Musical Composition: Man or Machine?
The Algorithmic Composer
Sam Agnew

The use of computers in the composition of music is a topic of controversial debate among scholars, musicians and composers as well as Artificial Intelligence experts. Music is generally thought of as an artistic field of human expression, which is seemingly far from being related to Computer Science or Mathematics. Contrary to popular belief, however, these fields are much more related than they seem. As exemplified by composers such as David Cope or Bruce Jacob, algorithms and computer programs can be used to create new pieces of music, and are effectively able to emulate this form of expression previously exclusive only to humans. In this research, I intend to explore the use of computers and algorithms in the composition of music. I will touch on some of the key issues that are often linked with the use of computers as compositional tools, and I will explain these breakthroughs in the field of Artificial Intelligence and Creativity as well as in Musicology and other related musical fields. Computers process information algorithmically, through a finite set of logical instructions. This algorithmic process can be extended to realms outside of the computer as well, exemplified by musicians such as John Cage and his use of randomness to compose music. Further breakthroughs in the field of Artificial Intelligence range from producing new musical works or imitating the style of well-known musicians, to simpler applications such as producing counterpoint for a given melody, or providing variations on a given theme. I will discuss all of these subjects and examples as well as the philosophical implications of what it means for a "machine" to create music. Music is one of the things that is thought of as being exclusive to our kind, and the fact that there exists mechanical, deterministic machines that can create pieces of art shocks people and causes a great deal of distress in the world of scholarly musical debate.

This paper will specifically focus on determining to what degree computers are able to compose music and algorithms are able to model human creativity through musical expression. Other discussion will include debate concerning whether or not a machine can produce a piece of music that is indistinguishable from a human composition in the same style. There is also the philosophical question of determining who the "real composer" of a piece of music is, if that music was algorithmically composed, and whether or not the music is authentic. I will also define the boundaries of the filtering that a musical composer does to the output of a computer program, and whether a "line" can

be drawn to determine when a person is genuinely composing music, or simply recognizing music, created by the programmer, that he or she enjoys, and editing it. Contrary to popular belief, machines are fully capable of composing music because music is an art form that naturally lends itself well to the act of computation.

It is assumed that music is innately human and unscientific, but in fact, algorithms are anticipated in many forms of early music. Many aspects of this musical composition are very straightforward and logical, lending itself quite well to algorithmic implementation. Some composers have subconscious patterns that they follow every time they create a new composition. This implies that certain aspects of musical composition are algorithmic even for human composers. This is nothing new, and music has been this way for hundreds of years. Even during the eleventh century music was composed algorithmically, although composers of the time were most likely not aware of the mathematical relationships governing what they were doing. Guido of Arezzo, a musical theorist of the medieval era, used an algorithmic method to derive plainchant from text through ascribing pitches to vowel sounds in order to fit the music to the text (Alsop 90). This was even before algorithms were commonly thought about outside of specific uses in Mathematics. Other musical concepts used by many composers over the centuries, such as tonalism and serialism, are also highly algorithmic and deterministic (Alsop 90). The mathematical nature of many aspects of music theory allows for smooth transition to computational processes. Musicians also subconsciously, and sometimes consciously, utilize algorithmic search spaces, sets of possible musical options in this case, when they are "engaged in the act of composition," and this is a highly mechanized thought. In this aspect, the use of a computer program would provide greater musical possibilities because "the practice of algorithmic and/or computer aided composition certainly does provide a much greater search space from which to reflect upon the place and value of optimization in their work" (Davismoon 238). This is useful in "optimizing" a piece of music, although in this case it would be more beneficial for a human and a computer to collaborate on a composition because optimization throughout an entire piece of music would seem boring to the human ear. However, this creates a great opportunity for human composers to work together with computers in order to enhance their musical creations. This collaboration goes even further because similarly to how algorithms were used to derive plainchant as previously discussed, algorithms may also be used to create polyphony to coincide with and add to that plainchant.

