

# Open Problems in Evolutionary Music and Art

Jon McCormack

Centre for Electronic Media Art,  
School of Computer Science and Software Engineering,  
Monash University, Clayton 3800, Australia  
jonmc@csse.monash.edu.au  
<http://www.csse.monash.edu.au/~jonmc>

**Abstract.** Applying evolutionary methods to the generation of music and art is a relatively new field of enquiry. While there have been some important developments, it might be argued that to date, successful results in this domain have been limited. Much of the present research can be characterized as finding ad-hoc methods that can produce subjectively interesting results. In this paper, it is argued that a stronger overall research plan is needed if the field is to develop in the longer term and attract more researchers. Five ‘open problems’ are defined and explained as broad principle areas of investigation for evolutionary music and art. Each problem is explained and the impetus and background for it is described in the context of creative evolutionary systems.

## 1 Introduction

*Education is not the filling of a pail, but the lighting of a fire.*

— W.B. Yeats

Music, art, and indeed creativity in general are defining traits of the human condition. Moreover, they are one of the primary reasons why we consider the richness of living experience to be more than just one of survival and reproduction, even though, ironically, they may have biological origins and purposes [1, 2, 3, 4]. In recent years, evolutionary computing (EC) methods have been applied to problems in music composition and art.

The premise in creating and researching evolutionary music and art (from here on EMA) is that creative problems are *non-trivial*. Creativity is considered a positive and sought-after trait in all human cultures [5]. Various hypotheses have been put forward for this, for example musical ability has been hypothesized as the result of sexual selection [6]. Moreover, creativity encompasses a broad scope of tasks in terms of psychology, problem solving, judgment and action [7].

Research in EMA has covered a variety of problems in aesthetics, creativity, communication and design. Broadly speaking however, much of the research and results have been ad-hoc; common methodologies have included: ‘use technique  $X$  from complexity research to make images or music’; ‘use aesthetic selection to evolve  $X$ ’; ‘devise a suitable fitness function to automate the evolution of  $X$ ’; and so on. Certainly there have been many successes using such strategies. However, even in this

glib set of scenarios there is a sense that these strategies are simply combinatory products of a set of well-explored ideas from other disciplines (admittedly in a different context).

The aim of this paper is to define a small set of ‘open problems’ in EMA. The goal is not to be critical of previous work, but to provide a well-defined set of challenges for the EMA research community. Such an approach has been successful in other disciplines [8].<sup>1</sup> After defining an important distinction in modes of research, the remainder of this paper presents five open problems specific to EMA and discusses some current research and background for each of them.

### 1.1 Evolution and Art

Before introducing the set of open problems, a key difference in research goals needs to be introduced. In classifying approaches, I make a primary distinction between *(i)* research where the resultant music and artwork is intended to be recognized by humans as creative (i.e. art) and *(ii)* research which explores the concept of creativity in general.

The first case is seemingly more straightforward so we will examine it first. In this case the results generated by the system and/or methodology are intended for human appreciation as art. That is, they exhibit properties that humans recognize as displaying some form of creative intention or aesthetic judgment by the creator (we conveniently ignore whether by person or machine). People who use such techniques might call or consider themselves ‘artists’ in addition to (or as opposed to) ‘researchers’.

The second case is different. Here, creativity is considered in a more open context, that is, it is not limited to being recognized by people as creative or aesthetic. In broad terms, creativity is not found exclusively in human behaviour. Bowerbirds, for example, create elaborate aesthetic constructs that serve no direct survival advantage, rather act as displays to attract mates. In the simulation context, a similar parallel exists in artificial life research. In Langton’s seminal paper [9], the concept of ‘life-as-it-could-be’ is introduced. Life-as-it-could-be represents a broader set of living systems, beyond the ‘life-as-we-know-it’ life observed on Earth. This broader definition of life admits the possibility of other kinds of life, radically different than what we currently know as life. For example, in Fredrick Hoyle’s classic novel *The Black Cloud* the Earth is visited by a radically different life-form that is unaware of the destruction it is causing on Earth until it is contacted by astronomers. Of course, this is science fiction, and it has been argued that life-as-it-could-be can never be that different from life-as-we-know-it because we could never recognize it as life [10]. Even more speculative is the idea that life-as-it-could-be might create its own art: *art-as-it-could-be*. That is, artistic products or systems created and analyzed by synthetic autonomous agents. The second case is dealt with in more detail in Section 2.4.

