In what ways would music composed by a computer intelligence differ from human compositions?

Curtis Ullerich
Philosophy of Cognitive Science

Music has a rich and varied history that extends back over a millennium in the Western tradition alone. Throughout this time it has undergone systematic advances, introducing new complexities and variations at every turn. Music is one of the most mystic aspects of our culture, bypassing rationality and accessing our deepest emotions. The substantially more recent science of artificial intelligence first intersected music several decades ago in the 1950s as generative programmed music emerged for the first time [15]. The science of artificial intelligence can often be as cryptic as our relationship with music, making their combination in the form of generative music and AI composition philosophically and cognitively ambiguous. Through the exploration of this combination, though, we may unravel a self-woven web of assumptions about our relationship with music and perhaps uncover that we exhibit a more mechanical processor of the beauty of music than we may care to think, or discover the creative abilities of machines we traditionally see as cold and calculating. While work has been and is being done in the field of AI composition, many questions remain unanswered. Because this exploration takes the young science of artificial intelligence into camp with an art that is not scientifically understood on many levels, it requires philosophizing in tandem with empirical research. With this combined approach we may discover how our own cherished works of past and present differ from music composed by an artificial intelligence.

More precise sensory input thresholds would cause scores to be more nuanced and specific.

The ability of mechanical sensors to discriminate between differences in musical parameters can be much higher than that of human sensory organs. In the case of android composers, their presumably more precise sensory recognition would allow and even stimulate the inclusion of more exact and literal musical gestures. Western music notation is an extension of what Western culture finds musically pleasing and necessary to convey direction to a performer. Western notation, though, is a "gross oversimplification" of actual musical performance [1]. Though this travesty is not exclusively Western, examining it provides a good example of the rift between what composers write and musicians play.

The Western system allows for notation of several parameters of performance; namely pitch, rhythm, tempo, dynamics, and duration. For example, one minim on the line above the treble staff at 120 beats per minute, taken explicitly, calls for the pitch of 442 Hertz to be played for one second. In reality, what is played depends on various qualitative factors and stylistic traditions. Most musicians would play the note with vibrato, (intentional pitch oscillation) to "colour" and "shape" the note. The amplitude of this oscillation will even vary throughout the duration of the note. Even the concept of a crescendo varies from one composer or performer to another. Artistic interpretation is a huge part of musical communication. A crescendo, for instance, may be realized through a linear, parabolic, exponential, polynomial, or any number of other volume changes. An android composer would have the sensory resolution necessary to make such nuanced directives meaningful.

The very existence of rubato (from the Italian “stolen”) as musical idea demonstrates the lack of dictated nuance available to composers in the Western rhythmic notation tradition. Rubato, musically interpreted as “stolen time” signifies to a musician that the marked tempo is not to be taken in strict consideration, thus rendering the notation system as written inadequate. Performance analyses reveal that a performance can vary "up to 50% of the notated metrical duration in the score" in a stylistically correct interpretation of a piece [5].

Human music cognition effectively takes maximal input, which is defined by the physical limitations of the listener's ears. The capabilities of mechanical sound processing are far beyond human abilities, meaning that if an android used maximal input, human performance would be imprecise and limited in range. This would lead to android compositions that specify this level of precision. Such a development in notation could potentially make music unplayable by human performers, which would lead to a schism in human/android composition. Of course, androids would also have mechanical limitations, so a similar limit of compositional precision would occur within android culture. Human performers would reasonably reject android compositions structured this deterministically because they would not allow for interpretation and therefore would lack freedom in performance. Indeed, robot performers have been created that reflect a human amount of dexterity and provide extended precision and ability in each parameter of music [13].

Differences in composition due to sensory input would likely manifest in more ways than parametric nuance. The specific resolution of frequencies could be treated differently by an android processing sound input. Humans perceive melodies as pitch invariant, in that a tune can be sung in any key and remain unchanged in our minds. Mechanical frequency sensors would allow or perhaps even cause an artificial intelligence to discriminate between keys or octaves in this way. In fact, there are a number of invariances in human music cognition; a tune replicated with variance in pitch, tempo, or amplitude (as these are the most easily understood parameters by those without a background in music theory) will retain its musical identity to a listener. It is possible that changes in any of these parameters would cause an android musician or composer to conceive of the tune as fundamentally musically different [12].

A computer's inherently different culture (as compared to that of a human being) would change the parameters and measurements of those parameters that are considered to be fundamentally musical by a human listener.

Cultures define music based on the parameters as previously discussed. Each culture places different importance on various parameters and even represents and realizes them in different ways. Current Western music, for example, uses a twelve-tone system for pitch in which a range of sound frequencies is divided logarithmically into twelve notes known as an octave. Arabic music, however, is based on a 24-tone system, also divided logarithmically. Western music places emphasis on melody, or the ordering of sequential pitches, whereas African music largely focuses on rhythm over melody. Western music is stereotypically based on the "pitch/duration paradigm" in which the frequency, length, and order of notes are seen as the most important aspects of a piece of music [4].

