Algorithmic Music Perception: Differentiating Human from Machine Composition

A Document

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

degree of

Master of Music in Music Technology

Education Track

By

Joseph A. DiVita

University of Valley Forge

Phoenixville, Pennsylvania

May 2017

Algorithmic Music Perception: Differentiating Human from Machine Composition

A document submitted to the graduate school

in partial-fulfillment of the requirements

for the degree

Masters of Music in Music Technology

by

Joseph A. DiVita

APPROVED BY
SUPERVISING COMMITTEE:

____________________________________ ____________

Supervisor: C. Floyd Richmond Date

____________________________________ ____________

Committee Member: William F. DeSanto Date

University of Valley Forge

Phoenixville, Pennsylvania

May 2017

Abstract

In recent years, musical works composed by computers have caught the attention of the public and have sparked debate on whether computer-based composition could advance beyond the robot-like expressions typically associated with machines to a level comparable to humans. This study will examine the historical and recent developments and will test whether people of varying musical backgrounds can successfully distinguish between works composed by humans and those composed by computers, especially as the computers implement a variety of modern compositional algorithms. The results of this quantitative study determined that a group of test subjects was able to correctly identify whether a piece of music was composed by a computer or a human 48% of the time. Furthermore, human compositions were correctly identified at a rate of 57%, while computer-composed compositions at 40%. Instrumental experience and age had little to no effect. Levels of music theory training showed a weak, but positive correlation. More modern and complex algorithms that employed a combination of AI systems proved to be more successful.

Table of Contents

List of Tables

Chapter 1 - Introduction

Background

Although preceded by earlier pioneering efforts, computer-based algorithmic musical composition has been most thoroughly explored and developed since the late 1950s. It is a type of computer creativity and often considered a form of artificial intelligence.

On a whole, according to Hiller (1959), Loy (1989 and 1981) and Winsor (1987), algorithmic composition "does not relate to the manner in which sound is actually produced, edited, or broadcast. It covers the choice of what sounds will go where in a composition" (Cope, 1997, p. 192). Algorithmic composition does not necessarily have to employ the use of a machine, but in this paper, the term will always refer to music composed by a computer or similar device.

Algorithmic music systems fall into two primary categories: indeterminate (stochastic) and determinate methods. Indeterminate methods produce results using random functions where some portion of the output is left to chance. Determinate procedures are those "where results are fixed by the algorithms and remain unchanged no matter how often the algorithms are run" (Edwards, 2011, p. 3). An interesting indeterminate system being explored today is generative music - a term originally used by English musician Brian Eno to refer to music that is created by a system (algorithm) and always changing. This has shown great promise in the field of video game design where the music is written in realtime, controlled by the players' actions. Chapter 2 , page 22 will identify and examine additional subtypes of determinate and indeterminate systems.

Ethical Issues

Fiametta Ghedini, a researcher at a company called "Flow Machines" which develops artificial intelligence systems, has said "people have a bias against all the creative things produced by a computer," noting how one website called the company's work a "dire warning for humanity." (Needham, 2017) Many people believe that computer-based algorithmic musical composition is cheating due to the fact that a machine is making decisions about the musical content rather than the human composer. For this reason, it has often been misunderstood and looked upon negatively or with fear. This is illustrated by a review given in 1989 of a concert of music composed by a machine known as "*Experiments in Musical Intelligence"* in the style of C.P.E. Bach. The local reviewer gave his assessment two weeks prior to the performance arguing that "too few people were aware of the music of the human version of C.P.E. Bach to be confused with digital imitations. The reviewer admitted to having not heard a single example of output from *Experiments in Musical Intelligence*, adding that he did not *ever* want to hear such output" (Cope, 1999, p. 80) Some common reasons given for these views are that the compositions lack "soul" or feeling, are lifeless, and they stand out from compositions that are completely human-composed. If this is true, the listener should be able to tell the difference aurally between computer-generated compositions and those written with other tools such as pencil and paper.

Patricio da Silva affirms that the fear and uneasiness surrounding the acceptance and appreciation of these tools lies on the unwillingness to understand their purpose and design. (Silva, 2003)

Many other people believe the problem with this prejudiced thinking is that it stifles the advancement of the art form. Michael Edwards (2011) has stated:

"Without wishing to imply that instrumental composition is in a general state of stagnation, if the computer is the universal tool, there is surely no doubt then not to apply it to composition would be, if not exactly an example of Ludditism [sic], then at least to risk missing important aesthetic developments that only the computer can stimulate and facilitate, and which other artistic fields are already taking advantage of." (p. 7)

If technology is a tool that humans use to make their work easier, we should be utilizing it like any other tool at our disposal. Ghedini explains the intended use for these machines is to "augment creative possibilities without substituting them." While British researcher Geraint Wiggins believes the "social impact will be that artificially intelligent composing companions will be stimulating us to try new things." (Needham, 2017) David Cope, a world-renowned, respected author and algorithmic composer says, "he cannot see anyone arguing convincingly that the notes gathered on a page of music contain a [principle of life, feeling, thought, and action. (Webster's, p. 1278, 1991)] No, the soul we perceive when we hear a deeply moving musical work, if "soul" is even the right word, is *our own soul*" (Cope, 2001 p. 91).

In theory, a computer following an algorithm is no different than a human doing the same, apart from the fact that the process is completed much quicker. Edwards (2011) writes:

"Much of the resistance to algorithmic composition that persists to this day stems from a basic misunderstanding that the computers compose the music, not the composer. This is, in the vast majority of cases where the composer is also the programmer, simply not true. As Curtis Roads points out, it takes a good composer to design algorithms that will result in music that captures the imagination.” (p.7)

Earliest Computer Efforts

The first machine developed to create music automatically through melodic generation was invented by Harry Olson and Hubert Belar in the 1940s at RCA laboratories. This machine took its input from a statistical analysis of twelve Stephen Foster melodies and produced imitations of his songs. The machine, simply called “The Music Composing Machine,” would output a few bars at a time. Olson and Belar would “play this for audiences and ask them who the composer was and they would immediately say Stephen Foster. But when you ask them which composition of Stephen Foster, they couldn't say” (Heyer, 1975). Then in the 1960s, two major scholarly papers were published that marked a tipping point in the field of artificial intelligence: “Pattern in Music” by Herbert Simon and Richard Sumner (1963) at Carnegie-Mellon University and “Linguistics and the Computer Analysis of Tonal Harmony” by Terry Winograd at MIT (1968). Both publications are recognized as the earliest studies that deal with machine recognition of musical concepts, a critical and integral preliminary process to machine composition. Before any machine can produce musical output, it must first be programmed to process the language of music. Simon and Sumner’s (1963) writing examines form through pattern and sequence, while Winograd’s (1968) compares music with language and systematic grammar. These efforts broke down music into a rudimentary set of rules and patterns that machines could recognize, code and store to be referenced when compiling output.