This use of computation is overwhelmingly apparent in the development of certain forms of polyphonic counterpoint due to the algorithmic nature of such musical techniques, and this is exemplified by the work of Adiloglu and Alpaslan's work on machine learning as it applies to two-voice counterpoint composition. This translates so well to computation because the rules of counterpoint state defined sets of instructions explaining how two musical lines can interact. These two researchers created a computer program that produces pieces of first species counterpoint, which is described by Johann Joseph Fux in his rules of counterpoint, as being the most deterministic form of counterpoint (Adiloglu, Alpaslan 301). The purpose of their research was to develop a new method of representing notes using computers. The program successfully learns the behavior of the input, and does not utilize intervals that are considered "hollow" or "boring" to the human ear, and it also generates a cadence, which is extremely common in this type of polyphonic music. How effective a machine is at producing music depends on several things, and as Adiloglu and Alpaslan have discovered through this research, "a good knowledge representation results in a strong learning capacity which is crucial for successful applications" (301). As shown by this program, music that has a set of rules to follow translates even more easily to computer composition than other styles, and this is true for other aspects of the compositional process that are more naturally algorithmic as well.

Musical form can also be considered formulaic or deterministic, which may seem obvious upon further thought because the very idea of form, and how a piece is structured, is subject to specific rules, which transition perfectly to the use of algorithms. Some composers do this intentionally, and others subconsciously fit their pieces of music to specific forms. (There are cases where a piece of music does not follow any particular form, or even has no order whatsoever in some Avant-Garde pieces, but these cases are not quite as relevant to the current discussion, as chaotic form would be even easier for a computer to simulate simply by lack of structure analysis, and if an algorithm produces this type of structure it may even be due to laziness of the computer programmer!) Musical form may be perhaps one of the most structured aspects of music, and this concept is expressed by many scholars on the subject of algorithmic composition. Historically, much writing on musical form has taken place and has also focused on compiling standard models (Collins 104). These templates make the structuring of music by computers that much easier, because they already exist and do not have to be calculated by the machine. Nick Collins describes the algorithmic nature of musical form

quite extensively, and discusses how it "reveals many areas of enrichment for algorithmic composition" (104). These specific models of musical structures relate so well to the computational process because they "support parsing and chunking of information" which " promotes their mechanizability for the algorithmic composer" (Collins 104). This concept of "chunking" and parsing information is natural to how humans and computers both process, and this creates another big similarity between the two. The similarities between the human brain and the computer allow for a transition of human creativity to an algorithmic process. These similarities make room for even more crossover between the way humans and computers both think, and this is still true for the process of composition.

Humans love to examine the natural world, which is the entire basis of scientific discovery including that of mathematics and computers, and due to their direct relationship, the use of computers and algorithms is an effective method in helping us express this natural world. Numbers have been used to represent natural phenomena for ages, and some of these same things are expressed through music. This creates an intimate relationship between musical composition and mathematical representation. Creating music using these mathematical representations comes naturally to both humans and computers. This is an intuitive thing for composers to do because "music can represent a reflection of our environment and our imagination. Therefore, there are many good reasons for creating music from the same numbers we use to understand our world," and it is an effective form of expression because "composing music from numbers is a symbolic expression of the interconnectedness that unifies our world" (Middleton and Dowd 134). These numbers are readily processed by computers. This concept is commonly used among algorithmic composers. The use of natural series of numbers, such as the Fibonacci sequence and Markov-chains, in applications to music has produced several breakthroughs in algorithmic composition. Several programs discussed later, such as the counterpoint program designed by Adiloglu and Alpaslan as well as the Jazz improvisation machine produced by Gillick, Tang and Keller, utilize these same mathematical series to create music. Using these numbers for expression is almost the exact opposite of what would be expected of a composer. However, because it is a "machine" that is composing the music, the creative limits that a human would normally face are not placed upon the machine. Although the machine has its own "creative" limitations, due to simply executing logical instructions, it is more well suited to get the job done, and this is because these "musical systems embody particular music theories,

110

and these theories can extend to whichever...arrangements of musical objects their author desires" (103). The theories that these musical systems are representative of, are the same deterministic aspects of music previously discussed that translated so well to computation. The use of these mechanical characteristics of music theory, such as the concept of a key, scale, musical form, tonalism, serialism or harmony, are all well-defined structures for a computer to process. These are all aspects of music that can be processed algorithmically, and there exist many more situations in which computers easily adapt to musical composition. The problem lies in the characteristics of music that are not so readily transferable to implementation in an algorithm.

