2015-04

Artificial Intelligence in<i>Organised Sound</i>

Miranda, Eduardo

<http://hdl.handle.net/10026.1/6521>

10.1017/s1355771814000454

Organised Sound

Cambridge University Press (CUP)

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

Artificial Intelligence in Organised Sound

Eduardo R. Miranda and Duncan Williams

Interdisciplinary Centre for Computer Music Research (ICCMR)

Faculty of Arts and Humanities

Plymouth University

Plymouth PL4 8AA

United Kingdom

E-mails: eduardo.miranda@plymouth.ac.uk; duncan.williams@plymouth.ac.uk

Artificial Intelligence is a rich and still-developing field with a number of musical applications. This paper surveys the use of Artificial Intelligence in music in the pages of *Organised Sound*, from the first issue to the latest, at the time of writing. Traditionally, Artificial Intelligence systems for music have been designed with note-based composition in mind, but the research we present here finds that Artificial Intelligence has also had a significant impact in electroacoustic music, with contributions in the fields of sound analysis, real-time sonic interaction and interactive performance-driven composition, to cite but three. Two distinct categories emerged in the *Organised Sound* papers: on the one hand, philosophically and/or psychologically inspired, symbolic approaches, and on the other hand, biologically inspired approaches, also referred to

as Artificial Life approaches. The two approached are not mutually exclusive in their use, and in some cases are combined to achieve ‘best of both’ solutions. That said, as *Organised Sound* is uniquely positioned in the electroacoustic music community, it is somewhat surprising that work addressing important compositional issues such as musical form and structure, which Artificial Intelligence can be readily applied to, is not more present in these pages.

1 INTRODUCTION

Artificial Intelligence (hereafter, AI) concerns the development of human-like intelligence, typically in software. The ultimate goal of many AI approaches is to produce optimal or super-human solutions; this might be to augment, or indeed to entirely replace the role of composer, performer or listener. In the context of electroacoustic music, this might means assisted composition, human-like performance rendering of scores, sound organisation and/or advanced synthesis control. Philosophical questions are often raised by those working with AI in music. For example, in the context of the above applications, a question that is often asked is this: Who is the composer if AI has been used in the creation of new music?

Organised Sound is well placed as the leading platform for electroacoustic composers, sound designers, sonic artists and the like to explore such type of questions in their own practice. This article surveys the approaches to AI that have been taken by researchers in the pages of *Organised Sound* through the last twenty years. The reader should be advised that although we are aware

of other significant progress in the field – for example, AI has many applications in musical analysis and music education – the objective of this exercise is to evaluate what the journal has captured regarding the developments in this important field of music, and as such these developments are mainly focussed on composition and/or interactive performance.

We will assume that the reader has some knowledge of the most prominent AI terminology. Therefore, for the sake of brevity we provide details of working processes only when they are absolutely essential to the context of the paper. The timeline is loosely chronological, and sees two main camps emerge amongst the approaches. We examine these in more detail and compare the approaches, and suggest the likely next steps in this fertile and evolving field.

2 THEN AND NOW

“Composers, musicians and computer scientists have begun to use software-based agents to create music and sound art in both linear and non-linear ... idioms, with some robust approaches now drawing on various disciplines.”

(Whalley, 2009)

Much of the work found in the timeline of AI in *Organised Sound* can be classified by approach. For example, work towards compositional models using AI begins with symbolic approaches, using machine learning or human input to determine rules for the creation of tonal music. Such systems have

appeal to composers who are familiar with symbolic approaches to music representation; such as for instance, composing using traditional classical music notation. Conversely, connectionist approaches, based on ‘black box’ neural networks with no symbolic musical representation, have seemingly not made huge inroads in the pages of *Organised Sound*. Although the chronology is not strict, broadly speaking work using AI such that the system can self-program begins to surface later in the timeline, and often feature some comparison or combination of the two. In any case, such systems include distributed autonomous agents, genetic algorithms, flocking or swarming simulations, and neural networks. Essentially, here goes anything that falls under the banner of simulating some sort of biological process, or living beings, with cognitive learning and evolving behaviour, above and beyond that of an initial structural rule-set. In this paper, we consider that the former category is a symbolic approach to AI in music, whilst the latter is a biologic approach.