An android culture could focus musically on purer manifestations of sound, such as white noise or Schoenberg's twelve tone technique. White noise is a sample of sound that contains equal distributions of all frequencies, and the twelve tone technique is a compositional strategy in which all twelve tones in an octave are represented equally in order to avoid emphasizing any one key or set of keys. Because of the ease with which computational systems can represent pitches, it is plausible that androids would first (if not predominantly) focus on pitch oriented music. The harmonic series present in naturally produced sounds has a mathematical beauty that is readily represented logarithmically. Because of this inborn harmony between mathematics and music, androids could reasonably prefer such easily understood systems over the complex waveforms involved in processing and rendering the sound samples used in acousmatic music, for example.

While humans tend to appreciate musical depictions of nature (as in the birdsong present in Vivaldi's "Spring"), androids may find metallic and machine-like sounds to be closer to their hearts, both metaphorically and physically. Androids may even find human compositions to be a curiosity and deconstruct their musical form, effectively running the gamut of musical form throughout history in reverse, beginning at the present and moving back until arriving at the first musical endeavor of chant. Androids would possibly disregard all current compositions and begin with pure atonal music, subsequently developing a completely new systems of pitch, if not new musical parameters altogether.

It is reasonable to suggest that the development of culture requires self awareness, and many will argue that this theory is moot because androids do not exhibit this characteristic. While current artificial intelligences are not deemed to be self aware, the plausibility of such a development in the future validates the above speculation on this subject.

AI compositions would not be creative or original, whereas human compositions are.

Perhaps the easiest compositional model to create is that which is based on randomness. Computers are very good at employing pseudo-random techniques, and when applied to a statistically analyzed database, simple recombination of basic patterns is a natural first step in creating a composing machine [17]. However, this leads to the criticism of whether or not the machine is actually composing music or just rehashing pieces from its corpora. This challenge is somewhat ambiguous because of the recombinant nature of composition. While each composition can be considered to be unique, components from previous pieces can be found in any composition. In fact, without such borrowing and replication, musical compositions would have no cohesive form by which to create a style of music. Without such recombination music as we know it simply couldn't exist.

Musical creativity often manifests through the idea of learning rules in order to break them tastefully. It is this kind of novelty that causes a listener's ears to perk up, because his or her musical expectations for the piece have been violated in an interesting way. Some models of music cognition (as well as general AI) have been built on the principle of craving novelty in order to foster a growing interest in learning and exploring the world. This kind of novelty can be programmatically introduced into an artificial intelligence mostly simply by assigning expected values to the outcomes of events. If the value of the outcome violates the expectation, then the intelligence considers that outcome to be novel. Musically, these values are predefined based on the rules of note leading and counterpoint. The organization of Western music notation lends itself handily to this kind of novelty analysis. However, novelty is distinct from creativity.While an outcome may be considered to be novel, it may not be creative. Creativity, while also novel, has a more stringent qualification [11]. Novel ideas and outcomes are oft repeated, whereas a truly creative idea is original as well as interesting. Describing creativity in this way branches into subjectivity, which itself is a key factor in the uncertain realm of creativity.

Because of this extra difficulty in being creative, the problem of programming a creative AI becomes much more difficult. Novelty, especially in music, is relatively simple, but creativity seems, at least superficially, to be a more human endeavor, perhaps because of the ambiguous classification of "creative" ideas. Such a criticism is handily retorted, however, by the epistemic challenge of whether or not human composers actually create new music, or just recombine and realize new pieces from musical memory or other environmental input. Again, this deals with the same distinction between novelty and creativity. The layers of obfuscation around the human creative process provide a similar layer of difficulty as the problem of pushing beyond novelty and into creativity in artificial intelligence. With this challenge, it seems that androids would be as adept at composing creatively as human musicians.

AI compositions would not be different than human compositions.

Significant work has been done in the field of music intelligence by University of California, Santa Cruz professor David Cope. He developed his original brainchild, Experiments in Musical Intelligence (EMI; anthropomorphized by Cope as Emmy), in 1981 in order to overcome a case of composers' block. EMI began as a strict rule-based algorithm that generated Bach-like four-part chorales. The generated music sounded like music theory exercises, as Cope did not program the behavior of intentional error into Emmy [9]. Human composers traditionally create interest in music by deviating from standard rules of music theory and violating listeners' expectations in a novel way [1]. Because Emmy did not take this into account in composing, listeners found her works to be predictable and uninteresting [2]. After implementing the ability for such deviation, Emmy's output improved musically. Cope went on to develop Emmy into a "data-driven" recombination program that took corpora of composers' collected works and produced new pieces in their style [6]. Cope played various pieces created by Emmy in concerts alongside Bach originals. Many of the audience members couldn't tell the difference between the Emmy and Bach compositions. This is philosophically significant, because as far as composition is concerned, this means that Emmy passed the "musical Turing Test" [2][9].