---

<sup>1</sup> At least in terms of getting the authors of the paper numerous citations. By posing difficult problems, most have not been solved, ensuring the longevity of the paper. This strategy is employed in this paper as well.

## 2 The Open Problems

This section introduces the open problems. A discussion surrounding each is also presented.

### 2.1 Searching for Interesting Phenotypes

A common practice in EMA is one of *search for interesting phenotype*. In this scenario, the artist or programmer designs some form of parameterized system. The system generates output, typically in the form of sound or image. In most cases, the number of parameters is very large, making an incremental or ordered search of the entire parameter space intractable. Hence the use of other search techniques such as genetic algorithms or aesthetic selection.

In this mode of EMA there are two primary considerations:

1. the design of the generative system and its parameterization;
2. the evaluation of the fitness of phenotypes produced by the system.

In the case of aesthetic selection, the fitness evaluation is implicit, being performed by the user of the system. I will return to the second consideration in a later section, for now let us examine the first point in more detail.

The well-known system of Karl Sims generated images using Lisp expressions evolved by aesthetic selection [11]. In essence these expressions were a combination of basic arithmetic operations and standard mathematical functions such as trigonometric and fractal functions. Even with a limited number of such expressions, the range or gamut of possible images is extremely large. However, it turns out that all of the images produced by such a system are of a certain ‘class’ — that is they all look like images made using mathematical expressions. While there might exist a Lisp expression for generating the *Mona Lisa* for example, no such expression has been found by aesthetic selection.<sup>2</sup>

Steven Rooke extended the aesthetic selection system of Karl Sims [12]. He did not change the basic methodology (evolving images created from expressions by aesthetic selection), rather he added a range of additional functions to further increase the gamut of possibilities. Certainly his images looked different and more complex than those of Karl Sims, but they were still of a certain class (images made using an expanded set of mathematical functions).

Indeed, in all uses of aesthetic selection the results produced are ‘of a certain class’, that is they exhibit strong traits of the underlying formalized system that created them (the parameterized system). A natural, but unsuccessful strategy has been to increase the scope and complexity of the parameterized system, giving an even larger gamut of possibilities in the phenotype. Systems of more than trivial complexity cannot be exhaustively searched. In all systems to date, this process is limited by the

---

<sup>2</sup> In one version of Sims’ system, scanned images could form part of an expression tree, allowing ‘real’ images to be manipulated and processed by the system. This does not change the problem discussed here, however.

creativity of the artist or programmer in that they must use their creativity to come up with representations and parameterizations they think will lead to interesting results. The search process has shifted up a level (from parameters to mechanisms), but it is still a search problem that needs to be undertaken by humans: it cannot (yet) be formalized, and hence, automated.

What is needed then is a system capable of introducing novelty *within itself*. The physical entities of the Earth were capable of such a task, in that they were able to create an emergent physical replication system. This was achieved from the bottom up, in a non-teleological process of self-assembly and self-organization. It was possible because atoms, molecules, genes, cells and organisms are all physical entities and part of the same system. Generative systems for EMA could use such a mechanism. This brings us to state the first open problem:

**Open problem #1:** To devise a system where the both the genotype, phenotype and the mechanism that produces phenotype from genotype are capable of automated and robust modification, selection, and hence evolution.

That is, a system that does not produce images of mathematical functions or biomorphs or any particular class of phenotype, due to a fixed parameterized representation. Rather, the genotype, its interpretation mechanism and phenotype exist conceptually as part of a singular system, capable of automated modification. Any such system must be ‘robust’ in the sense that it is tolerant of modification without complete breakdown or failure. A similar challenge has been posed in artificial life research for the evolution of novel behaviors [13].

It might be argued that the phenotypes produced by DNA are ‘of a certain class’ (i.e. biological organisms), however DNA is able to build organisms, which in the appropriate environment are capable of open-ended creative behaviour. These systems exploit dynamical hierarchies to achieve their complexity.