Studies

In the past, several studies have addressed both the effectiveness of computer algorithmic composition and human interaction with it. Donya Quick's 2014 doctoral dissertation "Kulitta: a Framework for Automated Music Composition," surveyed participants' reactions to compositions from three separate origins, two of which were computer algorithms. Her conclusions varied depending on which algorithm was used, but supported the idea that humans could not tell the difference between computer and human compositions. (Quick, 2014) David Cope's *EMI* (experiments in musical intelligence) has been tested in a study by Professor Douglas Hofstadter of the University of Oregon to see if listeners could discriminate human performances of the machine's work from those by Bach. This study affirmed the conclusions that humans could not tell the difference. There is also a current study in the form of a survey about generative music taking place at the IT University of Copenhagen whose aim is to "validate a new music generation technique." The study focuses on emotional responses to computer-composed music and its role (if any) in aural distinction of the pieces' origin. In 2016, researcher Nick Collins at Durham University published a study of automated composition's effectiveness in musical theater that was evaluated by both theater composers and theater audiences. The machine-composed work had a two-week run in London. (Collins, 2016).

Computer creativity, algorithmic music composition, and generative music are prolific areas of study and experimentation today. Although many corporations such as Google and universities such as Stanford, MIT, and The University of Montreal are doing experimental work in these areas, shortcomings still exist in their endeavors. As the related literature section will show, these studies fail to adequately examine the possible influence the demographics of the listener may play in the process of evaluation of computer-generated compositions. Variables such as age, musical background, and training have often been overlooked. The proposed study will attempt to mitigate this gap in current research. It will investigate through statistical analysis trends in the ability or lack thereof to discern computer-generated from human compositions based on specific demographics.

This study may shed light on the current widespread idea that computer-composed music is of lower quality than music composed without the aid of a computer algorithm. Todd and Loy summarize this viewpoint stating “laws of art are characteristically fuzzy and ill-suited for algorithmic description” (Todd & Loy, 1991, p. 212). It may influence scholarly discussion among others seeking to advance or understand the practice. Results of this study could help design engineers of music software to take new philosophical approaches to product development. It may impact the current intellectual music society to embrace the practice so that the current state of our technology can allow us to reach our potential as creators.

The purpose of this quantitative study is to explore whether individuals have the ability to distinguish musical compositions that have been written by a computer implementing an algorithm from those created by humans without a formal algorithm. The participants will receive twelve, fifteen-second audio samples. The musical selections for the study will be of similar genre (classical, instrumental) and either computer-composed or human-composed. The participant will answer in electronic survey form with either “computer,” “human,” or “don’t know.” The selected piece of music will be the independent variable, and the answer given will be the dependent variable. The study will also explore possible demographic factors that may influence the ability of participants to make correct identifications. These mediating variables will be: the age of the participant, if the participant has any experience playing an instrument, if the participant has experience playing more than one instrument, and to what degree they have studied music theory, if at all. The age variable will be defined by different ranges that include 10-15, 15-20, 20-30, 30-50, and over 50. The research will involve random participants of varying ages and musical backgrounds.

The study will attempt to yield answers to the following questions:

1. Can the participants determine the difference between computer-generated compositions and human-generated compositions?
2. Does the age of the participant have an effect on their ability to determine the difference between computer-generated compositions and human-generated compositions?
3. Does experience playing a musical instrument or multiple instruments have any significant influence on their ability to determine the difference between computer-generated compositions and human-generated compositions?
4. Does the amount of formal training in music theory have any significant influence on their ability to determine the difference between computer-generated compositions and human-generated compositions?

The (null) hypothesis in this study will be that the subjects will not be able to correctly identify whether a piece of music is computer-composed or human-composed at a better than random rate (50%). Therefore, there will be no statistical significance after analysis of responses ($p<0.05$). Furthermore, the mediating variables will have no significant effect on the subject's ability to distinguish the difference.

Chapter II - Literature Review

History

The Merriam Webster dictionary defines algorithm as "a step-by-step procedure for solving a problem or accomplishing some end especially by a computer." ("Algorithm") By this definition, algorithms can be implemented without the use of a machine. The earliest recorded use of algorithmic composition was "*Musikalisches Würfelspiel*;" which is German for "Musical Dice Games." The numbers on the dice were connected with pre-composed passages to be arranged by chance. This was quite popular with western European composers in the late 18th century but generally connected with Franz Josef Haydn and Wolfgang Mozart.

> One such Würfelspiel, attributed to Haydn, consists of two six-by-eight matrices containing the number one to ninety-six. The number eight represents the number of measures in a typical classical phrase, and the number six represents the number of possible outcomes of the throw of one die. These numbers are then keyed to ninety-six measures of music. Throwing a die sixteen times produces two phrases of eight measures each. All measures were created by Haydn with the foreknowledge of how they could connect with any of the six possible succeeding measures. Hence the 6^{16} or over 2.8 trillion possibilities generated from this *Musikalisches Würfelspiel* will all approximate Haydn's style. (Cope, 1997, p. 197)

According to Roads (1980) in the Computer Music Journal "Artificial Intelligence and Music" the first machine for composing music automatically was invented in 1821 by Dutch inventor Dietrich Nikolaus Winkel, who was also the inventor of the first successful metronome. This machine was called the "Componium." It would generate variations on a theme that was mechanically programmed into the device. The next developments would not come until the middle of the 20th century utilizing computers.

According to Manning (1980), the next composing machine was designed by Hubert Belar and Harry Olson in the late 1940's. This machine used statistics to quantify the melodies of twelve Stephen Foster songs. In 1955 at the University of Illinois, composers Lejaren Hiller and Leonard Isaacson used the Illiac computer to compose *The Illiac Suite*. James Harley reports that Hiller initially took an approach of filtering a random number generator through sets of general rules culled from particular musical styles to create order from chaos. The device would slowly progress from non-predictive behavior to relating the composition to previously "heard" material to anticipate outcomes. (Harley, 1995) The following year Hiller founded the Experimental Music Studio at the University to study and develop electroacoustic and computer music. (Roads, 1980)

In the early 1960s, while working for Bell Labs, computer music pioneer Max Mathews developed a programming language called MUSIC that underwent four revisions. This language focused mainly on direct synthesis. In Hubert Howe's (2009) article in the Computer Music Journal "My Experiences with Max Mathews in the Early Days of Computer Music" he details the dissemination through the computer music community and its potential for revolutionizing the emerging field. This work proved to be of the utmost importance at the time. Howe notes that the writing of Mathew's boss, John Pierce, suggests, "they did not regard music as a serious field of study" (Howe, 2009, p 43). This was a common sentiment that musical establishments had regarding computer music at the time. Howe describes how one of Mathews' contemporaries, John

Chowning, was “initially turned down for his tenure at Stanford” (Howe 2009, p.42) even though he invented frequency modulation (FM) synthesis in 1967, prior to his tenure case. Chowning would go on to invent digital reverberation and the simulation of moving sounds by the early 1970s. (Howe, 2009) Considering his commercial success and impact on the recording industry with these inventions, it seems absurd that Stanford music professors had such a disregard for computer music as they did not see it as a valid area of study, much like John Pierce at Bell Labs. There was a consensus among the academic music community that musical experimentation with machines was frivolous or bizarre.