Partly due to philosophical criticisms of Emmy within the music community, Cope decided to halt development of Emmy in 2003 and destroy her musical corpora to begin work on a new intelligence, dubbed Emily Howell, that would create music in "her" own style, learning and being trained by Cope's feedback [6]. Emily Howell has since produced a record of original music called *From Darkness, Light* [7]. The first concert of Howell's work, in which Cope neglected to reveal her as an algorithm, was well received. Audience members commented to Cope that the performance had been beautiful and moving, comments that unknowingly declared Emily Howell the victor in this round of the musical Turing Test [7].

The results of Cope's work demonstrates that music composed by an artificial intelligence does not differ from human composition at the level noticeable by a human listener. This is not to say that AI music *cannot* differ, but that an artificial intelligence can compose music of the same quality and characteristics of classical composers in the Western tradition.

One criticism of this claim is that neither Emmy nor Emily Howell were intelligent, but mere recombinance algorithms that did not display originality or creativity in their work. While this may be true of early manifestations of Emmy, Emily Howell was specifically designed to learn, adapt, and compose in her own style. Cope trains Emily Howell by selecting from her output the best works and providing a feedback loop by which her style evolves.

Because of this relationship Cope has with Emily compositionally, the claim that her output is an object of Cope's own intelligence is natural. However, this relationship is one of tutor and student, in which Cope trains Emily's developing mind. Conceivably, AI compositional technique could diverge once the need for human training is no longer present and an AI musical culture develops, as discussed above.

Many detractors of Cope's work say that generative compositions such as those by Emmy and Emily lack emotion and have no "heart or soul." This view tends to be held by those xenophobic of an artificial intelligentsia. Effectively, if this view is true, then it is impossible, by definition, for an artificial intelligence to truly compose music. Such a constrictive view of music, if accepted, appeases those xenophobic of advances in strong AI, but would seem to hamper the development of music itself by disregarding a relatively new and potentially insightful foray into music cognition and theory.

Computers could create music both in ways humans have never considered and in ways impossible to humans due to physical limitations [1].

What an AI *can* do is often unpredictable, but that does not mean the possibility or capability was not coded upon creation--we just do not necessarily understand what that program we are creating can do. Conway's Game of Life, in which complex pattern emerge based on very simple rules, is a good example of this [14]. Extrapolating from this example, an AI programmed to compose simple music has the potential to create new and complex pieces in unexpected ways.

With the advent of stereo music systems in the 20th century, the possibility of spatialization of music in recordings became possible. From a simple two channel recording, works composed in the electroacoustic medium can have dozens of distinct channels. Surround sound theatre systems are a more common manifestation of multi-channel audio. AI composers can have the on-board processing power necessary to realize works initially in multiple channels, a task currently only possible to human composers through the use of much computer processing [1].

Another recent advancement in AI in music is a direct result of the capability of computers to process a musical performance in real time on a scale not possible by human brains. A marimba playing robot named Shimon built at the Georgia Tech Center for Music Technology improvises alongside live musicians [8][10]. This puts time constraints on the processing required to generate new music and requires more refined methods of pattern recognition than purely compositional intelligences as it responds to input in real time. This real time interaction is the very motivation the researchers at Georgia Tech had for creating Shimon: They intended to use the computational power built into Shimon along with human aesthetic judgment to create new music, allowing the two improvisational approaches to rebound off of one another to supplement the creative process [16]. This new kind of composition would not be possible without human utilization of artificial intelligence.

The case of Shimon brings the issues of music cognition to the forefront in a way hardly relevant to most previous research in music intelligence, particularly that of Cope. The general historical framework of musical composition has been to set notes to a page for later performance. Rarely is it the case that the score remains unchanged from the original concept, and even more rare is the ability of a human composer to write out an entire piece in one thought with no or minimal corrections. The ability to process music this quickly and accurately requires many decisions by the designer of the intelligent system, particularly when scaling the input to thresholds similar to human limitations in order to allow collaborative performance.

Conclusion

While artificially intelligent composers seem to be able to create music in a remarkably (if not identically) human way, the most likely route for AI composers to go would be an expansion of musical style beyond what is physically or computationally possible by humans. Androids in particular, with sensory input and the capability for tactile interaction with music and their own compositions, will have a strong chance of diverging from standard composition.

