Jill Riddiford

The Musician and the Algorithm: The Role of Artificial Intelligence in Music Composition and the Creative Process

"Supposing... the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent."

Ada Lovelace 1842

For a Research Paper completed in Partial Fulfillment of the
Requirements of a Bachelor of Music Degree at the
Academy of Music and Performing Arts
2018

Abstract

As technology increasingly pervades our lives, the words, 'Artificial Intelligence' (AI) and 'algorithm', are rapidly spreading from computer science into the mainstream lexicon. Creative fields including music are not immune to the relentless march of automation. But for those who create and compose music, and for listeners of music, is there a point at which the use of computer programs becomes an anathema – will indeed music lose its soul? Many musicians and non-musicians alike may be surprised to learn that the use of algorithms in music composition has a long history. However, as computer processing power has increased and the programs have become more sophisticated, the interplay between technology and music has grown closer. The aim of this paper is to explore published research into algorithmic composition, focusing mainly on the last 30 years, and to review some publicly available online programs where various parameters can be manipulated to produce music compositions. As music appreciation is largely a subjective process, the issue of how to evaluate AI compositions is also relevant and will be examined briefly in this study. The research concludes that to be a useful tool in the music composer's kit, and to produce 'listenable' music, there needs to be a balance between algorithmic and human involvement in the process.

Acknowledgements

I would like to thank my supervisor, Dr Elizabeth Jones, for her guidance throughout this project and valuable suggestions on the research development. I would also like to thank the individuals who participated in the music listening experiment for their time, useful input and feedback.

Declaration

This work has not previously been submitted for a degree or diploma in any university and to the best of my knowledge and belief contains no materials previously published or written by another person except where due reference is made.

Signed Jill Riddiford *Date 27* April 2018

Table of Contents

Glossary

Aleatoric music	Music composition technique where some elements rely on chance, for example, Mozart (Troedson n.d.), John Cage (Jensen 2009).
Algorithm	A process or set of rules which specifies how to do something in a finite number of steps.
Artificial Intelligence (AI)	The simulation by computer systems of human intelligence including learning, reasoning and self-correction.
Artificial Neural Network (ANN)	A type of AI that attempts to mimic the way a human brain works by creating connections between elements or nodes. Information is fed in and then channeled straight through nodes into an output. The organisation and weights of the connections determine the output.
Cellular Automata (CA)	A process where cells in a grid change properties (specified by set of rules) and based on the properties of their neighboring cells, see *Figure 1*.

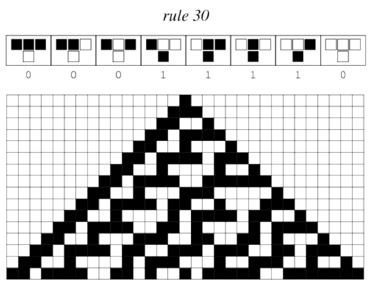

Fig 1. First 15 steps of a CA evolution.

The simplest type of cellular automaton is a binary, nearest-neighbor, one-dimensional automaton. There are 256 such automata, each of which can be indexed by a unique binary number whose decimal representation is known as the "rule" for the particular automaton. (Weisstein 2018)

Evolutionary developmental algorithms (evo-devo)	Unlike classical evolutionary algorithms where 'genes' are mapped directly to solutions, evo-devo recognises the significance of the development process and uses indirect encodings which can produce complex and novel outcomes.
Genetic Algorithm (GA)	Computation process that mimics the evolution process. Its key components are reproduction, mutation, crossover and selection.

Long Short-Term Memory Networks (LSTMs)	A special kind of RNN, capable of learning from longer term signals by way of a memory unit inside the network. The LSTM makes a decision by considering the current input, previous input and previous memory.
Machine Learning	A computer using algorithms to learn how to perform a task, for example, with an ANN. A subset of AI.
Markov Chain	Process where transitions occur based on probabilistic rules but where the possible future states are fixed, see *Figure 2*.

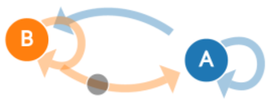

Fig 2. A simple two-state Markov Chain.

With two states (A and B), there are 4 possible transitions (not 2, because a state can transition back into itself) (Powell n.d.).

Recurrent Neural Network (RNN)	An ANN with short term 'memory' or feedback loops. Both current inputs and inputs from the recent past are considered.
Rule-based system	A set of 'if-then' statements which specify steps to be taken by a process. Also known as knowledge-based system.
Stochastic music	Music composition technique that uses probability calculations, for example, Xanakis (Brown 2005).
Turing Test	Developed by Alan Turing in 1950, a test of a computer's ability to exhibit intelligent behavior indistinguishable from that of a human.

A New World of Potential or Losing Music's Soul?

The main aims of this research project are first, to present a brief review of how algorithmic programs have been applied to the music composition process over history, and second, to evaluate some current Artificial Intelligence (AI) programs and their value to general users, alongside amateur and professional musicians.

Just as the World Wide Web evolved from a specific military application to global popular use, research and discussion of AI has spread from academic scientists to the business community, creative fields and increasingly to the general population.

The defeat of the Russian grandmaster chess champion by IBM's *Deep Blue* computer in 1997 (Broekhuyse 2017) marked a milestone in AI news reaching the mainstream media. However, it is only in the last few years that AI research has exploded, as shown in the Google Trends graph for the search term, 'Artificial intelligence' in *Figure 3*.

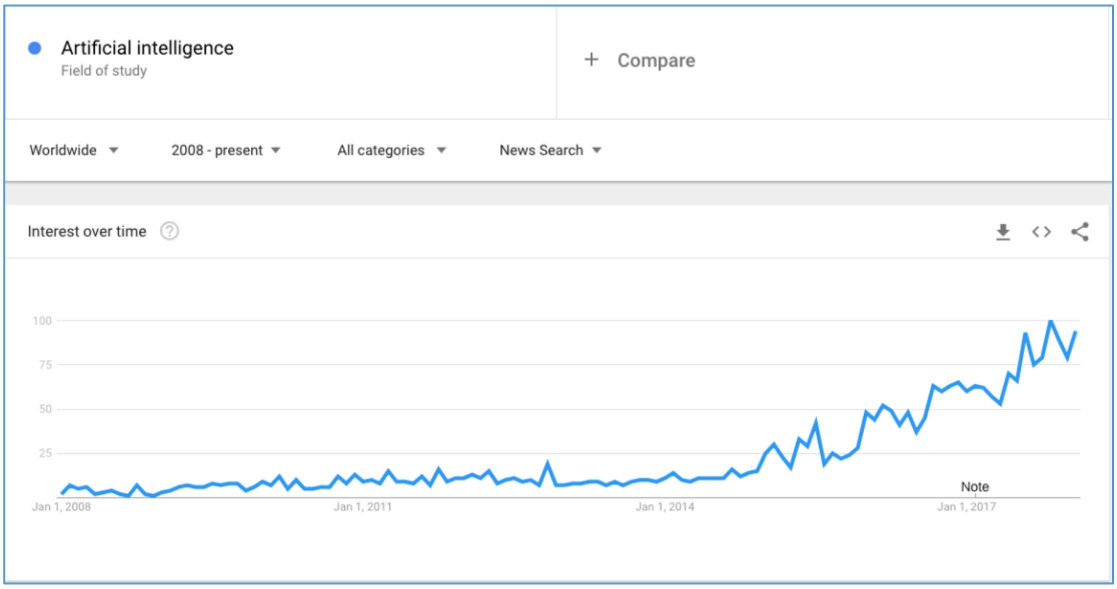

Figure 3: Google Trends statistics for 'AI' search term

As interest in AI spreads, stories in the popular media about applications of AI appear almost daily, for example see Savage (2017), 'Now computers are writing perfectly acceptable pop songs'.

While AI terminology is becoming more commonplace, many people still experience a cognitive gap because of the ongoing permeation of technology into nearly every aspect of life. This is likely to be particularly pronounced in areas which have traditionally been thought of as predominantly ‘creative’ such as music composing – “Is nothing sacred!”