2.1 Symbolic approaches

Symbolic approaches to machine learning in AI are synonymously referred to as traditional AI, or GOF AI: Good Old Fashioned AI. They are concerned with hard-coding a set of rules, which prescribe the behaviour of the machine. The choices for developing the rules are manifold, as Collins points out:

“In case it is still unclear, any algorithmic method might be applied, and this potentially includes all artificial intelligence techniques. The extent to which

such algorithms have yet to be harvested makes this an open research area; there are favourite techniques, controlled probabilistic expert systems being a typical route”

(Collins, 2008)

One such probabilistic example would be a series of rules or constraints whereby the machine is taught to generate note structures within various degrees of aleatoric likelihood. Casey (2001) presents and discusses a system for selection of sound based on automated classification and further for subsequent sound organisation. This system use a machine learning method to differentiate between musical sounds, genre, speech, and environmental sounds, and can be used creatively by matching target sounds to other sounds from a database. This forms the core of a sound generating method referred to as ‘concatenative synthesis’. The application of a particular type of probabilistic logic is illustrated by Casey: that of hidden Markov models, or HMM. HMM models are common in algorithmic composition tasks and Casey’s application to sound selection and organisation further illustrates they are also suitable for use in electroacoustic composition. HMM represent the likely behaviour of events where a given event can adopt any one of a range of states; in a note-based algorithmic composition these might be note sequences, rhythms and so on. Changes in states are known as transitions, and it is the probability of given transitions that can be used to algorithmically generate sequences of notes based on an analysis of the transitions between states in some selected source material. Visell (2004) introduced the

application of a spontaneously organising HMM to the analysis and synthesis of statistically-driven stochastic music, based on pattern theory analysis. This is an analysis technique that derives the principles for a generative model from specified source signals. In both cases, material that follows the general structural rules of the source signals can be generated whilst still allowing for variation and variability in the result. Further, Vissell also gives some consideration to Artificial Life alternatives, which we will discuss in the next section, specifically addressing the alleged advantage of the (HMM) system thus: “[the]... essential difference is that the standard artificial neural networks do not attempt to model directly the native domain of the signal (time, in the case of sound signals). Consequently, the possibilities for structural refinement based on the analysis of output relative to natural signals are more limited.”

(Visell, 2004)

Indeed, several of the papers we have surveyed document some comparison between these two streams of AI - that is, traditional/symbolic and Artificial Life/biological systems - but consideration of the capability of AI to address structural issues is less common. This is particularly surprising given the well-documented applications of AI to music information retrieval, and something which might well be addressed in the future as a practical research question by those working in the field of electroacoustic music.

Other complex probabilistic systems of sound selection and organisation can also be found, such as the fuzzy logic based approach presented by

Eigenfeldt and Pasquier (2010). Fuzzy logic allows for degrees of reasoning in the AI, rather than exact Boolean values or a series of gate-based logics for decision making. Eigenfeldt and Pasquier also present a method for sound organisation with a self-organising map, abbreviated as SOM, which is essentially an artificial neural network. In this system, perceptual proximities in sounds timbres are assessed on the Bark scale: 24 critical bands of frequencies which are correlated to various psychoacoustic responses; perceived brightness, sharpness, and so on. This analysis provides values for similarity that are used to build connections in the SOM, which can then be navigated in a real-time process of sound organisation. Nevertheless, these rules remain pre-defined, and do not evolve autonomically. Once the rules are established and the system has been given the input parameters (in the examples above, a database of sounds), musical results can then be evaluated by the user. Again, autonomic evaluation (machine learning, genetic algorithms and the like) are not present in traditional symbolic approaches.

2.1.1 Evaluating music produced with AI

Clearly, seemingly simple rule-based systems are able to create new sequence of musical notes or ordered sounds, and in the case of aleatoric rule-sets, a near-infinite amount of variety. But how can we evaluate the music produced by these approaches? A traditional way to evaluate the success of any artificial intelligence system is the so-called 'Turing Test'. This test allegedly evaluates whether an AI system has created material, which is indistinguishable from material created by a human. The original Turing Test was developed to evaluate computer generated text, but the more popular

procedure, which is commonly applied in musical AI is much simplified. The test is often adaptable to music simply by asking listeners the following question: Do you think this piece of music been composed by a human or by a machine? However, beyond the evaluation of rule-based success in symbolic AI systems, 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? Does the authorship rest with the rule-maker? The rule-programmer? A casual user who chooses new seed material as an input for the system? Or does this duty begin at the selection stage of the process? Is the author in fact the decision maker who evaluates the generated material? Or is the author in fact the machine itself?

Philosophically, Jacob summarizes the evaluation of success in such AI as:

“... how to program a computer to differentiate between ‘good’ and ‘bad’ music. The philosophical issues reduce to the question who or what is responsible for the music produced?” (emphasis: original author)

(Jacob, 1996)

We might then conclude that the issue of aesthetic musical quality in the creation of AI assisted music is moot. Implicit in Jacob’s summary is that we

have seen the development of AI in practical terms reach a level where it is possible to model the knowledge base of a human composer, and that the questions regarding such work are solely to do with authorship, authenticity, and creativity. When discussing his own AI-based approach to real-time composition, Eigenfeldt (2011), addresses this more directly when he suggests that the act of designing the complexity of interactions between agents is a compositional act in itself.