The first modern developments in generative algorithmic composition by machines were made by focusing on the challenges of quantifying music in a way that could be effectively “understood” by a machine. In 1967 Herbert Simon and Richard Sumner explored how humans perceive music through the idea of pattern recognition. The concepts that they explored would later be applied to machine learning. Doris Lora (1979) summarized the findings with “Human beings seem to have a strong inclination to impose patterns on random auditory stimuli (noise). Simon and Sumner found that people hear orderly sequences in random noises and tend to treat elements that do not fit their imposed patterns as ‘exceptional’ ” (para. 11). Using these concepts, a system was developed to program machines to identify and replicate sequences of sounds.

The following year at MIT, Terry Winograd (1968) made great strides in computerized language recognition with the scholarly document, “Linguistics and the Computer Analysis of Tonal Harmony.” Winograd built upon an idea presented by Allan Forte (1962) who says “The similarity of music to language at many levels has long been a subject of discussion, and studies are now beginning which use some of the basic ideas of linguistics in order to better express the structure of music" (Winograd, 1968, p. 4). This study brought the idea of context recognition to automatic music learning and was the genesis of "grammar learning" in artificial intelligence. This first venture was primitive due to the inability to create applications that could "automatically translate text in one language to another due to the fact that all the subtleties of expression in any language cannot be translated literally" (Manning, 1980, p. 121).

As Roads continues his overview, he points to a programming language called EUTERPE developed by Stephen Smoliar at MIT which generated Gregorian chant, medieval polyphony and Bach style counterpoint. Smoliar's work with EUTERPE began in 1967 and continued into the 1970s. Roads also makes mention of a program called "SIM-SIM" which was written in 1973 at the University of California (Berkeley) that simulated jazz improvisation. The advancements of that decade made it possible for machines to engage in some level of musical cognition in the form of pattern and context recognition. This coupled with computers' superior problem-solving capabilities by the early 1980s led to a wide expansion in areas of focus in the field, incorporating such developments as generative sound synthesis, and machines programmed to test humans' cognitive abilities through music. (Roads, 1980)

In 1981 David Cope began work on his landmark computer program known as "EMI", an acronym for "Experiments in Musical Intelligence." Cope describes EMI as "an analysis program that uses its output to compose new examples of music in the style of the music in its database without replicating any of those pieces exactly" (Garcia, 2015, para. 9). This endeavor started as an answer to a case of writer's block. He intended to use what knowledge he had gathered about programming and artificial intelligence up until that point to create a machine that would compose in his own style. At the time, David Cope struggled to identify his own style so he set out to have the machine compose in other styles in a method he describes as “data driven” (Garcia, 2015, para. 11) The machine would compose based on data analysis alone and through this the machine developed its own “style.” Upon observing the output of EMI, he would adjust his own style of input to achieve the results he was looking for. This whole process took eight years to refine.

The 1980s saw an introduction of simple algorithmic composition components in consumer grade electronics, the most prolific being the early keyboards manufactured by companies such as Casio and its contemporaries. The machines would playback pre-determined, variable accompaniment patterns such as arpeggiations based on the input of one or two keys.

In the 1990s, David Cope went on to design the computer program known as “Emily Howell,” a program that can accept verbal and audible feedback from an audience or programmer and adjust the compositions accordingly. He is quoted as saying:

> The program produces something and I say yes or no, and it puts weights on various aspects in order to create that particular version. I’ve taught the program what my musical tastes are, but it's not music in the style of any of the styles—it's Emily's own style. (Cheng, 2009, para. 11)

That same decade saw new introductions in the consumer market as well with programs like “Band-in-a-Box” and separate hardware “arranger” units, to be used to compliment solo performers. These devices would have selectable “styles” consisting of a specific instrumentation playing pre-determined patterns that could be triggered via MIDI and were often built-in to entry-level keyboards. These arrangers have expanded into a very sophisticated line of products with hundreds of preset and customizable paradigms modeled after popular songs or artists.

In the twenty-first century we have seen a massive expansion of the technology associated with algorithmic composition in terms of application, access, and availability. What once took massive computers the size of washing machines to generate can now be accomplished with notebook or tablet machines, and even smart phones. Some of the more powerful programs that have gained traction with consumers are *SuperCollider* and Common Music's *Grace* for Windows, Mac OS and Linux. A stand-alone application and DAW plug-in for both Windows and Mac OS is called *Liquid Notes*. Applications exclusively for windows are *Easy Music Composer* and *Maestro Genesis.* (Moss, 2015, para. 6) Moss adds that the following algorithmic composition programs designed to work in real-time have appeared: "*Impromptu* and various other audio plugins and tools help VJs and other performers to "live code" their performances, while OMax learns a musician's style and effectively improvises an accompaniment, always adapting to what's happening in the sound" (Moss, 2015 para. 14). Educational software such as Music First's "Sight-Reading Factory" have incorporated algorithmic composition for generating unique excerpts based on user specified parameters for students to practice. The program can write infinite scores for ensembles as well as single note lines of varying lengths and play them back. The user can dictate the length, rhythmic resolution, meter, pitch range, instrumentation and dynamic markings (if any) the machine should use. This technology may be questioned by some for creating art, but has proved its value in the classroom.

Types and Processes

Martin Supper (2001) believes it is important to make a distinction between algorithmic score synthesis and computer sound synthesis when referring to computer music. Score synthesis refers to the process of a machine determining the design of the elements in a musical work. Sound synthesis is the generation of simple sounds through methods such as subtractive, granular, additive, or FM synthesis. The following processes explored in the remainder of this section will deal with Score synthesis exclusively.

There are many different types of systems that create compositions using a variety of technical approaches. George Papadopoulos and Geraint Wiggins (1999), British researchers, categorized these systems in the following manner for the AISB Symposium on Musical Creativity: mathematical models, knowledge based systems, grammar systems, evolutionary methods, systems which learn*,* and hybrid systems.

"Mathematical models" are made up of stochastic processes - machines that output results based on random combinations of variables and their probabilities. One popular example of the concept is a Markov chain. The Markov chain is a stochastic process that has additional imposed rules or restrictions. Its main advantage is that a lack of complexity facilitates good real-time application. Random number generators would also fall within this category. Disadvantages of stochastic processes are the need for a human to analyze many pieces for probabilities and any deviations from the norm are often overlooked resulting in a lower quality of expression. (Papadopoulos & Wiggins, 1999)

“Knowledge based systems” use sets of rules or constraints. They have the ability to “reason” and can explain their choice of actions in a sense. Laske (1989) states the methodology is “an orderly sequence of steps leading from the elicitation to the analysis of knowledge, onto its modeling in some implementation-independent form, and, via system design, to implementation in the form of a knowledge base- a very long life cycle” (p.7). They have numerous downsides. One is that they cannot reason based on instinct or intuition because those things are not easily quantifiable into rules or facts. The compiling of the knowledge base is time consuming, and tedious, especially when applied to music. They lack flexibility. Their development and use require an “expert,” usually in addition to the programmer, to advise and clarify concepts. Their results are often “rigid” as the exceptions to the rules cannot be programmed or they become too complicated. (Papadopoulos & Wiggins, 1999)