How do we evaluate the merit of algorithmic music? To assist in this primary question, a literature review is presented, including a brief history of music algorithms and focusing mainly on journal papers published since the early nineties.

Any discussion of music algorithms soon becomes inseparable from questions around the creative process, aesthetics, and the acceptance, or otherwise, of technology’s influence. Therefore, these aspects in the research will also be reviewed.

As there are many different types of algorithms that can be applied to music composition, we also need to consider whether some are more valuable than others in helping to create the optimum conditions for creative output.

There is also the practical consideration of using online programs. We see that the composition tools that are freely available online vary in terms of complexity and usefulness. Experiments are carried out using examples from some these programs as well as existing algorithmic works. The resulting compositions are then reviewed by a small group of professional and amateur musicians.

Combining Musicianship with Computer Science

Long before the digital computer age, musicians have used various formal instructions and processes to make compositions. One of the earliest examples was Mozart's Dice Game (first published in 1793), which involved, "assembling a number of small musical fragments, and combining them by chance, piecing together a new piece from randomly chosen parts". (Alpern 1995)

The Dice Game is an example of aleatoric music, or music relying on chance (Brown 2005). Brown's colleague Troedson (n.d.) reproduced the Game's process in a computer music program. jMusic. He notes that the strict harmonic structure of the elements, and some 'cheating' by Mozart whereby some bars had limited options rather than being chosen by chance, produces a high degree of similarity in the resulting compositions. (Troedson n.d.)

John Cage was an avant-garde composer in the 20th century who also used chance techniques to compose music including the Chinese oracular book, the *I Ching*, or *Book of Changes*. The title of Cage's major work, *Music of Changes* (1951), was a reference to this. Jensen (2009) notes that Cage, "sought a balance between the rational and the irrational by allowing random events to function within the context of a controlled system":

> "The amount and type of composerly control varied from piece to piece, but these fundamental ideas of control and randomness remain constant throughout Cage's indeterminate and chance derived output." (Jensen 2009)

As the digital computer age took hold from the late 1950s and with the exponential increase in computing power since then, the types of algorithmic composition techniques have expanded significantly. The potential for more complexity and creativity in digital music programs was perhaps first heralded in the 1830s when Ada Lovelace, a mathematician and colleague of inventor Charles Babbage, (credited with conceiving the precursor of the modern digital computer, the Analytical Engine), foresaw possibilities of creativity in automated music composition that went much further than numerical manipulations. (Loughran 2016)

> "Supposing, for instance, that the fundamental relations of pitched sound in the signs of harmony and of musical composition were susceptible of such expression and adaptations, the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent." (Alpern 1995)

The first instance of digital computer-aided composition is widely recognised as *The Illiac Suite for String Quartet* by Hiller and Isaacson in 1956 which was generated using rule-based systems and Markov chains (Fernández 2013). The composition process used a generator/modifier/selector approach. First a technique was used to generate raw materials as a base composition, then various techniques were applied to further manipulate the material, and finally selection rules were applied to choose suitable material.

Iannis Xenakis was another key pioneer in computer music who used stochastic techniques to control almost every aspect of a musical work including, "the overall duration and structure…individual note attributes and sound generation…" (Brown 2005)

Edwards (2011) notes that Xenakis had no compunction adapting the output of his algorithms as he saw fit, whereas Hiller believed that if the output of the algorithm was deemed insufficient, then the program should be modified and the output regenerated.

A literature review shows that the main algorithmic methods generally fall into five categories:

1. **Rule-based or knowledge-based**
2. **Stochastic – using probability or randomness**
3. **Neural Networks**
4. **Genetic, or evolutionary, algorithms**
5. **Cellular Automata**

The application of different types of algorithms introduces varying degrees of uncertainty and potentially, creativity into the composition process. But whatever technique or combination of techniques chosen, ultimately it is the composer who makes the creative decisions which will result in a completed piece of music.

As Brown (2005) notes, an important consideration when composing music is achieving an appropriate balance between predictability and surprise. When writing music with a computer, the composer needs to, “decide the nature of the uncertainties and their boundaries.”

Edwards (2011) also comments, “Much of the resistance to algorithmic composition that persists to this day stems from the misguided bias that the computer, not the composer, composes the music.”

However, are some programs more suitable than others for generating music?

Fox and Khan (2013) compared three types of algorithms to compose music: stochastic generation via Markov chains, routine planning (rule-based), and genetic algorithms. They found that:

> “…the stochastic algorithm is at a disadvantage because it does not apply any explicit strategy to either follow music theory or compositional strategies that make a song listenable. The main detractor of the songs generated by the genetic algorithm is the overly random nature of the notes. The planning approach does not suffer from either of these problems but may lack in originality because it is impacted the least by randomness.” (Khan 2013)

Brown (2004) compared rule-based and genetic algorithms for generating melodies and found that a particular combination of techniques can produce melodies that are well formed and, “in many cases display some elegance and novelty”.

Quintana (2013) applied an evolutionary developmental (evo-devo) approach with a computer classical music composer *Iamus*, to produce music which according to Berger was, “qualified by professional musicians and composers as indistinguishable from those written by avant-garde composers.” (Quintana, 2013)

The same technology was also applied to composing popular music with the computer system *Melomics109*.

Algorithms based on recurrent neural networks (RNNs) is another approach that has been explored by researchers. Todd (1989) was one of the first to explore this 'connectionism' technique whereby a network learns aspects of musical structure, is given examples of music with those structural aspects, and then uses what it has learnt to compose new pieces of music.

Later, Eck (2002) updated the RNN approach to use long short term memory (LSTM) cells and applied this architecture to improvise blues based on a short recording. Since 2016, Eck's LSTM approaches have been applied to drum pattern generation, melody generation and polyphonic music generation on Google Brain's *Magenta* project. (McDonald 2017)

The final category of algorithm identified earlier is Cellular Automata. With their ability to produce complex patterns and 'lifelike' behaviour, Cellular automata models can be used to create artificial life forms within a computer program.

Miranda (2003) has explored how these models can display emergent behavior, that is, produce novel music compositions. In discussing Miranda's work, Johnson (2002) notes the contrast of this "evolutionary musicology" with the main achievement of other AI models to date, which has been, "restricted to mimicking the style of existing composers, either with a set of AI rules, or by learning a composer's style with a neural network."

Computer scientist, Stephen Wolfram (2011), who has studied cellular automata extensively, thought that the history of using programs to produce music had mostly resulted in music that was "either too robotic or too random" and that his discoveries seemed to offer new possibilities because they showed that, "…even with the rules of a simple program, it was possible to produce the kind of richness and complexity that, … we see and admire in nature." (Wolfram 2011)

His team’s research led to the creation of an online algorithmic music generator, *WolframTones*.

Increasing experimentation by researchers and the development of new types of algorithms is producing a greater potential to produce ‘real’ sounding music. But is there a missing component – can a computer program emulate the performance of a human composer or musician, that is, incorporate the dimension of emotion via expressivity?

Whether it is a concert pianist playing a Beethoven sonata, a professional vocalist interpreting a jazz standard or an electronic music composer manipulating the loudness and timing of MIDI notes, expressivity is an integral component of the music experience.

Coutinho (2005) notes that musicians have, “intuitively linked expressivity with irregularity within certain boundaries”, and quotes Polish composer Paderewski who referring to a score’s expressive marks, stated, “emotion excludes regularity” and “Chopin played from his heart. His playing was not rational, it was emotional.” (Coutinho 2005)

Recognising the importance of expressiveness, researchers are incorporating this dynamic into music composition algorithms. For example, Simon and Oore (2017) at the *Magenta* project have presented an, “LSTM-based recurrent neural network designed to model polyphonic music with expressive timing and dynamics”.

Whether it’s imitating existing patterns or changing variables of a random program, after reviewing the work of some notable researchers and comparing the processes used, it is apparent that human and computer inputs in algorithmic music composition are inextricably linked.

The next section looks briefly at the challenge of evaluating computer music and then reviews some of the composition programs available.

Evaluating and Experimenting with AI Music

Will an AI robot prefer music composed by humans or computers? Anon

As algorithmic music becomes more prevalent, one of the key question arising is – how do we evaluate these compositions?