Eigenfeldt is not alone in wishing to stress the ownership and authorship of the music when AI is involved, though Dahlstedt feels more conflicted in this regard:

“I have a slight feeling I did not write that music, and yet I am quite sure no one else did. I designed the algorithm, implemented it and chose the parameters, and still I feel alienated.”

(Dahlstedt, 2001)

Certainly many involved in algorithmic composition can find themselves in agreement with either end of this scale of ownership, though the real-time nature of these particular systems is something of a special case. It nevertheless highlights that the difference between structural and performing rules is perhaps a smaller one in the field of *Organised Sound* than one might expect in more traditional music composition, where AI is often employed solely to create human-like performances of music sequenced or scored in an otherwise traditional manner. One conclusion we might draw is that the whole,

regardless of parts-composition and parts-performance, is of central importance to those of us working with electroacoustic music.

2.1.2 Imitation of style with symbolic AI

Let us consider a number of other applications for this type of symbolic AI in such music. When the rule-set can be derived from another input, for example an existing piece of music rather than being pre-determined in some other fashion (as in the HMM examples given above), the effectiveness of the learning may be judged as a measure of successful imitation in the output.

Systems for mimicking a composers style by training the rule-set in this manner exist and have been used successfully, as documented in Ron Geesin's review of David Cope's *The Algorithmic Composer* (Orton, 2000) wherein *Alice* (Algorithmically Integrated Composing Environment) is able to extrapolate rules from source material (and thus, compositional 'style' from material contained in the source database) without the need for the composer to specify a rule-set in advance:

"Cope warns [that] the user should not imagine that composing with Alice is necessarily easier than composing without its aid. The choice of the musical material for the database ... is critical ... A poorly matched database can only give poor results."

(Orton, 2000)

In this example the database of source material clearly becomes an important part of the generative music process, and implicitly, the evaluation of success is in the ear of the beholder.

Another real-time example of style imitation by means of AI can be found in the automatic generation of a musical accompaniment. However, learning in real-time requires a more complicated approach to the AI than symbolic approaches alone can afford. For example, Cunha and Ramalho (1999) proposed a system for generic automatic accompaniment achieves good results by combining a symbolic approach with a neural network. As with the work with SOM by Eigenfeldt and Pasquier, we consider that Cunha and Ramalho's system also falls in to the second category of AI: that of biologically inspired, Artificial Life approaches.

2.2 Biologic approaches

Biologically inspired AI include systems using neural networks, distributed agents, genetic algorithms, and flocking simulations, all of which have made their way into the pages of *Organised Sound*. One of the fundamental differences between these and the symbolic approaches documented above is in the learning process. Unlike symbolic approaches, these systems can often continue to adjust their rule-sets, potentially developing further without continued human intervention. To some extent this gives a way to address the issue of creativity that symbolic approaches found philosophically challenging

Dahlstedt (2001) introduced a system called as *MutaSynth*, which fosters a way to explore interactive composition by modelling basic evolutionary processes through sounds. Here, operators referred to as genetic modifiers are applied to create mutations and variations from parent sounds, before the user selects the outputs they have a preference for. This preference is analogous to a fitness function in the field Evolutionary Biology, and as such can be automated with AI in evolutionary models. Indeed, this is an issue that is important to consider in any such approach.

Whalley (2004) describes two possible ways to consider the issue of fitness: evolutionary systems and intelligent software agent systems, with the goal of developing a intelligent machine capable of having a musical, interactive conversation. Whalley settled on intelligent agents, which are devices that can make informed decisions, move within a network, and learn in response to their environment over a software-based evolutionary system. In order for this approach to be conversationally interactive, the system must be able to listen and respond appropriately to human input, as well as to initiate conversations of its own accord. Each interaction the agent experiences will thus enhance its own learning. Musical parameters including tempo, dynamics, and other acoustic features (e.g., panning, audio effects, etc.) are then mapped to performance gestures. An interesting aspect to such a system is in the continuous exchange of ideas between human users and the AI, unlike systems which only allow for human interaction at the beginning or the end of the process, such as, setting parameters, selecting source materials,

evaluating results, and making aesthetic decisions about good or bad, and so forth.

2.2.1 Self-organisation

Self-organisation implies a degree of cognitive ability, in the case of multiple agents to interact, respond, and create structure on a localised level.