“Grammar systems:” These systems, most notably David Cope’s EMI (which happens to fall under multiple system types), extracts rules and commonalities called “signatures” from compositions to emulate certain styles. They are good for emulation, but have significant drawbacks. Their ability to generate a large quantity of work is usually offset by the fact that a significant portion is of questionable aesthetic quality. The act of analyzing works must be done by a human and is time consuming. (Papadopoulos & Wiggins, 1999)

“Evolutionary methods” are processes that are modeled after biological evolution employing concepts such as natural selection, reproduction, mutation, and recombination. The most common evolutionary method is known as the “genetic algorithm”. They are efficient for dealing with large search spaces. A search space is the entire possible range of solutions or combinations that can be output by the algorithm. However, they are simplistic in their output, due to their method of making choices based on paths of least resistance, which can result in decreased quality. (Papadopoulos & Wiggins, 1999)

"Systems which learn" do so by examples. They are also divided into two categories: artificial neural networks and machine learning. These systems have advantages over knowledge-based systems in that they can learn exceptions to rules and loosely resemble human brain functions. They have many disadvantages. The whole process is highly symbolic and these machines generally fail at capturing musical features such as phrasing or tonal functions. David Cope's EMI also falls within this category. (Papadopoulos & Wiggins, 1999)

Hybrid systems, such as generative music systems, use a combination of artificial intelligence techniques. The advantages are that they combine the positive attributes and can minimize particular shortcomings of other systems on their own. The disadvantage of hybrid systems is that they are usually complicated. The design implementation and testing is often time consuming. (Papadopoulos & Wiggins, 1999)

Examples of systems and their functions

This section will describe three examples of algorithmic composing systems, each based on a different combination of the technical tools described above.

Dirk Povel (2010) designed an expert system called "Melody Generator" to use in a study to test "the adequacy of a melody generating device based on specific theoretical ideas about the construction of tonal music" (p. 684). In the study, Povel described how this knowledge-based system works to create melodies from scratch. The project was based on two assumptions: that the conception of music is a psychological phenomenon, and that tonal music is conceived within the context of time and pitch. The first assumption states that the perception and organization of music is produced purely in the psyche. That is to say that outside of the human mind, specific collections of sound have no meaning. The second assumption focuses on the more technical musical elements. Time is configured by meter (imposing constraints as to when events occur). Pitch is configured by key and harmony (imposing constraints as to what frequencies may occur). Within that context, tone sequences are generated using a set of construction rules specifying the organization of single pitches in a hierarchy designed from the top-down (larger parts to smaller parts).

Rhythm, meter, subdivision, and accentuation are deconstructed on a perceptive level and grouped according to metrical stability.

> The term metrical stability denotes how strongly a rhythm evokes a meter. We have seen that the degree of stability is a function of how well the pattern of accents in a rhythm matches the pattern of weights of a meter. (Povel, 2010, p. 685)

It is stated that this correlation can be quantified by a correlation coefficient, and when it falls below a critical level, the meter would be lost. This becomes the framework on where notes will begin in time.

Pitch hierarchy is defined in three levels: key, harmony, and tones (melody). A system of weights and constraints is also discussed for the levels of key and harmony as well as the relationship between meter and key. Melody is discussed in greater depth and broken down to a framework or skeleton, which can be based on a chord model or scale model. (Povel, 2010)

The application is fed a series of parameters for the time, key, and harmony dimensions, which it creates first. Afterward, melody is then constructed within a flexible framework of parts and subparts, divided into bars, divided again into beats and perhaps even further subdivisions, depending on the parameters dictated by the user. These are then treated as "slots" where pitches can be placed based on the hierarchy of key and harmony. The entire process is done in four stages: Construction, editing, (re)arrangement, and transformation which are controlled by the user. The transformative stage allows the user to have the machine alter certain parameters of the generated melodies to create variations. (Povel, 2010)

Sano & Jenkins (1989) proposed a neural network designed upon the workings of the human ear to replicate pitch perception. Jenkins breaks down the biological ear into parts, focusing particularly on the middle and inner ear. The middle ear acts as an impedance matching device, transmitting the pressure variations to the cochlea where they grow in amplitude, peak and collapse. High frequencies collapse at one end, and low frequencies on the other. The frequency detection is performed by hair in the inner ear that is part of the cochlea. This hair acts like linear band-pass filters. There are two theories of how pitch discrimination works. The place model, which states that the frequency perceived is directly related to which hairs are stimulated and the periodicity model which states that neural signals sent are proportional to the frequencies gathered. The network proposed would work on the place model using single, complex tones. The network will identify pitch and octave information but no timbre. (Sano & Jenkins, 1989)

Sano and Jenkins (1989) describes in intricate detail the mathematics of the stages of modeling the nerve output of the cochlea and reducing it to semitone "buckets." The process produces 463 "Just Noticeable Difference" (JND) buckets for a frequency range of two octaves. This number is then reduced to 24 (semitone) buckets. It was deduced that there is a considerable amount of imprecision in the stimulation of the hair. A single sine wave signal with a frequency of 1047 Hz (C6) caused hairs cells to respond over the range of bandwidth from 998-1102 Hz. Because of this, the network uses both non-overlapping and over-lapping fan-in processes during the main stage of "pre-processing."

The next main stage is the synthetic mode unification process, where the buckets are tied together with their upper harmonics. After another stage containing a process labeled normalization and further reduction into pitch classes, the tonal data is fed into a back propagation neural network for discriminatory identification (perception).

Lukasz Mazurowski (2015) designed a mathematical system that created what was described as "background music" through the use of mini-models. The system he created was called GEBMS (Generative Electronic Background Music System). The overall algorithm can process music in an analytic, transformational, or generative manner. The analytic functions "tend to reduce the potential data size and the general music predisposition of the representation by extracting specific features" (p. 1). Transformational algorithms alter the information and the generative algorithms output a sequence based what was deduced from the input. A set of secondary functions determines the context, which is described as "the surrounding information that influences the computation of an algorithm and therefore and algorithmic music system" (p. 1).

The mini-model contains a subset of data made up of class (monophonic, chord, or drum), note matrix, an Onset Positions Vector that describes where notes can be generated in the matrix, Combination Motifs Matrix (CMM) which is defined by the user as rules of possible pitch sequences. From this information, patterns are generated and compared to one another for similarity and are eventually selected for inclusion in the final composition. Following this process a transformational rule is applied that is defined by the user as well as a probability table which leads to the final output.

Attitudes, Criticism, and Testing

Algorithmic composition has the tendency to polarize people into two camps: those who view it negatively, and those who view it like any other tool at a composer's disposal. Those who view it negatively cite reasons that it is cheating, produces inferior results, or lacks emotional depth. Many "traditional composers" frown upon the use of a computer to aid in compositional decision making because they think it would be unlike music that the same composer would have created without the machine's influence. This argument is certainly invalid if the composer has designed and implemented their own algorithm modeled on their own process of composing. Muscutt and Cope (2007) speculate that the potential for high output leads listeners to have less interest in computer aided algorithmic compositions.