## 2.2 The Problem of Aesthetic Selection

Aesthetic selection of images carried the promise of being able to search for the most beautiful or interesting phenotypes in any parameterized system. In practical terms however, it can only perform a limited search within a certain class of phenotypes, not all possible phenotypes that can be generated by the system. Therefore, the methodology itself tells us little about creativity in general, and does not really offer *the* most beautiful or interesting images from any system.

This limitation of aesthetic selection leads us to ask why it does not achieve its goals and what other methods might be better. Aesthetic selection has several problems:

1. Population size is limited by the ability of people to perform subjective comparisons on large numbers of objects (simultaneously comparing 16 different phenotypes is relatively easy, comparing 10,000 would be significantly more difficult). In the case of visual phenotypes, the available display size may also limit the number and complexity of phenotypes that can be simultaneously shown in order to perform subjective comparison.

2. The subjective comparison process, even for a small number of phenotypes, is slow and forms a bottleneck in the evolutionary process. Human users may take hours to evaluate many successive generations that in an automated system could be performed in a matter of seconds.
3. Genotype-phenotype mappings are often not uniform. That is, a minor change in genotype may produce a radical change in phenotype. Such non-uniformities are particularly common in tree or graph based genotype representations such as in evolutionary programming, where changes to nodes can have a radical effect on the resultant phenotype. This problem is not limited to EMA applications and has been widely studied in the EC community.
4. The size and complexity of genotypes is limited. In general, simple expressions generate simple images. Complex images require more resources to compute and in a real-time system genotypes that consume too much time or space are usually removed before they can complete. In general, it is difficult to distinguish a genotype that takes a long time to do nothing (such as a recursive null-op) and one that takes a long time to do something interesting (this is analogous to the halting problem). Fractal and IFS functions are often found in aesthetic image systems, as they are an easy way of generating complexity in an image with minimal time and space complexity. The problem is that this is not a general complexity, but a fractal one, with characteristic shapes and patterns.

These limitations are indicative of why we can't find the Lisp expression that generates the *Mona Lisa* by aesthetic selection – the human doing the selecting is limiting population size and diversity to such an extent that the genetic algorithm has little chance of finding anything more than local sub-optima. Moreover, the generation scheme, its mapping and complexity, is limited by representation and resources.

Such difficulties have led researchers to try to devise schemes that remove some or all of these limitations while still providing the ability to find interesting phenotypes within the parameterized system's gamut of possibilities. One approach has been to change the interface and selection relationship between user(s) and phenotype [14] rather than removing human subjectivity from the process completely. However, this technique while successful for the situation in which it was devised is not generally applicable to all aesthetic selection problems.

Genotype-Phenotype mapping has also been researched. One interesting approach has been to evolve genotypes that represent some computational process, which is itself generative. That is, the genotype specifies the process of construction and then the construction process builds the phenotype. As the construction process itself is evolvable rather than fixed, more complex outcomes are possible [15].

To address the problems of subjective fitness evaluation by humans, a different approach has been to try to formalize the fitness function, so it can be performed by computer rather than human. This introduces the second open problem:

**Open problem #2:** To devise formalized fitness functions that are capable of measuring human aesthetic properties of phenotypes. These functions must be machine representable and practically computable.

Aesthetics, while well studied in art theory and philosophy, has yet to be fully understood by science. While there have been some noble attempts to measure aesthetic properties, many consider the proposition itself doomed to failure. The mathematician G. D. Birkhoff famously proposed an ‘aesthetic measure’, equal to order divided by complexity. Birkhoff defined ways of measuring order and complexity for several different categories of object, including two-dimensional polygons and vases. While somewhat successful for simple examples, it failed to capture aesthetic qualities with any generality, being described more as a measure of ‘orderliness’ [16].

Neuroscientists have also studied human aesthetic response in order to gain understanding about what makes us consider things beautiful. Ramachandran proposes ‘ten laws of art which cut across cultural boundaries’. These include ‘peak shift’ where exaggerated features exemplify learned classifications, grouping, contrast, isolation, symmetry, repetition, rhythm, balance and metaphor [17].