Loughran and O'Neill (2016) proposed that evaluation systems relying on a "Turing-style discrimination test… may be selling the computational systems short":

> "With the ever increasing power of computational machines, why should we limit these new intelligent systems to a human level of creativity we barely understand ourselves?" (Loughran and O'Neill 2016)

Miranda and Williams (2003) note that while the Turing Test has been adapted to music simply by asking listeners whether they think a piece of music been has been composed by a human or by a machine, there are other considerations:

> "… the use of AI in music composition tasks raises several further aesthetic and philosophical questions. How do we determine what is good or bad when evaluating the output of such systems? And indeed, who is the author? Aesthetic issues are far from universal and are not readily evaluated in a systematic or repeatable way. How do we determine authorship in the case of creating new music with such systems?" (Miranda and Williams 2003)

The subject of applying AI to music creation also exposes somewhat of a culture clash between science and music. Fernández (2013) notes that AI and the arts, "speak different languages and have different methods and goals, creating great difficulties in the collaboration and exchange of ideas between them".

Bearing out this observation, a web search shows that a number of composition tools require some knowledge of computer science and coding/programming expertise, are unavailable to the public or are only accessible at research institutions. Other programs are available to experiment with online, refer to Table 1, or can be downloaded and integrated into a digital audio workstation such as *Orb Composer*.

A sample of online programs is outlined in *Table 1* and a brief discussion of the key features of each follows.

AI program	Type of algorithm
Amper Music https://www.ampermusic.com/	Possibly neural networks
Flow Machines http://www.flow-machines.com/	Machine learning Markov chain (Jordan 2017)
Jukedeck https://www.jukedeck.com	Neural networks (Mathieson 2018)
Magenta (Google Brain) https://magenta.tensorflow.org/	Machine learning – RNN-LSTM
Musical Algorithms http://musicalgorithms.ewu.edu/index.html	Various including Markov chains
WolframTones http://tones.wolfram.com	Cellular automata

Table 1: Examples of composition programs and algorithm type

Amper Music is an AI music startup founded in 2014. Its composition interface offers two modes, Simple and Pro. Users first select between styles – hip hop, cinematic, classic rock, modern folk, 90s pop, then select a corresponding mood, for example, hip hop suggests cool, chillout or reflective. The difference with the Pro option is that the composer can select further variables including instruments, tempo and key.

Flow Machines appears to be closed to public users. On its website the company says its AI tool is, “an intelligent assistant able to help you composing songs in any style, automatically or interactively”.

Jukedeck is an earlier startup which has some basic online composition tools available. In a similar fashion to Amper Music, it provides a choice of 13 genres, then depending on the genre selected, the user can select a choice of mood and instrument type, see *Figure 4*.

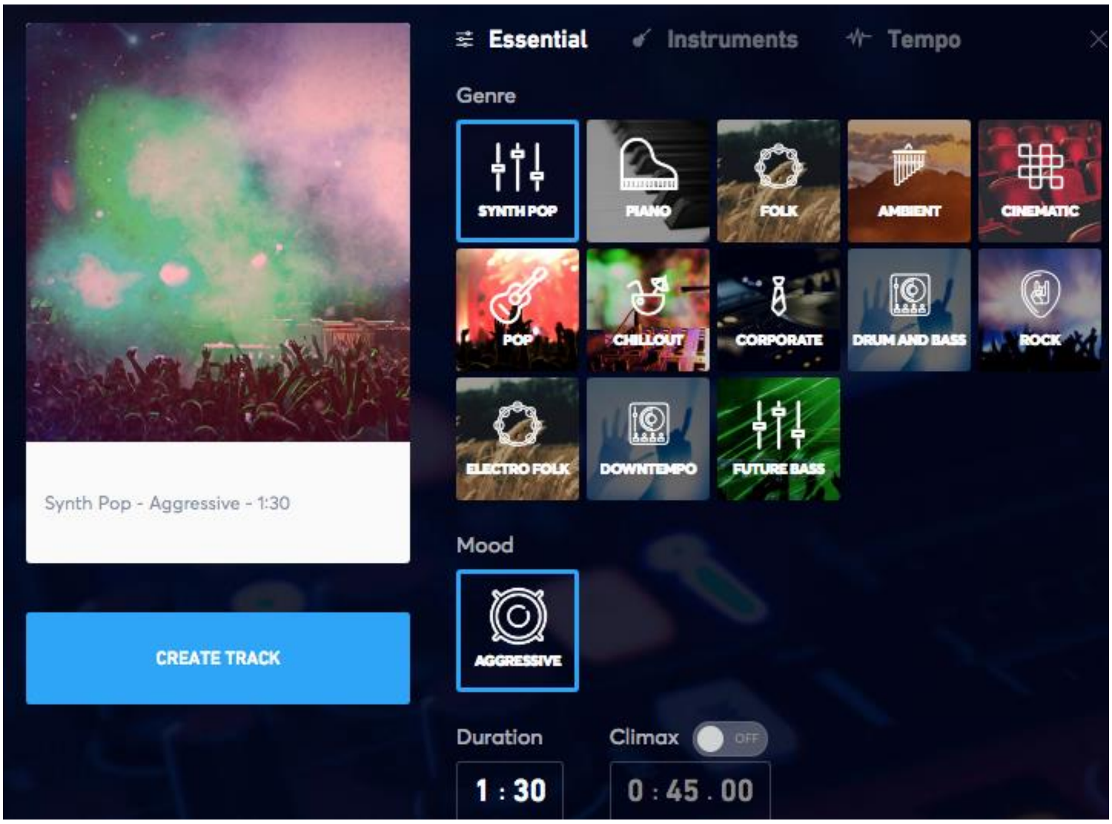

Figure 4: Example of composition template in Jukedeck

Magenta is a research project aimed at developers, "…exploring the role of machine learning in the process of creating art and music". *Magenta* was started by researchers at Google but uses an open-source "machine learning framework" and development platform. (Magenta 2018)

Unlike the high-level, formulaic user interface of programs like Amper Music and Jukedeck, *Musicalgorithms* provides a more granular approach. The online program has 8 algorithms to choose from, including Markov chains. However, users then need to input a range of values which require a more indepth knowledge of how the algorithms work, see *Figures 5 and 6*.