> Bruce Jacob (1996) defines "two distinctive types of creativity: - the flash out of the blue (inspiration? genius?) and the process of incremental revisions (hard work). Not only are we years away from modeling the former, we do not even

begin to understand it. The latter is algorithmic in nature and has been modeled in many systems both musical and non-musical." (p. 1)

By this reasoning, algorithmic composition with the aid of a computer is utilizing a tool to accomplish the work faster and no less valid than algorithmic composition with the aid of a pencil and paper. Jacob raises the question about the difference between music composition and music recognition.

> When an algorithm written by a composer produces music that is not exactly what the composer would produce, the composer filters it – he culls out parts that conflict with his own personal taste. Is this composition? How minimal must the changes be in order for the algorithm-produced music to be considered composed (by the algorithm or the composer)? How extensive must the changes be before ones role as algorithm – designer moves from composer to editor? At what point is the composer merely *recognizing* music that he likes? Is there a difference? (Jacob, 1996, p. 2)

The next logical, hypothetical question is: what difference does it make if it is done by a machine, or done by a human? David Cope makes the point that "computers are not a requisite for algorithmic composition" (Muscutt & Cope, 2007, p 12).

Jacob concludes that the success of using a machine to allow a composer to work more quickly depends on having a close match to the composer's own process and having a way to sort the desirable results. He has developed a system called "variations" that is designed to closely reproduce his own creative process and advises that composers must find their own path to quantifying their processes. "The problem of computer – aided music composition is best reduced from answering the nebulous *how do I create music* to answering slightly more concrete questions: *what steps do I use to compose music* and *what types of harmonic movement do I like*" (Jacob, 1996, p. 12). Through this exploration and matching of one's own processes ensures the computer is a tool and not a crutch.

P. Landa (1994) wrote about his experiences at the 1992 International CompuMusic Conference in San Jose California in respect to a panel discussion regarding algorithmic composition. Panelist Max Matthews played an example of an early experiment of his own involving permutations of a phrase of a popular song and although not meant to be taken so seriously, it became a reference point for many audience members to make arguments that algorithmic composition is worthless and futile. Landa recalled how in the 1970s, two journalists set out to abolish the idea of modern art by giving a chimpanzee some oil paint, brushes and a canvas. They submitted the resulting painting to an annual art exhibition in Stockholm. The painting was actually admitted by the jury and displayed next to other works by credible modern artists.

Landa surmised that the people in the ICMC session did not understand or pay attention to what impact directional procedures have on stochastic (random) processes the same way the two Swedish journalists did when they overlooked the consequence of directing the chimp and stopping them when the painting looked like a piece of artwork. In reality, it was the journalists who were the artists, and the chimpanzee was just a tool or intermediary. Composers throughout history have had the tendency to utilize whatever tools from the highest technologies are available.

Donya Quick (2014) discusses obstacles to assessing the output of automated composition machines. There is a problem of judging what "good" music is because of its highly subjective nature. For some styles of music, such as Bach chorales, it is possible to make these judgments by analyzing the score, but for newer styles of music, no rules exist beyond how humans respond to the pieces. Adding complexity to the matter is the quality of performance or delivery of the composition. Automated music systems unfortunately often trade quality for novelty. Those that produce more “creative” music also produce a great amount of garbage, and those that produce higher-quality results more consistently tend to sound more or less the same.

David Cope (1999) has offered some answers to these issues as he makes a number of points about how we react to computer-generated music. His first generation machine, EMI, is regularly criticized for having too many close quotations to the music it has been directed to base its composition on. He argues that all human composers allude to other musical works, including their own. Cope urges that there must be a redefining of terms and criteria. Many listeners fail to view the compositions as new, original music, trying to connect the compositions with their other experiences of what they consider computer music. There seems to be a disconnect about how to categorize the results of his computer-generated compositions as evidenced in a 2015 interview with computerhistory.org.

> I spent almost a year trying to get an actual record company to produce the music. It was really tough. Cope said, “I remember my greatest exasperation was, coming in on the same day, were two negative replies. The first said ‘we only publish contemporary music, and this, by our definitions, is not contemporary music, and then the other one said ‘we only do classic music, and this is not classical music’, so I said ‘then, what is it?’”(Garcia, 2015, para. 19)

One common criticism of EMI's output is that is lacks any trace of humanity. Cope refutes these arguments by stating that humans composed the music that is in its database, the algorithm and code was written by a human and humans listen to and evaluate the output. His response to critiques about his machine lacking "soul" is that the "soul" experienced by listeners to any music is that of their own; it is a function of the listener. When asked if he ever composes without using algorithms, his response was that he does not believe one can compose without using at least some.

Quick (2014) designed a machine called *Kulitta* for algorithmic composition and tested the "humanness" of its output in an experimental study. *Kulitta* addresses music composition as a process of problem solving. Quick identified two problems: the vertical, which deals with the exponential nature of all possible combinations of pitches being sounded from an 88-key piano, and the horizontal, a similar concept dealing with rhythm. The way her machine, *Kulitta* deals with these problems is through the process of "musical abstraction." This process breaks large tasks down to smaller ones and allows the solutions to gradually emerge into finer levels of detail.

Her study consisted of a survey that presented the participants with 40 short pieces of music and the participant would rate their level of confidence that it was written by either a human or a machine. To avoid the issue of style bias, the participants were first prepped by being shown examples of music composed by both humans and computers in contrasting styles of atonal and classical chorales. The phrases the participants voted on (stimuli) came from three sources. The first source was Quick's program, a random-walk generator and three chorales written by J.S. Bach. Two versions of *Kulitta* compositions were used, one set using hand-built grammars, and one set trained from Bach Chorales. The experiment was run online using Amazon Mechanical Turk as a source of participants. Data from a total of 237 participants was obtained that had various levels of education, musical training, and a mix of composers and non-composers. They were to rate each piece on a scale of 0-6, whereas 0 was absolutely human and 6 was absolutely computer. Demographics were also collected on age and gender, nationality, ethnicity, first language, and highest level of education and musical background. (Quick, 2014)

The results of the experiment found that all comparisons were statistically significant ($p<0.01$). Bach's scores averaged closest to zero (absolutely human), the random generator averaged closest to six, *Kulitta's* scores fell in the middle with the average on the human side. Music training and theory training showed no correlation with the average scores given to each composer. Although Quick's study did collect information on musical background and demographics such as age, it was only applied to her particular machine.

Chapter III – Design and Implementation

The purpose of this study is to determine if humans can accurately distinguish computer-aided algorithmic musical compositions from those that are entirely human-composed. The target population (*n*) was from North America, age 16 and older, and had various forms of musical training and experience. The 250 participants were randomly selected volunteers that took the survey during a three-week span. They were pulled from a variety of sources including in person requests by the researcher and requests made on various online communities and forums such as Reddit, Facebook groups, and The Academic Forum.