One of the greatest challenges to the field of artificial intelligence in musical composition is the fact that there are also other features of musical creativity that are harder to model for computers. Human creativity can be defined in two different contexts, the spontaneous, "genius" inspirational creativity, and the contrasting process of iterative revisions, which can be described more simply as "hard work." The former approach is nearly impossible to model algorithmically, especially with current technology and scientific understanding, because we are not aware of the psychological processes that cause this inspiration. However, the latter "blood sweat and tears" approach to composition is much easier to model algorithmically, and this has been done numerous times. The way human composers create music is generally through a long, iterative process. This process resembles algorithmic computation, and is able to be recreated and imitated by means of computation because it is closely related to the way computers "think." (Jacob 2) This second form of creativity is much more well suited to computers because of the defined structures that the composer follows. This has been exemplified previously by the counterpoint program, and Jacob's own work is also based off of this fact. Bruce Jacob designed a computer program that simulates a "theme with variation" type of composition. A given theme of music is entered in the program, and it then computes variations of that theme. After "composing" a new variation of the input theme, the computer then analyzes it to see whether it is good or not, and then continues with the "composition." This very much resembles the "hard work" method of creativity. However, both forms of creativity have been modeled in some way by the work of David Cope, who is arguably the most important in terms of composition, and his work is certainly the most well known and controversial.

The work of David Cope, specifically his computer program "Experiments in Musical Intelligence," referred to as EMI, shows characteristics of both types of

creativity described by Bruce Jacob, even though the second "genius" type of creativity is only exemplified as an after thought, as I will discuss. David Cope created EMI in order to assist himself in finishing an opera he was commissioned to do. The program's objective is to analyze a database of musical information, consisting of thousands of pieces of music from many different composers, and to recognize patterns in the music in order to produce compositions in a similar style to those in the database. Cope's intention was to have a database of his own compositions so the computer program could produce a piece of music in his own style without him composing the music himself, thus conquering his "writer's block." However, what resulted was something much more than a Cope opera. He soon found that the program could take compositions from almost any given composer, and produce a new piece of music in the same style. Some pieces were indistinguishable from that of the original composers. EMI has produced new music in the style of many composers ranging from Bach to Scott Joplin. The algorithmic processes of EMI are very representative of the first type of creativity described by Jacob. By analyzing different compositions within a database, and finding patterns in the music in order to create similar pieces, the program demonstrates exactly what Jacob was talking about when he referred to creation through "hard work." This seems to be a pretty intuitive concept though. The real surprise is that Cope's program indirectly exemplifies the second form of creativity as well, although not to the extent that we can say computers would be able to model this creativity. By modeling music that was composed in this way, the fact that EMI produces music in the exact style of Beethoven, or other composers, shows that computers are able to model the "genius" form of creativity. This is simply because the compositions in the database can be described in this way, and a piece that derives from those represents the same aspects of creativity, although indirectly. This may be the most popular and controversial example of this type of modeling of computational creativity, but there is more work in the field that yields somewhat similar results.

Work has also been done to model and replicate Jazz solos in the style of specific musicians, and this has proven effective at representing the types of creativity described by Jacob as well as managing to exemplify the second type of spontaneous creativity better than EMI does, although it still does not fully display that type of creativity. A team of researchers, Jon Gillick, Kevin Tang and Robert Keller, has created a computer program that successfully replicates improvised Jazz solos. This seems to be contradictory to what a Jazz solo even is. Improvisation would seemingly be impossible

to model and implement algorithmically because when a human performer improvises, it is "on the spot." The player does not think so much about the music in his or her head, due to the spontaneous nature of the performance. However, "although a given jazz performer might not be aware of how he or she does improvise, it seems reasonable to say that ideas of what one is able and willing to play can be captured in the form of patterns or, more generally, some form of grammar" (Gillick et al 57). These patterns are discovered in a similar way to how EMI discovers musical patterns among compositions. Specific Jazz solos are transcribed, input into a database, and then analyzed by the computer program in order to produce similar improvisations. The output then represents an improvised solo in the same style of a particular Jazz musician, such as Charlie Parker or Miles Davis, in the same way that EMI produces music in the style of other composers. The produced solos can also be altered to fit any backing chord progression, in order to display the versatile nature of Jazz soloing, and how it can be modeled by a computer. Although the effectiveness of the results varied, a large amount of the improvisations produced were considered to be very close in style to the original solos. (Gillick et al 65) Several different groups of people were questioned, and most have said that the solos produced by the program are somewhat close representations of each Jazz musician. Some solos the machine produced were better than others, and were nearly indistinguishable in style to the intended composer, but other produced solos seemed to lack "direction." This can be fixed with further programming because the work is not yet complete, but still remains relevant to the discussion. Particularly, the longer solos were of lesser quality because they lacked structure throughout their entirety, but some were still of high quality. This shows that creativity can most certainly be modeled by computers in a number of different ways. This specific example demonstrates the second type of creativity much more than EMI does, and this may be simply because of the nature of improvised soloing. Most of how the algorithm works displays the "hard work" method of creativity, but the fact that these pieces are improvisations shows that it is possible to model spontaneous creativity in some ways, although not by any means entirely at this point in time.