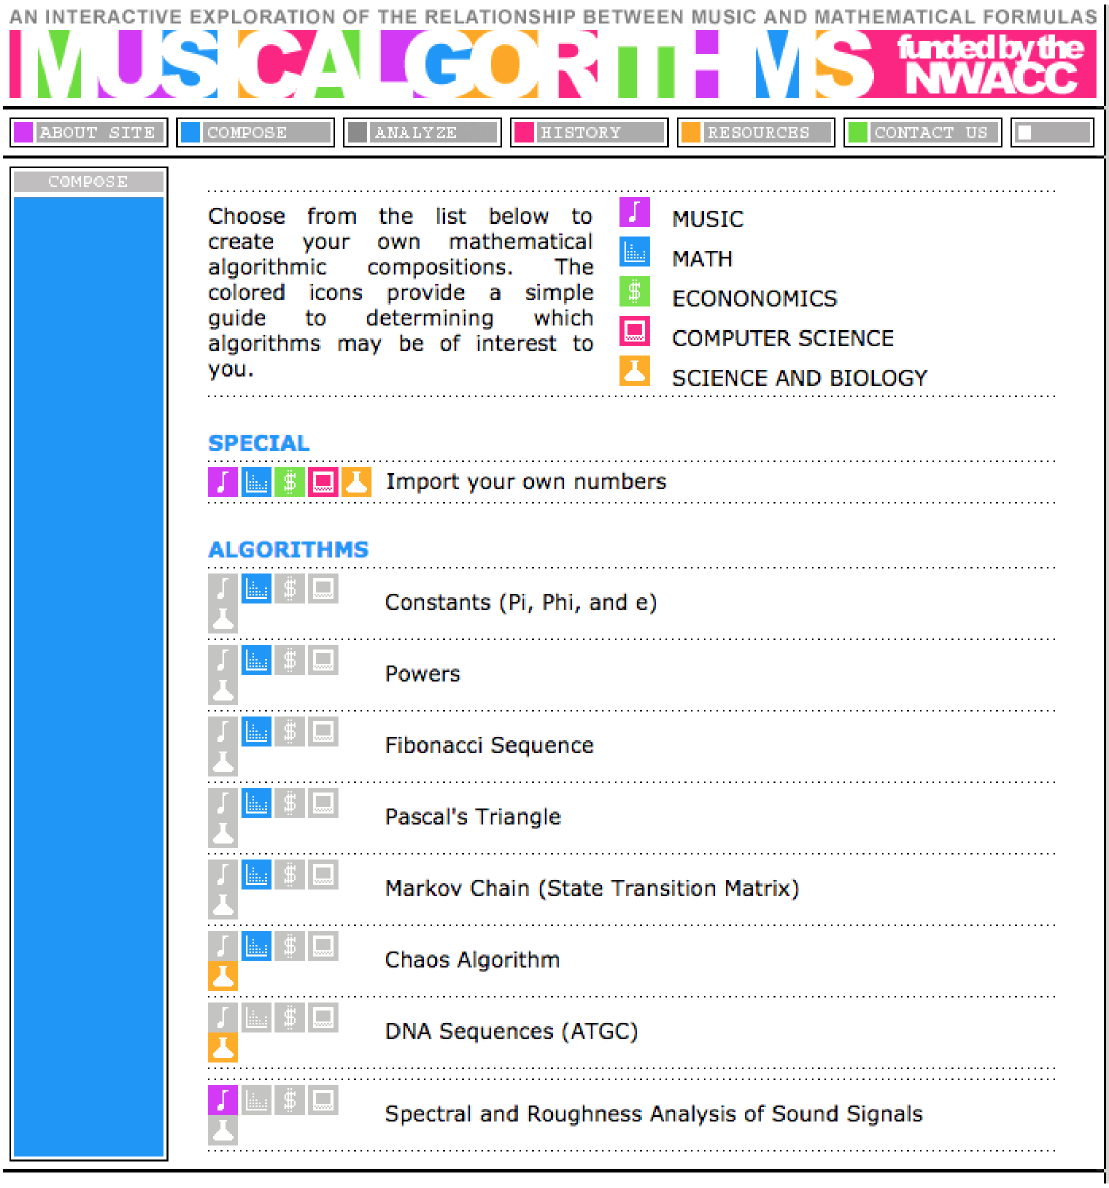

Figure 5: Choice of composition starting points in Musicalgorithms

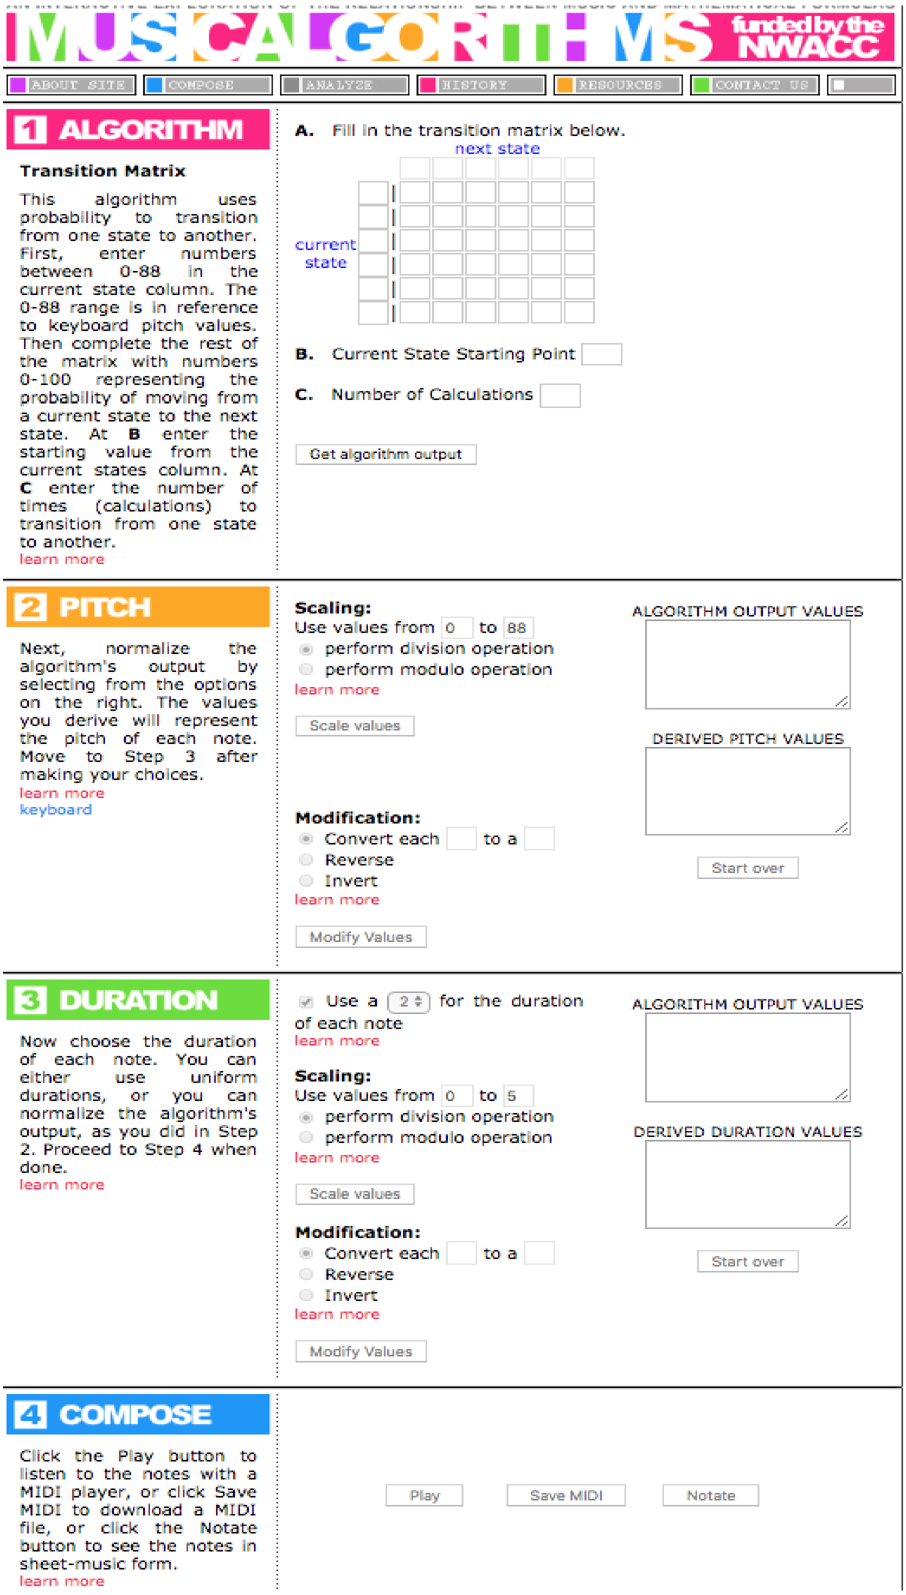

Figure 6: Example of parameter setting choices in Musicalgorithms

WolframTones presents the user with 15 music styles to start a new work. A choice of 100 instruments and synths can be used in 10 instrument slots and a vast number of scales and modes can also be selected. The algorithm itself is varied by selecting from a list of 'Rule types'. However, users don't need to have detailed knowledge of the underlying algorithm, and can create an endless variety of compositions through trial and error. Also, the Play window provides a good representation, see *Figure 7*, of how the composition evolves and if it will be useful or not.