A survey was used to gather data because it is an effective and efficient collection procedure due to its low cost, rapid turnaround time, its ability to gather information on multiple variables at once, and its ability to provide a high level of generalizability. The survey was given online using embedded audio in Google forms. In order to minimize bias due to underrepresentation of certain age groups or those without internet access, the researcher contacted some participants via telephone or face-to-face to administer the survey in person. Validity and reliability of the testing instrument has not been established through large-scale normalization. Every effort has been made to ensure the simplicity of the questions so that the validity and reliability are self-evident.

The participants were directed to the survey through a hyperlink where they were given instructions on completing the survey. In question one they indicated their age from the ranges offered: 16-20, 21-35, 36-50, and 51 or older. In questions two and three, the participants responded if they currently play an instrument, if they play multiple instruments, with a yes or no answer. In question four, the participants responded if they have had any music theory training from never, basic, high school level, or college and beyond.

These participants were then given twelve 15 second excerpts of music, six composed by a human and six composed with the aid of a machine implemented algorithm (questions five through sixteen). The participants answered each question with whether they believed that the piece was composed by a human, a computer, or that they did not know. The participants played each example. Each example was able to be played multiple times, but they had to enter an answer by clicking on the button corresponding with their response before moving on to the next example.

The copyrighted music being presented is being used under the fair use provisions of the US copyright law, which states:

> § 107 . Limitations on exclusive rights: Fair use
> Notwithstanding the provisions of *sections 106* and *106A*, the fair use of a copyrighted work, including such use by reproduction in copies or phonorecords or by any other means specified by that section, for purposes such as criticism, comment, news reporting, teaching (including multiple copies for classroom use), scholarship, or research, is not an infringement of copyright. In determining whether the use made of a work in any particular case is a fair use the factors to be considered shall include— (1) the purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes; (2) the nature of the copyrighted work; (3) the amount and substantiality of the portion used in relation to the copyrighted work as a whole; and (4) the effect of the use upon the potential market for or value of the copyrighted work. The fact that a work is unpublished shall not itself bar a finding of fair use if such finding is made upon consideration of all the above factors. (U.S. Copyright Law 2011)

The human-composed music contained excerpts from Steven Reich's "Music for 18 Musicians," J.S. Bach's "An Wasserflüssen Babylon, BWV 267," Franz Josef Haydn's "String Quartet in C, op. 74, no. 1," the Prelude from Camille Saint Saëns' "6

Études, Op. 52," Philip Glass' "Violin Concerto No. 1," and Fréderic Chopin's "Etude Op. 25 No. 12 (Ocean)."

The computer algorithmic music contained excerpts from David Cope's "Emily Howell Fugue," "Bach-style Chorale" and "Emmy Vivaldi," Donya Quick's "Kulitta Piano Etude #4," and "Chorale on Strings," and the first experiment from Lejaren Hiller and Leonard Isaacson's "Illiac Suite."

The twelve audio examples and questions were presented in the following order:

1. J.S. Bach's "An Wasserflüssen Babylon, BWV 267
2. Prelude from Camille Saint Saëns' "6 Études, Op. 52
3. David Cope's "Emmy Vivaldi"
4. Philip Glass' "Violin Concerto No. 1"
5. Donya Quick's "Kulitta Piano Etude #4"
6. Steven Reich's "Music for 18 Musicians"
7. Franz Josef Haydn's "String Quartet in C, op. 74, no. 1"
8. Lejaren Hiller and Leonard Isaacson's "Illiac Suite"
9. Donya Quick's "Chorale on Strings"
10. David Cope's "Bach-Style Chorale"
11. Fréderic Chopin's "Etude Op. 25 No. 12 (Ocean)"
12. David Cope's "Emily Howell Fugue"

The following table illustrates the variables in the study:

Variable Type	Description	Item on Survey
Independent Variable	Music selection presented	Questions 5-16 (computer/human)
Dependent Variable	Participants' answer to music questions	Questions 5-16 (computer/human/I don't know)
Mediating Variable #1	Age	Question 1 (10-20/21-35/36-50/50+)
Mediating Variable #2	Does participant play an instrument	Question 2 (yes/no)
Mediating Variable #3	Does participant play more than one instrument	Question 3 (yes/no)
Mediating Variable #4	Has participant studied music theory	Question 4 (Never/Basic/High school level/College or beyond)

Table 1: Variables of the Study

After the data was collected, the correct and incorrect responses for questions five through sixteen were tallied as well as the number of "I don't know" responses. The participants were categorized according to mediating variables. Tables two through four illustrate the survey responses.

A descriptive analysis of the means, standard deviations and ranges of the answers were presented. The results of independent T-tests on mediating variables 1 through 4 against the dependent variable were reported. The results of a correlation coefficient test on mediating variables 1 through 4 against the dependent variable were reported and $p<.05$.

Chapter IV - Results

The purpose of this study was to determine if humans can aurally distinguish music composed by a machine from music composed purely by humans. Additionally, it examined if there was any significant correlation between this ability, age, and musical background. The (null) hypothesis in this study is that there will be no significant difference in the ability of the participants to identify computer vs. human-composed pieces of music based on the demographics of age, musical instrument experience, or the study of music theory.

Out of a sample of 265 responses (n=265), the following tables illustrate their age, instrumental experience, and theory level.

Age Range	*n*	% of *n*	Average Score
16-20	19	7.2%	45.6%
21-35	62	23.3%	50.7%
36-50	131	49.4%	49.6%
51+	53	20%	43.6%

Table 2: Participant age

Instrumental experience	*n*	% of *n*	Average Score
None	97	36.6%	42.8%
one instrument	44	16.6%	51.5%
more than one instrument	124	46.7%	52.1%

Table 3: Instrumental experience

Theory level	*n*	% of *n*	Average Score
None	50	18.8%	42%
Basic	70	26.4%	45.4%
High school	55	20.7%	47.9%
College or beyond	90	33.9%	54.4%

Table 4: Theory level

The mean score for the entire sample was 48.3% with a standard deviation of 2.0, indicating the majority of participants fell within a small range very close to the center of the curve. This low average indicates that the participants were not able to tell the difference between computer-generated and human compositions. The following bar graph represents the total number of correct responses for each listening question (questions 5-16).