Birkhoff’s measure and subsequent aesthetic measures of its lineage focus on *measurable features* of aesthetic objects. These are commonly geometric properties, dimension, proportion, fixed feature categories, organizational structure, etc. The basis being that any such feature or property can be objectively measured directly. However, there are many things considered important to aesthetic theory that cannot be measured directly. Such features or properties are generally interpreted rather than measured, often in a context-sensitive way. For example, much has been made of harmonious proportions (such as the golden ratio) in nature, art and music [18]. While such measures are interesting and revealing properties of many different types of structure, they say nothing about the semantics of the structure itself. It not only matters that ancient Greek temples exhibit similar geometric golden ratios, but the context of their form in relation to Greek and human culture, the meaning and significance to the observer, and the perceptual physicality (the interpreted physical relation between observer and observed). It seems that such easily measurable properties are used at the expense of details that are more specific. That is, they are at a too high level of abstraction, where other important features and specific details are ignored. Scientific theories deliberately choose levels of abstraction applicable for physical laws to be ‘universal’. This has been a reasonably successful strategy for the physical universe. For aesthetic laws, however, there are not necessarily such direct abstractions or physical measures.

### 2.3 What Is Art?

In answer to the question ‘what is art?’ Frieder Nake proposed that anything exhibited in art galleries is art [19]. That is, in general terms, people (usually experts) feel that a work has qualities that deem it appropriate to be exhibited in a place recognized for the exhibition and appreciation of art. While there have been many exhibitions of ‘computer generated’ art and even more specifically EMA art, many of these works are primarily selected because they are created by computer, rather than because they are art.

If EMA art is to mature, it needs to become recognized as art for what it is, in addition to how it was made.

**Open problem #3:** To create EMA systems that produce art recognized by humans for its *artistic* contribution (as opposed to any purely technical fetish or fascination).

One might consider this a new version of the Turing test, where artistic outcomes of EMA systems might be compared alongside those done by humans. If the audience cannot tell the difference, or at least considers both worthy of the title ‘art’ then the test has been passed.

The idea of this test does not discount the possibility that EMA might have its own new aesthetic qualities or be part of a wider ‘movement’ in machine-based or generative art.<sup>3</sup> Indeed, Western art is characterized by continuing change and innovation, with movements and styles fluctuating in acceptance and popularity. However, human systems of art theory and appreciation do consider these factors (along with many others) in deciding what is art – so if EMA is really art these systems should be able to accommodate it.

## 2.4 Artificial Creativity

I now turn to the creative activity of artificial systems. As discussed earlier, this differs fundamentally from those systems designed to produce art that is recognized and appreciated by humans. Artificial creativity extends Langton’s idea of artificial life being ‘life-as-it-could-be’. Artificial agent and creature simulations are a common tool in artificial life research. More recently, some researchers have begun to look at creative behavior in artificial systems.

A number of definitions exist for creativity and creative behavior. In developing computational models of creativity, Partridge and Rowe require that creativity involve production of something *novel* and *appropriate* [21]. In addition, novelty may exist relative to the individual (Boden’s P-creativity), and for society or the whole of human culture (H-creativity in Boden’s terminology) [22]. For their computation model of creativity, Partridge and Rowe see novelty involving the creation of new representations through *emergent memory*.

Rob Saunders evolved artificial agents capable of ‘creative’ behavior using a co-evolutionary strategy of creative agents and critics [23]. Agents responded in terms of a psychological theory of interest to novel behavior. Most systems involved in the generation of novelty do so by appropriate recombination of basic primitives. This is Cariani’s *combinatoric emergence* where wholes that are more complex are constructed by combinations of irreducible primitives; the important point is that the total set of primitives and their function are fixed. In the case of *creative emergence* fundamentally new primitives enter the system [24]. Cariani and others have made a case that creative emergence is what we observe in nature. Clearly, this distinction relates to open problem #1, where we want the emergence of new primitives in our system, not just the combination of a fixed set.

**Open problem #4:** To create artificial ecosystems where agents create and recognize their own creativity. The goal of this open problem is to help understand creativity and emergence, to investigate the possibilities of ‘art-as-it-could-be’.

---

<sup>3</sup> For an interesting survey of historical precedents to generative art see [20].