Even though EMI and the Jazz improvisation program both produce output that is similar to the music produced by past musical "geniuses," this does not mean that it is easy, or even possible at this point in time, to model this type of creativity. Neither of these programs produced music that was "groundbreaking" in and of itself, because the only revolutionary aspect of the music was how it was created, and this has nothing to do

with the final product itself. Although both examples successfully displayed this type of creativity in the final results, both programs did so indirectly. The computer did not simply have a creative outburst in order to produce a musical masterpiece representing its inner emotions; it was following a specific set of instructions in order to produce the music. The algorithmic method of producing the resulting music was indistinguishable from that of creativity through "hard work," but the results showed something deeper because in the end the music was similar to that which is produced by "genius" creativity. This is why both examples show hints at the second form of creativity, but only through the results and not the methods. The first form of creativity-through-hard-work was very well demonstrated in both cases, and has proven to be a successful way to compose music algorithmically. It is true that the genius form of creativity cannot be modeled with our current understanding of science and the human brain, as well as creativity, but this may change with time because "As computers get more powerful, they become capable of handling increasingly complex tasks" (Jacob 4). Currently, computers are not by any means able to model or replicate "genius" creativity, and it is uncertain whether they will ever be able to do this or not. Advances in the field of artificial creativity will be able to shed any light on this subject, although due to the field's relative incompleteness, it does not seem likely that computers will be able to capture the essence of "musical genius" any time soon, and this is where the limitations of algorithmic composition lie.

Along with the view that work in the field of computational creativity is incomplete, there is fierce debate among scholars concerning whether or not algorithmic composition is a valid form of musical creation, and whether or not the music produced by it is real or not. Many scholars become infuriated to hear music created by computers because machines are thought to be too mathematical and mechanical to create elaborate works of art that spawn from the deepest of human emotions. Opposed to the common definition of music, the compositions produced by computers do not stem from human creativity, as other music is said to. Music is even commonly defined as being intentional and deriving from human creation, and this is one of the main things that make music what it is. One semester I took an Intro to Music class in which the professor formally defined music in this way stating that it is not just "organized sound" because in order for music to be music, human intention must be there. This would mean that random sounds generated from the environment would not be considered music, even if they somehow produce an organization of pitches, unless of course those environmental sounds are intended by some human as music, as shown with John Cage's infamous piano piece "4'33."

114

However, even if we assume this definition to be correct, it can still be argued that music produced by algorithms does indeed derive from human creation. This is because the algorithms themselves were created by a human, generally with the intention of producing algorithmic music. This clearly satisfies the presented narrow definition of music.

This type of discussion leads to many other difficult questions concerning who the "real" composer of the music is, and this particular question may stem from David Cope's work with EMI. The creator of the algorithm that produced the music may be considered the real composer, but it can be just as likely that, in the case of Cope's work at least, the original composer who's work is being modeled is the real composer, even though most of those composers have long since been deceased! Even the user of the algorithm has some influence on the composition of the final product through listening, and choosing which part he or she likes best (Jacob 3). Despite this dispute about who the "real" composer of the music is, it generally holds true that the "more closely an algorithm reflects a composer's methodology, the less question there is that the work is authentic and of the composer" (Jacob 3). Even this does not convince everybody, and the debate as to who the composer of a given piece of algorithmically composed music is remains inconclusive at this point in time. Due to certain opinions musicologists have toward computers composing music, the produced work may not even be considered music! However, this seems like a highly unlikely conclusion because it is solely based on the common biases regarding algorithmic composition.