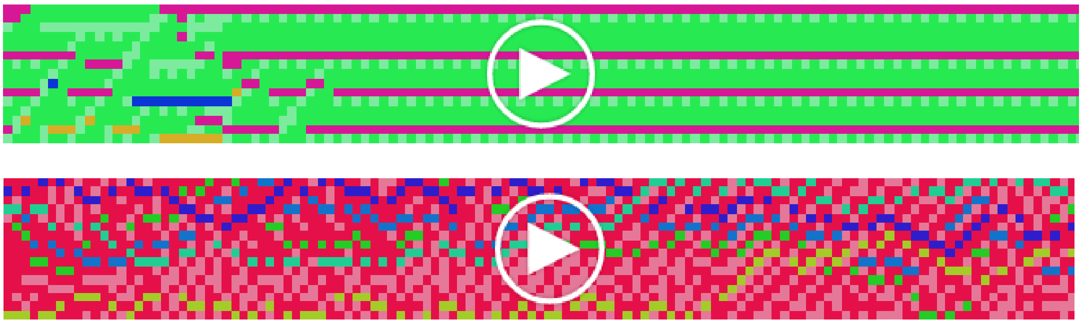

Figure 7: In the top, green example the melody is going nowhere while the bottom, red example has more variety and shows promise

A full example of the *WolframTones* composition template is shown in *Figure 8*.

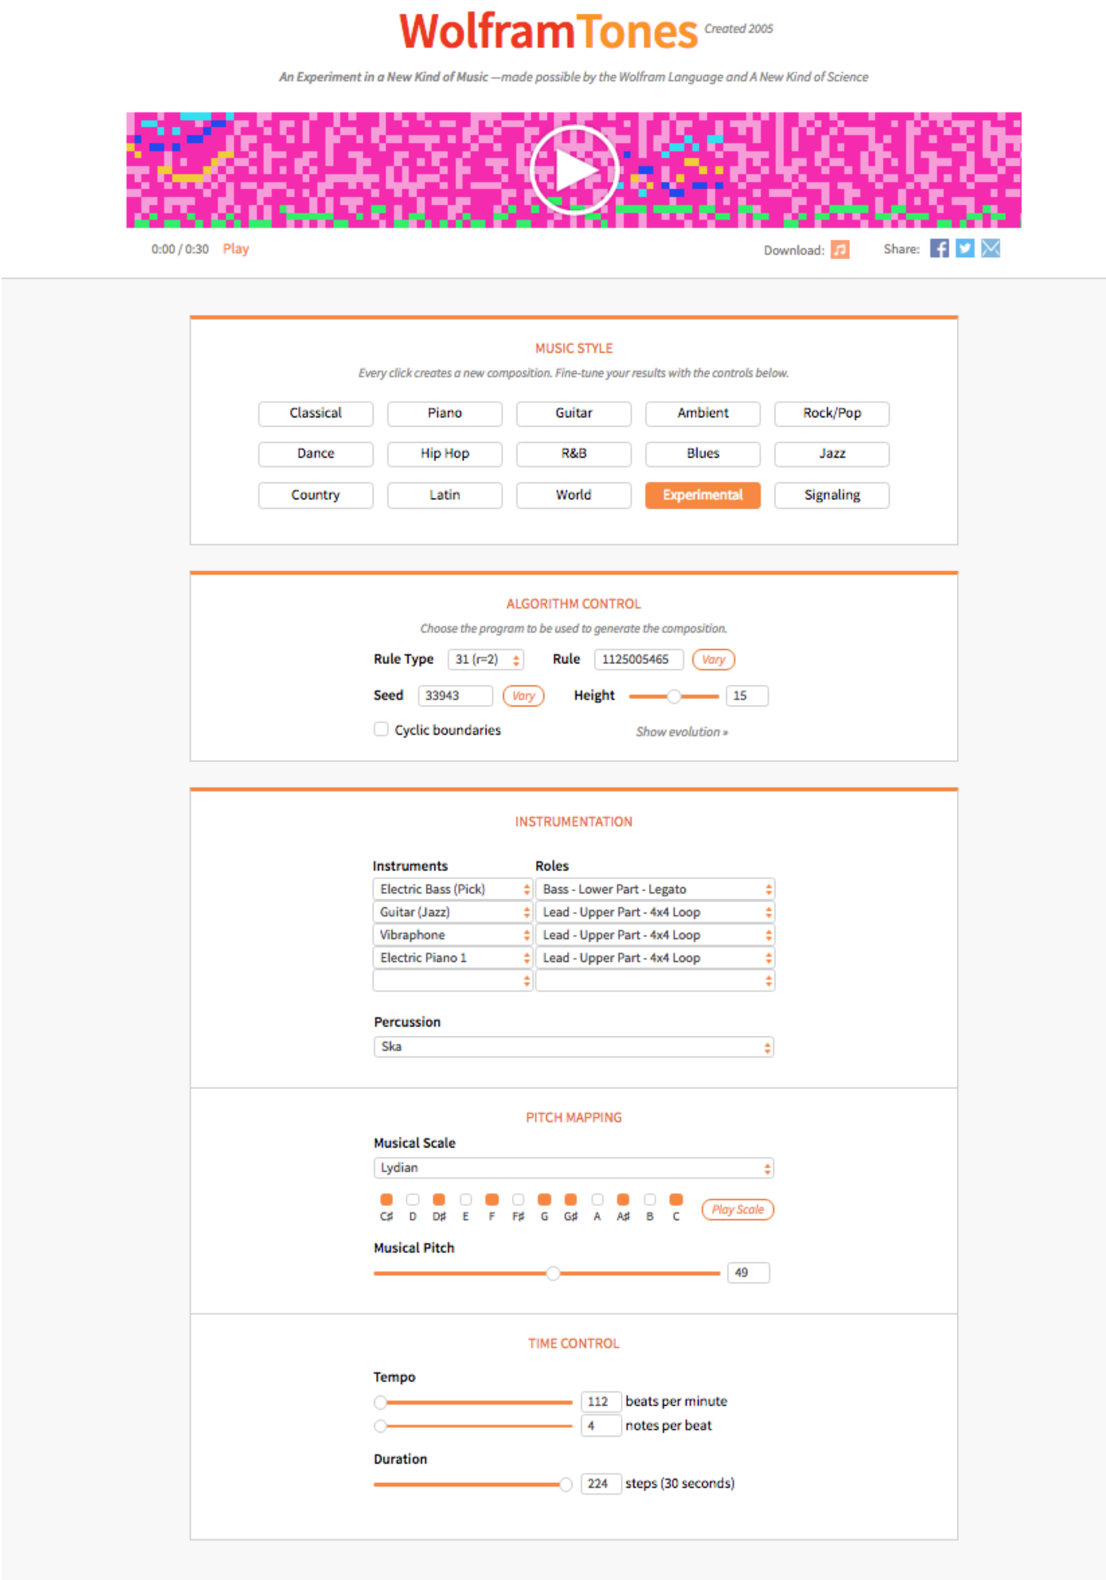

Figure 8: The WolframTones composition template. The example shown is the author's composition used as Sample 5 in the listening experiment.

Research experiment

To explore how different types of algorithms would be evaluated by listeners, five samples of algorithmic-generated music, in mp3 format, outlined in *Table 2*, were sent to five individuals for assessment. The listening samples had no description and were simply numbered 1 to 5.

Algorithm type	Composer/Designer	Genre	Experimenter
1. Aleatoric	Mozart: Dice Game	Classical	Troedson
2. Stochastic	Xenakis: *Achorripsis*	Experimental	Xenakis
3. Genetic algorithms-evo-devo	Vico, Quintana: *Melomics109*	'Popular'	Vico, Quintana
4. Neural Networks-LSTM RNN	Eck: Google *Magenta*	Blues	Eck
5. Cellular automata	Wolfram: *WolframTones*	Experimental	Author

Table 2: Five types of algorithmic music used in listening experiment

Source compositions are listed in Appendix 1.

Responses to the first four quantitative questions in the survey are outlined in *Charts 1 to 4*. The final question on the survey asked participants for feedback on their decision process, see Appendix 2.

Q1 In your opinion, do you think this piece of music was generated/composed by A) a computer/algorithm or B) a human composer?