It is worthwhile to remember that while we in the modern Western world have (generally speaking) one concept of music, there exist a multitude of cultural definitions of music already. Artificial intelligences and androids in particularly may very well find their own niche in the musical world distinct from anything we currently recognize. Because of the cultural rifts in music cognition between humans and mechanically advanced androids, we may find their music to be strange and even otherworldly. But before such developments are dismissed as abominations of modern art, we would do well to explore and attempt to appreciate already existing music such as the musique concrète of the French, music of the Aborigines, or the throat singing of the Tuvans. Expanding our own appreciations and knowledge of what already exists musically, as well as imagining and attempting to realize what could come to be may be our best hope to understand the potential music of the future and, indeed, develop it through research in music cognition and artificial creativity.

Works Cited

[1] Dobrian, Chris. "Music and Artificial Intelligence." Music.arts.uci.edu. 1993. Web. 05 Mar. 2011. <http://music.arts.uci.edu/dobrian/CD.music.ai.htm>.

[2] Blitstein, Ryan. "Triumph of the Cyborg Composer." Miller-McCune. 22 Feb. 2010. Web. 07 Mar. 2011. <http://www.miller-mccune.com/culture-society/triumph-of-the-cyborg-composer-8507>.

[3] Dannenberg, Roger B. "Artificial Intelligence, Machine Learning, and Music Understanding." Print.

[4] Wishart, Trevor, and Simon Emmerson. On Sonic Art. Amsterdam: Harwood Academic, 1996. Print.

[5] Desain, Peter and Honing, Henkjan. Music, Mind and Machine. Amsterdam: Thesis Publishers, 1992. p. 30.

[6] Cheng, Jacqui. "Virtual Composer Makes Beautiful Music—and Stirs Controversy." Ars Technica. 29 Sept. 2009. Web. 25 Feb. 2011. <http://arstechnica.com/science/news/2009/09/virtual-composer-makes-beautiful-music and-stirs-controversy.ars>.

[7] Ahmed, Murad. "Emily Howell, the Virtual Composer Making Waves in the Computer World." The Times. 22 Oct. 2009. Web. Mar. 2011. <http://technology.timesonline.co.uk/tol/news/tech_and_web/article6884631.ece>.

[8] Van Buskirk, Eliot. "Robots Pass Musical Turing Test | Listening Post." Wired.com. 21 Nov. 2009. Web. Mar. 2011. <http://www.wired.com/listening_post/2008/11/no-way-robot-ja/>.

[9] Orca, Surfdaddy. "Has Emily Howell Passed the Musical Turing Test?" H Magazine. 22 Mar. 2010. Web. Mar. 2011. <http://hplusmagazine.com/2010/03/22/has-emily-howell-passed-musical-turing-test/>.

[10] Streit, Valerie. "Robot Musician Improvises, Jams with Humans - CNN." Featured Articles from CNN. 29 Apr. 2010. Web. Mar. 2011. <http://articles.cnn.com/2010-04-29/tech/robot.musician_1_humans-and-robots-shimon-music-technology?_s=PM:TECH>.

[11] Boden, Margaret. "SEHR: Creativity and Unpredictability." Constructions of the Mind 4.2.Stanford University. 4 June 1995. Web. 01 May 2011. <http://www.stanford.edu/group/SHR/4-2/text/boden.html>.

[12] Dorrell, Philip. "The 6 Symmetries of Music." Philip Dorrell's Home Page. 26 Aug. 2005. Web. Mar. 2011. <http://www.1729.com/lists/SymmetriesOfMusic.html>.

[13] "Toyota Robots Play Trumpet and Violin In Tokyo (video) | Singularity Hub." Singularity Hub | The Future Is Here Today…. Web. 25 Mar. 2011. <http://singularityhub.com/2011/02/24/toyota-robots-and-humans-make-beautiful-music

-together-video/>.
[14] Gardner, Martin. "The Fantastic Combinations of John Conway's New Solitaire Game "life"" Scientific American 223 (1970): 120-23. MATHEMATICAL GAMES. 4 Apr. 2010. Web. Mar. 2011. <http://ddi.cs.uni-potsdam.de/HyFISCH/Produzieren/lis_projekt/proj_gamelife/ConwayScientificAmerican.htm>.
[15] Hass, Jeffrey. Introduction to Computer Music. Vol. 1. Indiana University. Web. Mar. 2011. <http://www.indiana.edu/~emusic/etext/digital_audio/chapter5_digital.shtml>.
[16] "GTCMT » Shimon – The Perceptual and Improvisational Robotic Marimba Player."GTCMT. Web. 01 May 2011. <http://gtcmt.coa.gatech.edu/?p=628>.
[17] Cope, David. Computer Models of Musical Creativity. Cambridge, MA: MIT, 2005. Print.