There exists a plethora of stereotypes associated with computational composition, many of which result from a misconception of what algorithmic composition is, as well as assumptions that the reason these flaws exist is due to compositional aspects rather than features of the performance. Some consider the music produced by computers as being considered "emotionless," "cold," or "mechanical." When people hear the music that they are told is produced by a computer program, they are most likely specifically listening for these aspects because that is what they would expect to hear. This view can be attributed to the placebo effect. This also goes with how we think of computers and how we think of music. Computers have no emotion because they are just logical mechanical devices, and when we combine this thought with music, we think of emotionless music. Often, the creator of a computer program that produces output, music in this case, also programs a method to hear that output immediately. This is usually a MIDI playback of the output converted into MIDI sheet music, which tends to be by far

the most effective, and quickest, playback method. It allows the composer to hear the results right away rather than transcribing it to sheet music every time. When computers "compose" music, they do not just magically start playing the music. They produce a completely logical and numerical output that must then be interpreted and converted into something coherent and meaningful. The pitch, amplitude, and duration of a note are all converted to numerical equivalents because that is how computers process or "think," and by converting these numerical equivalents into a MIDI format or something similar to it, the programmer allows his creation to play the music it composes immediately after composing it. This requires additional computer programming, but saves time in the long run. As many people know through experience, almost anything played through a medium such as MIDI will sound mechanical and emotionless, precisely because it is being performed by a computer.

Even the greatest works of Beethoven or Mozart will sound somewhat mechanical, because of the way we will perceive it when we hear it, and this is entirely due to the mechanical performance of the music. When people address this music as emotionless, they are often talking about the performance aspects rather than the compositional aspects, at least in most cases. Computers are not experts at performing music by any means, at least not at this point in time, but that does not mean they are unable to become master composers. When a human musician performs the music composed by a computer, he or she gives the same amount of emotion to a given piece of music as he or she would with any composer. A large portion of the emotion in this aspect comes from the performer, and how the music is being played by that performer. It is true for a variety of music that a great deal of the emotion in a given piece stems from the performance of the piece. A wonderful example of this would be from David Cope's creation, Experiments in Musical Intelligence (EMI). EMI's very first composition was a piece imitating the style of a Bach chorale. When Cope originally heard the output of the program, he thought it sounded completely horrible. This was likely due to the mechanical, emotionless sound that is expressed through a raw MIDI playback. After time had passed, he had a choir perform the music, and it suddenly gained a huge variety of emotions. The song had gone from cold and emotionless to pristine and otherworldly. The very same music with a different performance medium conveyed much more emotion than it had previously because the humans performing the music provided it with these emotional characteristics.

Computers do not have emotions, at least not in the same way that humans experience them, but this does not mean that music produced by computers is incapable of displaying aspects of human emotion. It seems apparent that, although machines do not have emotions, the music that is produced by some of the previously discussed examples is remarkably similar to music produced by humans. This alone implies that some of the things computers do seem to reflect emotions that are not there. Despite the lack of emotions in the compositional process, the end result still seems emotionally driven to human listeners including myself. Extensive research in the study of computers and emotions has shown that "machines already have some mechanisms that implement (in part) the functions implemented by the human emotional system," although this is not the same as the way humans experience emotions (Picard 10). This may be true, but even though a machine may implement aspects of emotion, this does not mean the machine has a "soul" (Picard 11). The fact that these machines have no soul also appears to be a common criticism of algorithmic composition. This is not a valid criticism because as we have seen, not all music derives from human emotions or the soul, and this is true even for music composed by humans. The formulaic characteristics of music seem to refute the thought that all music composition stems from the human soul. Perhaps the fact that a machine does not possess a soul is what makes the "genius" type of creativity so difficult to model. However, even if the lack of a soul represents difficulty in modeling genius creativity, computers are still remarkably capable of composing music through different methods, and this contributes to the ability of this music to remain indistinguishable from human compositions. This speaks to a computer's ability to appear as if it has emotions to a human being even though it does not. They are capable of the display of emotions, but not truly having these emotions.