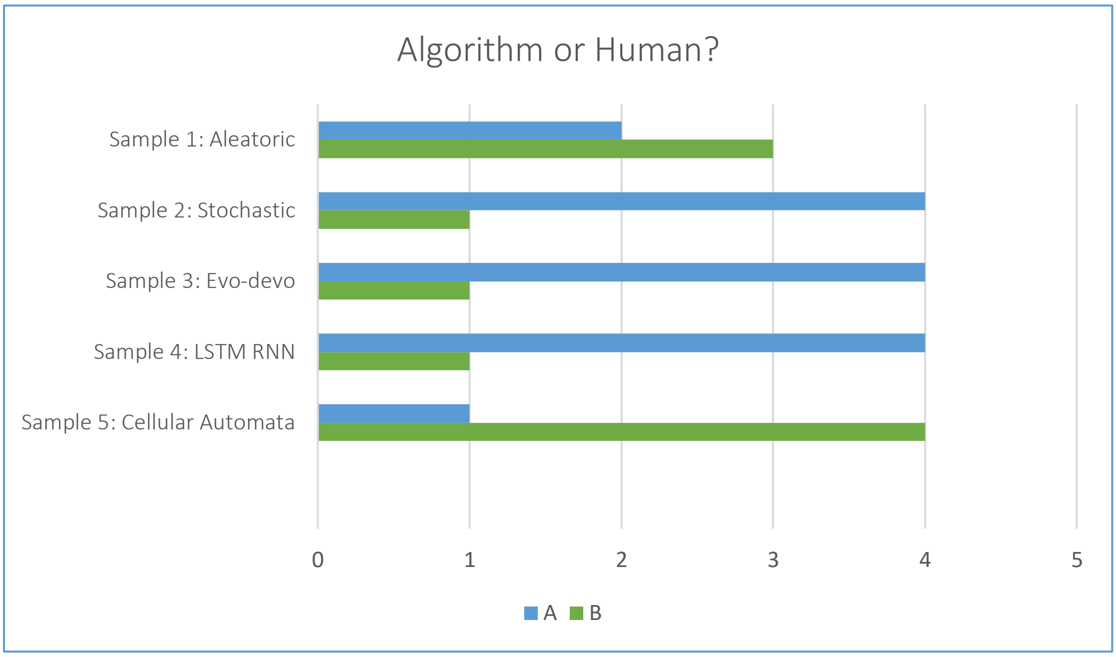

Chart 1: 4 out of 5 respondents selected algorithm for samples 2, 3 and 4. A human composer was the most popular choice for Samples 1 and 5.

In *Charts 2, 3* and *4*, survey respondents are colour coded 1 to 5 on the x axes.

Q2 On a scale of 1 to 10, how difficult did you find it to decide if it was A or B? (where 1 is easy, 10 very difficult)

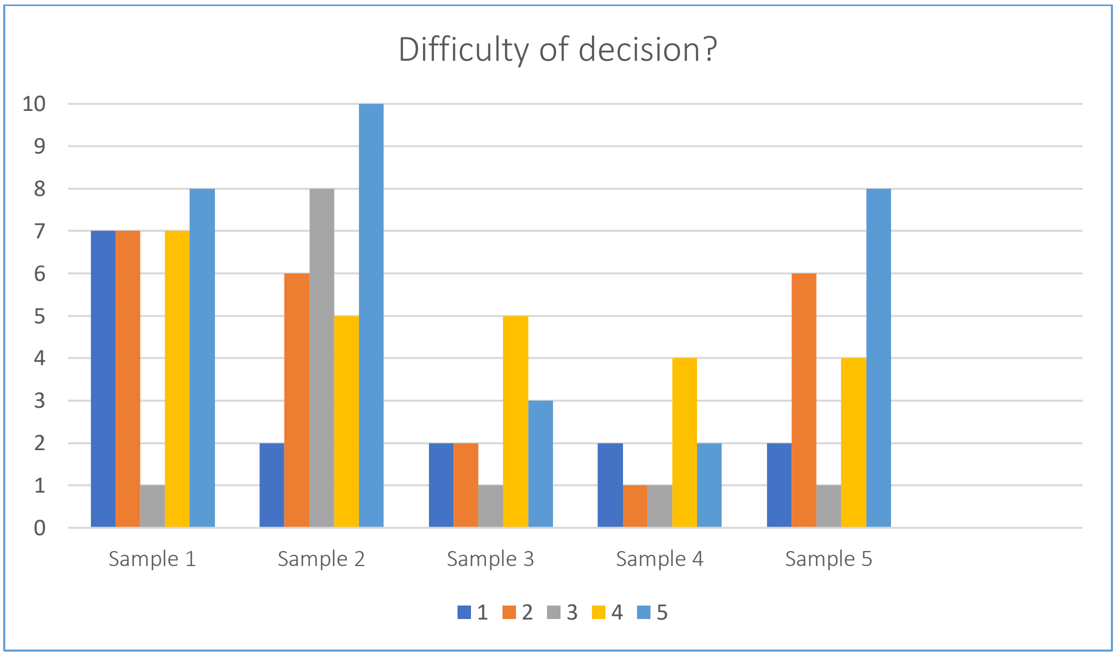

Chart 2: Samples 1 and 2 presented the most difficulty in deciding. However, there was a wide range of differences among listeners.

Q3 On a scale of 1 to 10, how structured or random do you think the composition is? (where 1 is very structured and 10 is completely random)

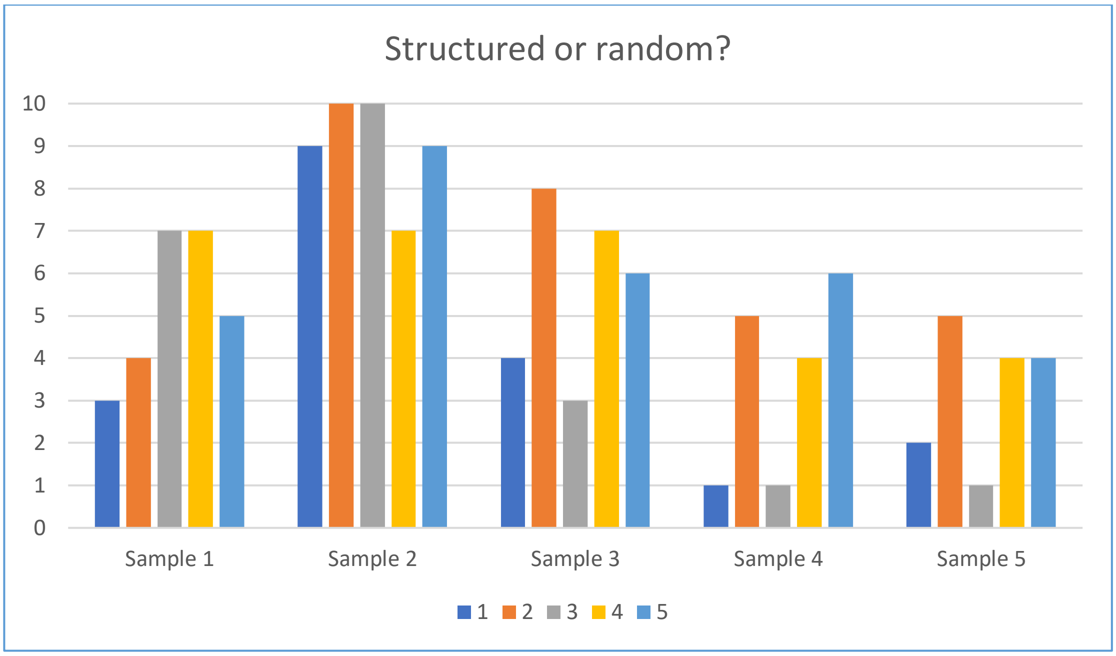

Chart 3: Sample 2 was assessed as the most random, while Samples 4 and 5 were thought to be the most structured.