Blackwell and Young (2004) described their own system for creating self-organised music by interpreting musical parameters from swarm dynamics, as might be exhibited by flocks of birds, herds of animals, or groups of co-operating insects. Swarms are modelled by local interactions between agents or particles, rather than a higher-level control. Moreover, users can interact with the particles of the swarms to influence their behaviour. Again, this shows the use of biologically inspired AI to create a system that can adjust its behaviour in real time, in continuous response to human input.

2.2.2 Imitation of styles with neural networks

As we briefly mentioned in the symbolic approaches category, another use case for such a system is that of the creation of automatic musical accompaniments for a human performer. Cunha and Ramalho (1999) described their application of a neural network for this task. Their system, which has already been trained in harmonic development, is capable of generating real-time accompaniment to songs it has not previously been exposed to by means of a prediction model. Neural networks are well

documented in AI as rough models of biological neuronal function and can accommodate a high level of complexity in their functioning. Neural networks developed in response to music - for example, developed in response to specific source material - would present a conceptually different solution to symbolic approaches for generating probabilistic rule-sets. The distinction is that the neural network develops connections rather than strict rules, which may give a unique perspective to systems operating outside of the note-based approach to music creation often taken in the work we survey here. Nevertheless, Cunha and Ramalho note that by the addition of a rule-based tracker to their neural network predictor, the performance of the resulting, hybrid model, was improved.

2.2.3 Genetic Algorithms and other combined approaches

Brown (2004) directly compared the aesthetics of melodies produced by both the symbolic (rule-based) and biological (Genetic Algorithm) approaches (Brown, 2004), and suggested that a combination of techniques yielded the most 'aesthetically appropriate' (sic) musical results. Other applications of Genetic Algorithm techniques to algorithmic composition have also been explored elsewhere in *Organised Sound* by Collins (2002), and Manzolli and colleagues (Manzolli et al. 1999). The former looked into providing control of sound synthesis parameters, and the latter into generating and evaluating chord progressions. Metaphorical musical genes, coding musical or sonic phenomena, are mutated and then evaluated via a fitness function. In the case of Collins' synthesis-driving system, the fitness function ultimately

remains the choice of the user, which the author refers to as “the fitness bottleneck of the human decision-maker”. Manzolli and colleagues also acknowledged the fitness function, but instead, they created a statistical function based on an analysis of existing memories, showing another way to incorporate the probabilistic rule-based approaches used in the symbolic AI stream.

3 CONCLUDING REMARKS

Whilst consciously remaining non-exhaustive in this paper, we nevertheless find that AI has had a definite presence in the pages of *Organised Sound*, with both symbolic rule-based systems, and biologically inspired Artificial Life approaches being used to create new work by the electroacoustic community, as well as a number of combined approaches which document good results. AI gives a rich pool for those interested in algorithmic composition to develop new systems and indeed to evaluate the musical effectiveness of their output.

The philosophical questions raised by the use of AI in creating music are also not overlooked in these pages, though the traditional questions, which might be used to evaluate the success of AI in such applications are perhaps less relevant to the electroacoustic community. The electroacoustic community seems less sensitive to issues of authentic human performance of classical music and more concerned with the aesthetic results, which might be obtained. The issue of whether or not the material generated is readily

distinguishable from human output when carrying out such processes entirely by hand is seemingly not relevant. Thus we find that the training of the AI, from rule demarcation to source material selection still constitutes the process of composition.

Given that *Organised Sound* has become arguably the foremost journal of the electroacoustic community, it is surprising that it lacks papers tackling the problem of overall musical form in composition, though some such work exists, for example David Cope's aforementioned *Alice* system used a musical phrase classification algorithm to enable the generation of music with formal structure and coherence in the compositions. AI has been well used as an analysis tool to determine and describe musical structure by means of structural representations or acoustic analysis. So, we might speculate that the absence of other such structural analysis by AI is because the electroacoustic community is not always so interested in directly addressing note-based music. Iannis Xenakis' UPIC system (Xenakis, 1996) for graphic scoring allows for structure in the linking of its pages, which become analogous to the score and the structure in note-based music. UPIC has already been shown in *Organised Sound* to be well-suited to learning applications (Nelson, 1997; Bourotte and Delhayé, 2013) so perhaps a method of training AI with UPIC as the interface would be welcomed by practitioners from the electroacoustic community.

More recently, Artificial Life approaches show that the application of AI to electroacoustic music creation has yet to reach maturity: it is still an open, and growing field of research.