Because of this, computers are also able to display emotions and appear to users as if they are capable of this, and they can also determine the emotions of humans based off of physiological feedback, which increases their ability to display these false emotions. Certain musical characteristics convey specific emotions to humans, due to cultural influence, and these emotions are displayed through specific properties of a piece of music. This makes the conveyance of musical emotion much easier for a computer to do because a composer does not require the specific emotions in order to produce music with those characteristics. Emotions in music are based off of our cultural perception of the given piece of music, so whatever emotions may be present in it are in fact reducible to specific traits of the music itself. This transitions to the computational process because a

computer would be able to produce these specific traits in the music easily. Through algorithmically composed music, computers are capable of maintaining "a controlled manipulation of the user's emotional state" because "experiments have shown that there are indeed representative physiological patterns for a user's attitude towards music which can be exploited in a music composition system" (Kim and Andre 269-270). Through analyzing the emotional content of humans in relation to characteristics of music, a machine would theoretically be able to produce music that displays specific emotions. This creates the illusion of having a compositional soul similar to that of a human. If algorithms are used effectively in the composition of music, it will be exceptionally difficult for a listener to be able to tell if the music has been generated algorithmically, or traditionally (Supper 52). By utilizing these specific features of music, the computer essentially begins to express emotions through music, even though the machine does not have any real feelings at all!

Despite all of the criticisms, stereotypes and biases against algorithmic composition, it has still proven to be a successful method for composing music. Several examples of programs that have successfully implemented algorithms in order to produce music, and a good deal of the music produced by these programs can be considered indistinguishable from music composed by humans. This truly shows that computers are capable for much more than what most people give them credit for, and that they can be used effectively in the arts, which is not generally thought of about computers. At this point in time, human composers are still more diverse than computers, but the fact that a computer can create thousands of pieces of music in only a few seconds or minutes simply by analyzing data and making calculations is amazing. This opens an entire new world of collaboration between humans and computers in musical composition because the use of algorithms and computers can greatly enhance a human's compositional power. This may be through optimizing a composer's current work, increasing a composer's options or even by helping the composer overcome severe writer's block as in the case of David Cope. This work has many more implications in the sciences as well. Research on algorithmic composition can contribute greatly to the field of artificial intelligence, and may help computer specialists learn how to better model the way humans think creatively. However, despite the many positive outcomes of algorithmic composition, it is destined to be "frowned upon" by traditional composers (Jacob 1). No matter how effectively a computer can model the human creation of music, traditional musicologists will always disagree with it. The fact that a machine can replicate works of art similar to the way

humans do offends a lot of people that spend their lives studying music. "There are questions one must successfully answer before one can use the algorithmic composition in a manner that escapes mudslinging" because algorithmic composition is destined for mudslinging simply based off of what it is (Jacob 13). It will be a very long time before the use of algorithms, and computers in the composition of music, will be taken seriously by music specialists outside of the world of artificial intelligence, but so far these machines have proven to be capable of being masterful composers.

Bibliography

Picard, R.W., 2000. What does it mean for a computer to have emotions?. Chapter to appear in Emotions and Human Artifacts, MIT press, Cambridge.

Adiloglu, Kamil, and Ferda N. Alpaslan. "A Machine Learning Approach to Two-voice Counterpoint Composition." *Knowledge Based Systems* 20 (2007): 300-09.

Alsop, Roger. "Exploring the Self Through Algorithmic Composition." *Leonardo Music Journal* 9 (1999).

Collins, Nick. "Musical Form and Algorithmic Composition." *Contemporary Music Review* 28.1 (2009): 103-114.

Davismoon, Stephen. "For the Future." *Contemporary Music Review* 28.2.

Gillick, Jon, Kevin Tang, and Robert M. Keller. "Machine Learning of Jazz Grammars." *Computer Music Journal* 34.3 (2010): 56-66.

Jacob, Bruce L. "Algorithmic Composition as a Model of Creativity." *Organised Sound* 1.3 (1996).

Kim, Sunjung, and Elisabeth Andre. "A Generate and Sense Approach to Automated Music Compositional." Proc. of 9th International Conference on Intelligent User Interfaces, ACM, New York, NY. 2004.

Middleton, Jonathon N., and Diane Dowd. "Web-Based Algorithmic Composition from Extramusical Resources." *MIT Press* 41.2 (2008): 128-35.

Supper, Martin. "A Few Remarks on Algorithmic Composition." *Computer Music Journal* 25.1 (2001): 48-53.