Q4 On a scale of 1 to 10, what was your emotional connection with the music? (where 1 is 'it left me cold' and 10 is for example, 'I found this really interesting/exciting or 'I appreciated the musicality of the piece')

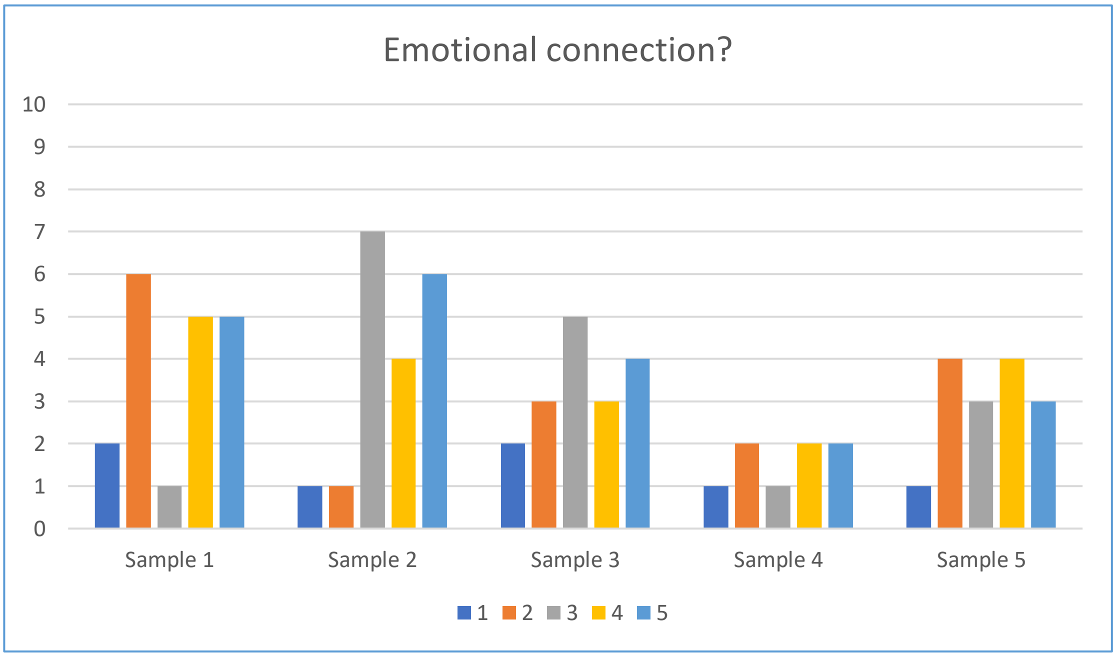

Chart 4: Few of the samples scored very highly on 'emotional connection'. Samples 1 and 2 had the highest scores, Sample 4 the lowest.

A few observations can be made from the results:

- No one selected ‘algorithm’ for all samples
- The degree of difficulty in deciding between algorithm and human composer varied greatly
- People don’t necessarily dislike ‘random’ music, i.e. Xenakis
- None of the pieces evoked a strong emotional connection and the lowest response was to music based on machine learning or neural networks

The music listening experiment also produced some interesting, candid comments (Refer to Appendix 2).

It is worth noting that the experiment only gave one example of algorithmic music in each category, so a second example could have produced different opinions.

Future Directions for AI and Music

This research paper has outlined a few of the significant developments of AI music in history, highlighted some current research in this area, and explored music produced by different types of algorithms and how listeners react to these compositions.

While music appreciation is a hugely subjective area, it is apparent that having the right balance of rules and unpredictability in the composition process is a key factor in producing music that is enjoyable to listen to. It is also clear that this is a collaborative process between humans and algorithms, at least for the present.

Hello World is an album announced as the first album composed with AI and used a cast of musicians to work with computer-generated scores to create the songs. (Hello World 2018)

Composer and curator of *Hello World*, Benoit Carré of *Flow Machines* said that while melodies, chords and sounds were suggested from the system, a human is always needed to stitch the songs together, give them structure and emotion, "There were many people involved in this…They gave their soul, their enthusiasm. I think that's the most important point of the album, in a way – that it's a very human one." (Marshall 2018)

However, music also has a more functional utility. Online platform *Amper Music*, for example, has designed its composition tools for applications such as music for, "video to gaming and other interactive technologies" and says that the, "massive growth of media requires a technological solution for music creation." (Donahue 2018)

Dredge (2017) quotes music industry consultant Mark Mulligan who suggests that this type of music is about sonic quality rather than music quality, "AI music is nowhere near being good enough to be a 'hit', but that's not the point. It is creating 21st-century muzak." (Dredge 2017)

Perhaps the machine learning approach being taken by many current researchers, may not the best alternative for producing truly creative, breakthrough music that is interesting or exciting to the listener?

However, is still early days for AI music research and very few musicians are likely to be working with AI music techniques currently – so there is considerable scope for new developments to emerge.

A final comment, based on the findings of this paper, is that it would be interesting to see more research in AI music created from other types of algorithms including stochastic or cellular automata.

References

Alpern A, 1995, 'Techniques for Algorithmic Composition of Music', Fall, p.1, Hampshire College

Broekhuyse P, 2017, Garry Kasparov vs Deep Blue: chess match lost to technology *The Australian* https://www.theaustralian.com.au/arts/review/garry-kasparov-vs-deep-blue-chess-match-lost-to-technology/news-story/779ca14fb487bcc309d66d50107e9e10 (accessed February 2018)

Brown A, 2004, 'An aesthetic comparison of rule-based and genetic algorithms for generating melodies' in *Organised Sound* 9(2) pp. 193-199

Brown A, 2005, *Making Music with Java: An introduction to Computer Music, Java Programming and the jMusic library,* p.207

Coutinho E, Gimenes M, Martins J M, Miranda E R, 2005, 'Computational Musicology: An Artificial Life Approach,' *School of Computing, Communications & Electronics*, University of Plymouth
http://pmc.usc.edu/PMJ/issue/4.1.01/paderewskirubato.html

Donahue L, 2018, Amper Music Raises $4M to Fuel Growth of Artificial Intelligence Music Composition Technology, *Globe Newswire*
https://globenewswire.com/news-release/2018/03/22/1444075/0/en/Amper-Music-Raises-4M-to-Fuel-Growth-of-Artificial-Intelligence-Music-Composition-Technology.html
(accessed April 2018)

Dredge S, 2017, AI and music: will we be slaves to the algorithm? *The Guardian*
https://www.theguardian.com/technology/2017/aug/06/artificial-intelligence-and-will-we-be-slaves-to-the-algorithm (accessed April 2018)

Eck D, Schmidhuber J, 2002, 'Finding Temporal Structure in Music: Blues Improvisation with LSTM Recurrent Networks', *IDSIA Istituto Dalle Molle di Studi sull'Intelligenza Artificiale*, p.1

Edwards M, 2011, 'Algorithmic Composition: Computational Thinking in Music' in *Communications of the ACM* 54(7) pp. 58-67

Fernández J D, Vico F, 2013, 'AI Methods in Algorithmic Composition: A Comprehensive Survey' in *Journal of Artificial Intelligence*, 48, pp. 513-582

Fox R, Khan A, 2013, 'Artificial intelligence approaches to music composition' in *Proceedings of the 2013 International Conference on Artificial Intelligence* 2 pp. 575–581

Hello World 2018
https://www.helloworldalbum.net/about-hello-world/

Jensen M G, 2009, 'John Cage, Chance Operations, and the Chaos Game and the "I Ching' in *The Musical Times* 150(1907) pp. 97-102

Johnson R C, 2002, Composer harnesses artificial intelligence to create music *EE Times*
https://www.eetimes.com/document.asp?doc_id=1145348 (accessed April 2018)

Loughran R, O'Neill, 2016, 'Generative Music Evaluation: Why do we limit to 'Human?', University College Dublin, p.1

Magenta: Make music and art using machine learning
https://magenta.tensorflow.org/ (accessed April 2018)

Mathieson S A, 2018, How artificial intelligence can aid and replace higher order human creativity *ComputerWeekly*
http://www.computerweekly.com/feature/How-artificial-intelligence-can-aid-and-replace-higher-order-human-creativity (accessed April 2018)

McDonald, 2017, Neural Nets for Generating Music, *Artists and Machine Intelligence*
https://medium.com/artists-and-machine-intelligence/neural-nets-for-generating-music-f46dffac21c0 (accessed April 2018)