It is not a trivial task to predict what AI might contribute to electroacoustic music in within the next 20 years of *Organised Sound*. But we suspect that AI informed and inspired by Biology will continue to evolve, in particular developments pertaining to Computational Neuroscience, where scientists are developing increasingly more sophisticated models of the brain. This research is providing us with better understandings of how our brain works. Such understanding is bound to result in new technological and also theoretical developments for music. Unfortunately, scientific progress at this front so far has been largely insignificant for music, as most of this progress has been on visual processing. The truth is that auditory processing turns out to be fiendishly more complicated than we had previously thought. Consequently, our current understanding of how the brain processes music lags far behind our understanding of other brain functions. Despite a fair amount of research that is being developed within the emerging field of Cognitive Neuroscience of Music, progress so far has been disappointing and profoundly irrelevant to musicians in general, and in particular to the electroacoustic community. One of the problems that we can identify here is that the great majority of scientists working in this field, and consequently, their respective peer reviewing community, lack knowledge of music, therefore rendering their experiments largely flawed. Still, we believe that the future is bright. A better understanding of how the brain listens to sound is bound to lead to new technology for the

analysis of electroacoustic music based on neurophysiologic models of our auditory system. For instance, in addition to today's cochleogram, which is based on how our inner ear analyses the spectra of sounds, in the future we might be able to build tools along the lines of a *thalamogram* (Miranda, 2010). This analysis tool would give information related to the activity of a functional region of the brain referred to as the thalamus. The thalamus plays an important role in controlling attention: it enables the brain to suppress information in order to focus on particular aspects of incoming sensorial information, including sounds. The thalamogram would reveal salient sound attributes that would be deemed more important than others in function of specific musical contexts or conditions. We would envisage the possibility of being able to specify such contexts as analyses parameters for simulating the focus of the thalamus under different contexts or conditions.

We surely need to see more musicians walking in the corridors of neuroscience laboratories if contributions of AI to electroacoustic music is to continue to be reported in the pages of *Organised Sound*.

References

Blackwell, T. and Young, M. (2004). "Self-organised music". *Organised Sound*, 9(2):132-136.

Bourotte, R. and Delhayé, C. (2013). "Learn to Think for Yourself: Impelled by UPIC to open new ways of composing". *Organised Sound*, 18(2):134–145.

Brown, A. R. (2004). "An aesthetic comparison of rule-based and genetic algorithms for generating melodies". *Organised Sound*, 9(2):191-198.

Casey, M. (2001). "General sound classification and similarity in MPEG-7". *Organised Sound*, 6(2):153-164.

Collins, N. (2002). "Experiments with a new customisable interactive evolution framework". *Organised Sound*, 7(3):267–273.

Collins, N. (2008). "The Analysis of Generative Music Programs". *Organised Sound*, 13(3): 237-248.

Cunha, U. S. and Ramalho, G. (1999). "An intelligent hybrid model for chord prediction". *Organised Sound*, 4(2):115–119.

Dahlstedt, P. (2001). "A MutaSynth in parameter space: interactive composition through evolution". *Organised Sound*, 6(2):121-124.

Eigenfeldt, A. (2011). "Real-time Composition as Performance Ecosystem". *Organised Sound*, 16(2):145–153.

Eigenfeldt, A. and Pasquier, P. (2010). "Real-Time Timbral Organisation: selecting samples based upon similarity". *Organised Sound*, 15(2):159–166.

Jacob, B. L. (1996). "Algorithmic composition as a model of creativity".
Organised Sound, 1(3):157–165.

Manzoli, J., Moroni, A., Von Zuben, F., and Gudwin, R. (1999). "An evolutionary approach to algorithmic composition". *Organised Sound*, 4(2):121–125.

Miranda, E. R. (2010). "Organised Sound, Mental Imageries and the Future of Music Technology: a neuroscience outlook". *Organised Sound*, 15(1):13-25.

Nelson, P. (1997). "The UPIC system as an instrument of learning".
Organised Sound, 2(1):35–42.

Orton, R. (2000). "David Cope, The Algorithmic Composer". Book Review.
Organised Sound, 5(2):111.

Visell, Y. (2004). "Spontaneous organisation, pattern models, and music".
Organised Sound, 9(2):151-165.

Whalley, I. (2004). "PIWeCS: enhancing human/machine agency in an interactive composition system". *Organised Sound*, 9(2):167-174.

Whalley, I. (2009). "Software Agents in Music and Sound Art Research/Creative Work: current state and a possible direction". *Organised Sound*, 14(2):156-167.

Xenakis, I. (1996). "Tutorial Article. Determinacy and indeterminacy".
Organised Sound, 1(3):143–55.