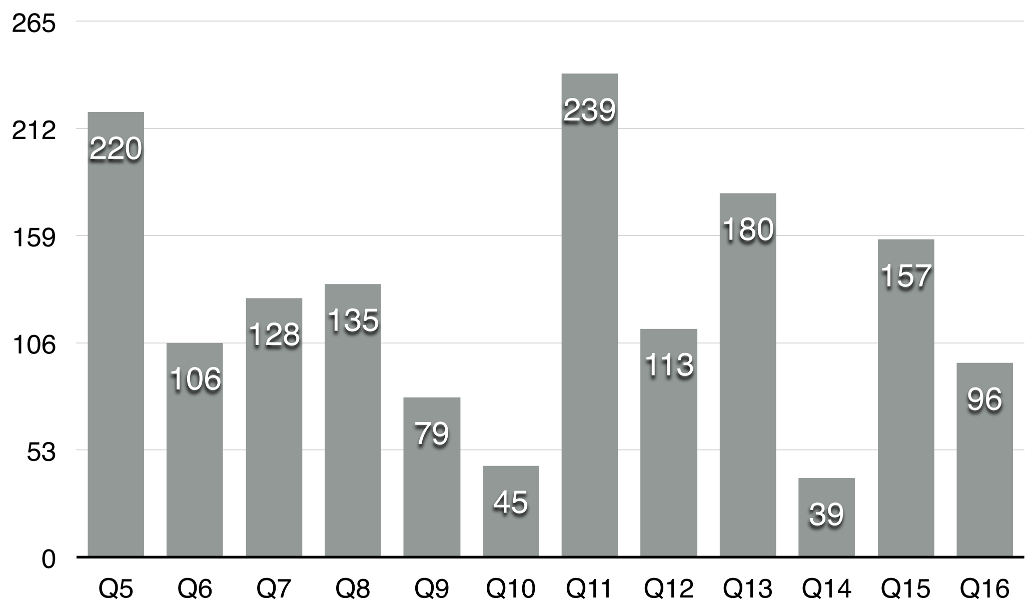

Table 5: Total correct responses

Description of Listening Questions:

- Q5: Human - *J.S. Bach's "An Wasserflüssen Babylon, BWV 267*
- Q6: Human - *Prelude from Camille Saint Saëns' "6 Études, Op. 52*
- Q7: Computer - *David Cope's "Emmy Vivaldi"*
- *Q8:* Human - *Philip Glass' "Violin Concerto No. 1"*
- Q9: Computer - *Donya Quick's "Kulitta Piano Etude #4"*
- Q10: Human - *Steven Reich's "Music for 18 Musicians"*
- Q11: Human - *Franz Josef Haydn's "String Quartet in C, op. 74, no. 1"*
- Q12: Computer - *Lejaren Hiller and Leonard Isaacson's "Illiac Suite"*
- Q13: Computer - *Donya Quick's "Chorale on Strings"*
- Q14: Computer - *David Cope's "Bach-Style Chorale"*
- Q15: Human - *Fréderic Chopin's "Etude Op. 25 No. 12 (Ocean)"*
- Q16: Computer - *David Cope's "Emily Howell Fugue"*

Table 6 shows that participants were able to correctly identify the six human compositions 57% of the time and the six computer-composed compositions 40% of the time. The fact that when computer-generated examples were presented to the subjects, 60% of their responses showed they could not definitively identify them would indicate that music composed by computers is quite successful at replicating humans. Additionally, one may speculate that this demonstrates some quality of pure human composed pieces that makes them more identifiable. This is in direct contrast to some criticisms explored earlier that it is the computer-composed music that stands out.

	Human	Computer
Total correct answers	902	635
Total incorrect answers	688	955
Percentage	57%	40%

Table 6: Total correct answers by composition type

Overall, the age group of 21-35 achieved the highest average score by 1.1 percentage point over the age group of 36-50 at 50.7% and 49.6%, respectively. The age group with the lowest average was 51+. All of the age groups' averages fell within 7.1 percentage points of one another, indicating that age has little to no effect on the ability to separate human from computer-composed pieces. The average between single and multi-instrumentalists were 51.5% and 52.1% respectively, indicating it has a negligible effect on the ability as well.

The average score for those who have never studied music theory was 42%, with basic theory at 45.4%, high school level at 47.9%, and college level and beyond at 54.4%. This does suggest a small positive correlation between music study and the ability to hear the difference between compositional origins. When subjected to a Pearson correlation coefficient test, the resulting value was 0.28, indicating the positive correlation does exist, though very weak.

The following table illustrates the results of one-way ANOVA testing for the age groups, instrumental experience, and theory levels.

Grouping	p-value
Age	0.068028
Instrumental experience	0.000016
Theory level	0.000034

Table 7: Demographic ANOVA test results

The results of these tests clearly show that that any difference in the average scores of the instrumental experience group and theory level group is not significant at $p<.05$. The null hypothesis must therefore be accepted. However, the results of the ANOVA test between age groups shows that the difference in average scores is significant at $p<.05$. This does not have an effect on the main hypothesis, but does show that the oldest participants had the most similar judgments, while the age group 26-50 had the most diverse. This may be due to it being the largest grouping by age in the sample. Table 6 shows the standard deviation for each age range.

Age Group	Standard Deviation
16-20	1.8
21-35	1.9
36-50	2.1
51+	1.6

Table 8: Standard Deviation by age group

Of the six computer-composed pieces in the study, the piece that was the most convincing by far was from the *EMI* system by David Cope, with only 15% of participants correctly identifying it as computer-generated. This Bach-style choral was performed by human voices and may suggest that people are influenced by the method of sound production. The piece that performed most poorly was the strings piece by the *Kullita* system. This piece was played back by synthesized strings, which to many listeners is obviously identified as such. The following table illustrates the algorithmic pieces used in the study, their system, overall score and method of performance in order most convincing to least.

Selection	System	Score	Method of Performance
Bach-Style Chorale	EMI	15%	Human
Piano Etude #4	Kulitta	34%	Sequenced and Synthesized
Emily Howell Fugue	Emily Howell	37%	Human
Illiac Suite	ILLIAC I	47%	Human
Emmy Vivaldi	EMI	49%	Sequenced and Synthesized
Chorale on Strings	Kulitta	68%	Sequenced and Synthesized

Table 9: Algorithmic example details and scores

All of the human-composed pieces were performed by humans with the exception of question 15 (Fréderic Chopin's "Etude Op. 25 No. 12") which was played by a sequencer using a synthesized piano. This did not seem to influence the listeners' opinions as much as the synthesized strings in Donya Quick's piece . It is interesting to note that the piano sample is more realistic, with more reverb, and there is a significant difference in the texture of the piece which masks some subtleties of the instrument's timbre. It is safe to say that many of the participants did not hear it as a synthesized or sampled instrument. This leads to the obvious, interesting question of how timbre may influence a listener's perception of a piece's origin. Would the same piece be assessed similarly if the selection was played with a primitive synthesized patch?

It is interesting that the four pieces that were correctly identified by more than half the participants were questions five, eleven, thirteen and fifteen. Three of these were human-composed and one computer. The human pieces were by Bach, Haydn and the Chopin piece mentioned earlier performed by machine. This suggests that either many participants knew the pieces or there are certain elements from them that have entered our universal sub-conscious similar to a Jungian archetype. The identification of the computer string chorale is most likely due to the synthesized strings discussed previously.

It would appear that the newer, more complex systems tend to be slightly more convincing. David Cope's *EMI* and second generation *Emily Howell* are both hybrids of multiple types of AI systems, minimizing limitations inherent to systems of one type.

Limitations of the Study

While completing this research some limitations in the design of the study became evident. The word “instruments” did not clarify treatment and inclusion of vocalists when asking for the number of instruments played.

This study does not explore whether the instrumentalists are presently regularly practicing musicians or how long ago the responder may have taken a music theory class. It may also be worth investigating whether gender has any impact on the responses or if experience in composition would have an effect.