Jordan L, 2017, Inside the Lab that's producing the first AI-generated pop album, *Seeker*
https://www.seeker.com/tech/artificial-intelligence/inside-flow-machines-the-lab-thats-composing-the-first-ai-generated-pop-album (accessed March 2018)

Marshall A, 2018, Is this the world's first good robot album? *BBC*
http://www.bbc.com/culture/story/20180112-is-this-the-worlds-first-good-robot-album

Miranda E R, 2003, 'Evolving Cellular Automata Music: From Sound Synthesis to Composition', *Sony Computer Science Laboratory*, Paris, France, p.1

Miranda E R, Williams D. 2015, Artificial Intelligence *in Organised Sound, Interdisciplinary Centre for Computer Music (ICCMR)*, Plymouth University, pp.7-8

Quintana C S, Arcas F M, Molina D.A, Rodriguez J D J, Vico F J, 2013, 'Melomics: A Case-study of AI in Spain' in *AI Magazine,* Fall p. 102

Orb Composer, 2018, The first Artificial Intelligence for Composers
https://www.orb-composer.com/ (accessed April 2018)

Savage M, 2017, Now computers are writing perfectly acceptable pop songs *BBC News*
http://www.bbc.com/news/entertainment-arts-41935971 (accessed March 2018)

Simon I, Oore S, 2017, Performance RNN: Generating Music with Expressive Timing and Dynamics, Magenta
https://magenta.tensorflow.org/performance-rnn (accessed April 2018)

Todd P M, 1989, 'A Connectionist Approach to Algorithmic Composition in *Computer Music Journal* 13(4) pp. 27-43

Troedson A, Mozart Dice Game n.d.
http://explodingart.com/jmusic/jmtutorial/MozartDiceGame.html (accessed March 2018)

Wolfram S, 2011, *Music, Mathematica, and the Computational Universe* http://blog.stephenwolfram.com/2011/06/music-mathematica-and-the-computational-universe/ (accessed April 2018)

References: Figures

Figure 1
Cellular Automata
Weisstein, E. Wolfram MathWorld
http://mathworld.wolfram.com/CellularAutomaton.html

Figure 2
Markov Chains
Powell, V. Setosa
http://setosa.io/ev/markov-chains/

Figure 3
Google Trends
https://trends.google.com/trends/explore?date=all_2008&gprop=news&q=%2Fm%2F0mkz

Figure 4
Jukedeck
https://www.jukedeck.com

Figures 5 and 6
Musicalgorithms
http://musicalgorithms.ewu.edu/index.html

Figures 7 and 8
WolframTones
http://tones.wolfram.com

Appendices

Appendix 1: Music listening samples

Participants were sent five music samples ranging in length between approximately 30 seconds to 2 minutes. They were asked to focus on the composition overall rather than the instrumentation used or the production quality.

The audio sources for the music examples are as follows:

1. Aleatoric: *Mozart Dice Game*
   http://explodingart.com/jmusic/jmtutorial/MozartDiceGame.html

2. Stochastic: Xenakis: *Achorripsis*
   https://www.youtube.com/watch?v=WasFTDq0dJI
   see also Stochastic Music:
   https://www.sweetwater.com/insync/stochastic-music/

3. GA evo-devo: *Melomics109*
   https://en.wikipedia.org/wiki/0music (Omusic 01)

4. LSTM RNN: *Magenta*
   http://www.iro.umontreal.ca/~eckdoug/blues/index.html (Composition 2)

5. Cellular Automata: *WolframTones*
   http://tones.wolfram.com/generate/GlPY20JC9ZeWXavOqaumvaqZM6hTL24FLO9CZQ4Y7dMV7

Appendix 2: Respondent comments

Q5 What factors led to you deciding your responses to the above questions, and any other comments?

The respondents' choice of A (algorithm) or B (human composer) is indicated after each quote.

1. Aleatoric

"Very hard to tell whether this was composed by a person or not. I'm sure a computer can do it. The rhythm was sometimes a bit random, which made me lean more towards the human composer. I was going to change my answer to A after listening to the other samples." (B)

"As far as composition is concerned, I could hear that this piece was decently structured yet I also felt that the outro was its weakest point. I could have enjoyed it more if there was a big ending or build up to something more fruitful.) (B)

"Of course I can't be sure what created it but I would hope a human composer would come up with something a little more moving and creative. It almost sounded like it was "playing within some parameters" and then doing anything within those parameters that ticked the box." (A)

"Phrasing, melodic harmonic relationship." (A)

"I wasn't sure if it was AI or human generated. I found the piece sweetly gentle but quite repetitive, thus my indecision regarding AI or human composer." (B)

2. Stochastic

"I think a program was used with random sounds thrown in together. Then again, a human could create this too." (A)

"This could have been composed by a human for sure as I have heard similar tasteless works over the years. If it was composed by a human it would be from one of which I would not admire or listen to by choice. It's just far too random and lifeless for me." (A)

"This piece was harder for me to pick. I thought it was possible it could have been a composition that had extended sections for instrumentalists to freely express themselves under the guidance of the conductor (like the crescendo in "A Day In The Life") but the more it went on I felt again that it sounded like a piece that was dynamically created from some set parameters. That said, I actually found it quite interesting to listen to and it conjured up cinematic images/moods for me." (A)

"Within the chaos there is significant structure possibly too difficult to create algorithm to reproduce." (B)

"Again, this was even more difficult to distinguish between human or AI composer. Primarily because the music could be interpreted as either Avant Garde and experimental or some AI generated random notes and sounds. I actually enjoyed this piece." (A)

3. **Genetic algorithm**

"It sounds like 3 tracks mixed together." (A)

"This piece is a wild card, one that could go either-way. However, there was no real thought or emotion throughout that says to me that it's definitely AI." (A)

"Right from the start I felt this piece had more of a human feel to it. An almost indescribable "design and direction" that I imagine a computer wouldn't be able to replicate. That doesn't necessarily equate to emotion though. For that I gave a middle score of 5. I appreciated what I felt were human qualities but Sample 2 I found more interesting in some places." (B)

"I have similar sounding compositions from AI programs." (A)

"This piece is quite bland and it seems to jump a little from section to section. It seems smooth and flowing but feels like something is missing. I think it is AI generated or a very simple "trick" by a human composer to fool me." (A)

4. Neural Network

"Sounds very easy - a simple, repetitive melody combined with some chords." (B)

"This piece reminded me what one hears when trying to get through to a telecommunications help desk. Nothing in it whatsoever for anyone to enjoy, admire or appreciate. I've seen Casio keyboards randomise demos with more body and soul." (A)

"Obviously a repeating jazz blues, I gave a 1 for structure but that's in regards to the accompaniment only. The melody/solo sounds completely soul-less as if dynamically created by a computer in response to "allowed notes and rhythms" parameters. Even when the "soloist" starts playing outside towards the end it still sounds manufactured and misses the inventiveness, passion and emotion that comes from a human composer/player." (A)

"Too disjointed." (A)

"This piece was really quite boring. It is definitely a computer-generated piece in my opinion. At first, I thought a little "Jazzy" in feel but that vanished quickly." (A)

5. Cellular automata

"A repetitive melody mixed with one or two other tracks." (A)

"This piece started off ok, it almost had a certain expectation that it would lead e on an audible journey yet not far into it everything about it just fell through the floor. It all the signs of an artists that lacked musical knowledge yet was using available tools to create 'something' even though it ended up to be not very pleasing to the listener." (B)

"Admittedly very unscientific, all I can say with this one is it sounds of human origin. The repeated regular/staccato vibe part, the cool kick drum groove, the distorted guitar etc all sound like a human composition. Actually, they sound like musicians jamming. Even with a specific, detailed algorithm I couldn't see a computer coming up with this. The only reason I gave the emotional connection a 3 is it didn't move me that much. Personal taste." (B)

"The phrasings too loose." (B)

"This piece was very difficult to discern regarding AI or human. It had a sense of structure and yet I could not make up my mind about its origins. If I could I would have chosen both A or B in question 1." (B)