple arrangements to their work, without an orchestration or music theory background.

Composers who are unlikely to be displaced by AI in the foreseeable future include media scoring and fine arts/concert composers. The former category entails extremely intricate processes like tempo mapping to markers, 'Micky Mousing,' or creating dynamic, nonlinear compositions for games.[20] While AI is already providing workflow optimization tools for media composers, such as automatic picture cut identification in digital audio workstations, nothing has come close to replacing human composers. The fine arts/concert composers in the latter category possess a strong propensity towards experimentation and contrarianism. Much like the invention of photographs helped inspire a move away from realism into impressionist art, I suspect these composers will adapt similarly to influence of AI in music.

---

[19] Generating Notes from Chord Symbols," Dorico Help Center, Steinberg, n.d., https://steinberg.help/dorico/v4/en/dorico/topics/write_mode/write_mode_arranging_tools/write_mode_notes_generating_from_chord_symbols_t.html.

[20] Micky Mousing is a technique that synchronizes musical figures with visual movements. For example, a character slipping on a banana peel being scored with an upward glissando and cymbal crash.

V. Conclusion

AI music composition has advanced exponentially in recent decades. Current technology is capable of generating compositions in a multitude of styles, contributing melodies, harmonies, and instrumental arrangements to existing work, and relating aesthetic or emotional information to musical materials. While AI may not be creative, depending on one's own definition, it is effectively producing creative outcomes. Furthermore, AI for music composition will likely impact creators of generic music, because of its quickly improving abilities to model popular music genres. Library music was given as a specific example for this. However, there is little indication that AI will substantially affect media or fine arts/concert composers. AI's potential is great, and while there is speculation as to its upper limit, it is safe to conclude that it will soon affect the lowest common denominator of musical creativity.

Bibliography

Cambridge Dictionary. s.v. "Creativity." December 20, 2023. https://dictionary.cambridge.org/dictionary/english/creativity.

Google DeepMind. "Transforming the Future of Music Creation." November 16, 2023. https://deepmind.google/discover/blog/transforming-the-future-of-music-creation/.

López de Mántaras, Ramón. "Artificial Intelligence and the Arts: Toward Computational Creativity." *The Next Step: Exponential Life*. Madrid: BBVA, 2016, pp. 4-8.

Marr, Bernard. "The Intersection of AI and Human Creativity: Can Machines Really Be Creative?" *Forbes*, March 27, 2023. https://www.forbes.com/sites/bernardmarr/2023/03/27/the-intersection-of-ai-and-human-creativity-can-machines-really-be-creative/.

Merriam-Webster Dictionary. s.v. "creativity." Accessed December 24, 2023. https://www.merriam-webster.com/dictionary/creativity.

Oxford English Dictionary. s.v. "artificial intelligence (n.)." December 2023. https://doi.org/10.1093/OED/7359280480.

Payne, Christine. "MuseNet." April 25, 2019. https://openai.com/research/musenet.

Steinberg. "Generating Notes from Chord Symbols." n.d. https://steinberg.help/dorico/v4/en/dorico/topics/write_mode/write_mode_arranging_tools/write_mode_notes_generating_from_chord_symbols_t.html.

Treitler, Leo. "The Early History of Music Writing in the West." *Journal of the American Musicological Society* 35, no. 2 (1982): 237-266.