Suggestions for further Study

This study may lead to another to investigate the bias (if any) listeners have when pieces are performed by humans versus machines in their evaluation. That study would use similar material, but the design would be focused on the method and quality of sound generation. Another study could investigate possible reasons for a listener choosing computer or human by asking questions about specific musical elements. This would be a study that would obviously favor participants with musical experience and background, as those with none would find it hard to understand certain questions or communicate their opinions. It is interesting to note that many of the participants made remarks that they “second-guessed” themselves regularly in this survey, which leads to the thought that most had developed a set of criteria for evaluation.

Final Remarks

Very few things limit the potential infinite variety of musical expression. Those being, the time it takes to create, edit and consume the work. Computers can drastically reduce the time required to create new works, and someday evaluate them. For now, programmers will still have to sort through the machine's output to cull what they find useful. The future may hold great promise in the field of computer-aided composition. Despite some people's fears, these machines will never render human creativity obsolete. However, composers of some musical forms may find less demand for what they do. The ability to create something greater with this tool than one could without it is far more rewarding in a progressive society. Eventually, perceived differences in human and computer compositions will diminish entirely. The two will be so intertwined that we will no longer separate one from the other. Furthermore, there will be no need.

Bibliography

Algorithm. (n.d.). Retrieved November 22, 2016, from http://www.merriam-webster.com/dictionary/algorithm

Cheng, Jacqui. (2009). Virtual composer makes beautiful music- and stirs controversy. Retrieved from http://arstechnica.com/science/2009/09/virtual-composer-makes-beautiful-musicand-stirs-controversy/

Collins, Nick. (2016). A funny thing happened on the way to the formula : algorithmic composition for musical theatre. *Computer music journal., 40(3)*. pp. 41-57.

Cope, D. (1999). Facing the Music: Perspectives on Machine-Composed Music. Leonardo Music Journal, 9, 79-87. Retrieved from http://www.jstor.org.dbs.valleyforge.edu/stable/1513481

Cope, D. (1997). Techniques of the contemporary composer (pp. 192-205). Belmont, CA: Schirmer Books.

Da Silva, Patricio. (2003). David Cope and Experiments in Musical Intelligence. Salt Lake City, UT: Spectrum Press.

Edwards, M. (2011). Algorithmic Composition: Computational Thinking in Music Retrieved from https://www.academia.edu/654039/Algorithmic_Composition_Computational_Thinking_in_Music

Garcia, Chris. (2015). Algorithmic Music- David Cope and EMI. Retrieved from http://www.computerhistory.org/atchm/algorithmic-music-david-cope-and-emi/

Harley, J. (1995). Generative Processes in Algorithmic Composition: Chaos and Music.

Leonardo, 28(3), 221-224. doi:1. Retrieved from http://www.jstor.org.dbs.valleyforge.edu/stable/1576078 doi:1

Heyer, M. (n.d.). HARRY F. OLSON: An Interview Conducted by Mark Heyer, IEEE History Center, 14 July 1975. Retrieved from http://ethw.org/Oral-History:Harry_F._Olson

Howe, H. (2009). My Experiences with Max Mathews in the Early Days of Computer Music. *Computer Music Journal, 33*(3), 41-44. Retrieved from http://www.jstor.org.dbs.valleyforge.edu/stable/40301045

Jacob, Bruce L. (1996) Algorithimic Composition as a Model of Creativity. http://drum.lib.umd.edu/bitstream/handle/1903/7435/OrganizedSound.pdf?sequence=1&isAllowed=y

Landa, P. (1994). Chimpanzees, Wind Chimes, and Algorithmic Composition. *Computer Music Journal, 18*(1), 9-10. doi:1. Retrieved from http://www.jstor.org.dbs.valleyforge.edu/stable/3680514 doi:1

Laske, Otto. (1989). The State of the Art of Computer Music. *Computer Music Journal, 13*(2), 7-8. doi:1. Retrieved from http://www.jstor.org.dbs.valleyforge.edu/stable/3680036 doi:1

Lora, Doris. (1979). Musical Pattern Perception. Retrieved from http://symposium.music.org/index.php?option=com_k2&view=item&id=1857:musical-pattern-perception&Itemid=124#x14

Manning, P. (1980). Computers and Music Composition. *Proceedings of the Royal Musical Association, 107*, 119-131. Retrieved from http://www.jstor.org.dbs.valleyforge.edu/stable/766120

Mazurowski, Lukasz (2015) Generative Electronic Background Music System, Proceedings of the International Conference on Numerical Analysis and Applied Mathematics, Rhodes, Greece, 2014. Melville, NY: AIP Publishing

Moss, Richard. (2015). Creative AI: Computer composers are changing how music is made. Retrieved from http://newatlas.com/creative-artificial-intelligence-computer-algorithmic-music/35764/

Muscutt, K. & Cope, D. (2007). Composing with Algorithms: An Interview with David Cope. *Computer Music Journal, 31*(3), 10-22. Retrieved from http://www.jstor.org.dbs.valleyforge.edu/stable/40072590

Needham, Jack. (2017). We are the Robots: Is the future of music artificial?, Retrieved from http://www.factmag.com/2017/02/19/we-are-the-robots-could-the-future-of-music-be-artificial/

Papadopoulos, G., & Wiggins, G. (1999, April). AI Methods for Algorithmic Composition: A Survey, a Critical View and Future Prospe. AISB Symposium on Musical Creativity, 110-117. Retrieved from: http://www.doc.gold.ac.uk/~mas02gw/papers/AISB99b.pdf

Povel, D. "Melody Generator: A Device for Algorithmic Music Construction," *Journal of Software Engineering and Applications*, Vol. 3 No. 7, 2010, pp. 683-695. doi: 10.4236/jsea.2010.37078.

Quick, Donya. (2014). Kulitta: A Framework for Automated Music Composition (Doctoral dissertation). Retrieved from http://www.donyaquick.com/files/dissertation.pdf

Roads, C. (1980). Artificial Intelligence and Music. *Computer Music Journal, 4*(2), 13-

25. doi:1. Retrieved from http://www.jstor.org.dbs.valleyforge.edu/stable/3680079 doi:1

Sano, H., & Jenkins, B. (1989). A Neural Network Model for Pitch Perception. *Computer Music Journal, 13*(3), 41-48. doi:1. Retrieved from http://www.jstor.org.dbs.valleyforge.edu/stable/3680010 doi:1

Supper, M. (2001). A Few Remarks on Algorithmic Composition. *Computer Music Journal, 25*(1), 48-53. Retrieved from http://www.jstor.org.dbs.valleyforge.edu/stable/3681634

Todd, P. M., & Loy, D. G. (1991). *Music and connectionism*. Cambridge, MA: MIT Press.

U.S. Copyright Law, 17 U.S.C. § 107 (2011)

Winograd, T. (1968). Linguistics and the Computer Analysis of Tonal Harmony. *Journal of Music Theory, 12*(1), 2-49. Retrieved from http://www.jstor.org/stable/842885