Computational creativity has largely relied on psychological theories of creativity. As neuroscience advances our understanding of creative behavior, this may lead to new models. The challenge for researchers in EMA is to convincingly demonstrate the autonomous emergence of agents capable of generating and recognizing novelty in their interactions.

### 2.5 Theories of Evolutionary Music and Art

Finally, any research involving music or art must be mindful of theories related to such practices from the disciplines themselves. Even studying these theories from an anthropological perspective is likely to shed light on the nature of creativity and aesthetics. Human culture and art is constantly changing and evolving – practices accepted today as art may not have received such acceptance in the past. Evolutionary and generative art is no exception. If this art is to progress, there must be critical theories to contextualize and evaluate it and its practitioners.

**Open problem #5:** To develop *art* theories of evolutionary and generative art.

It is important to distinguish between art theory and art criticism. Art criticism is based on how to evaluate art within some critical framework. Art theory is not like scientific theory in that its use for prediction or general explanation is minimal. There does not seem to be any laws of art that will predict artists' behaviors, or that explain the 'evolution' of art history by detailing what 'succeeds' in making a work beautiful or significant. For the products of EMA to be accepted as art, there must be some artistic theory that is associated with them. Some developments have begun in this area [25].

## 3 Conclusion

This paper has presented five grand challenges for evolutionary music and art research. There is little doubt that these are hard problems and will probably not be solved in the immediate future. However, there is no theoretical reason that prohibits their solution eventually. What is even more profound than the solution of the problems themselves, is the impact their solution will have on society and our understanding of ourselves and our creativity. This is certainly a worthy research agenda.

## References

1. Dissanayake, E.: *Homo Aestheticus: Where Art Comes from and Why*. University of Washington Press, Seattle (1995)
2. Dissanayake, E.: *What Is Art For?* University of Washington Press, Seattle (1988)
3. Pinker, S.: *How the Mind Works*. Penguin Press, Middlesex, England (1997)
4. Miller, G.F.: *The Mating Mind: How Sexual Choice Shaped the Evolution of Human Nature*. William Heinemann, London (2000)
5. Brown, D.E.: *Human Universals*. McGraw-Hill, New York (1991)



6. Miller, G.F.: Evolution of Human Music through Sexual Selection. In Wallin, N.L., Merker, B., Brown, S. (eds.) *The Origins of Music*, MIT Press, Cambridge, MA (2000) 329-360
7. Dartnall, T. (ed.) *Creativity, Cognition, and Knowledge: An Interaction*. Praeger, Westport, Connecticut; London (2002)
8. Bedau, M.A., McCaskill, J.S., Packard, N.H., Rasmussen, S., Adami, C., Green, D., Ikegami, T., Kaneko, K., Ray, T.S.: Open Problems in Artificial Life. *Artificial Life* **6**:4 (2000) 363-376
9. Langton, C.G.: Artificial Life. In Langton, C.G. (ed.) *Artificial Life*, SFI Studies in the Sciences of Complexity, Addison-Wesley (1989) 1-47
10. Bonabeau, E.W., Theraulaz, G.: Why Do We Need Artificial Life? *Artificial Life* **1**:3 (1994) 303-325
11. Sims, K.: Artificial Evolution for Computer Graphics. Proceedings of SIGGRAPH '91 (Las Vegas, Nevada, July 28 - August 2, 1991) In *Computer Graphics* **25**:4 ACM SIGGRAPH, New York, NY (1991) 319-328
12. Rooke, S.: Eons of Genetically Evolved Algorithmic Images. In Bentley, P.J., Corne, D.W. (eds.) *Creative Evolutionary Systems*, Academic Press, London (2002) 339-365
13. Taylor, T.: Creativity in Evolution: Individuals, Interactions, and Environments. In Bentley, P.J., Corne, D.W. (eds.) *Creative Evolutionary Systems*, Academic Press, London (2002) 79-108
14. McCormack, J.: Evolving Sonic Ecosystems. *Kybernetes* **32**:1/2 (2003) 184-202
15. Kitano, H.: Designing Neural Networks Using Genetic Algorithms with Graph Generation System. *Complex Systems* **4**:4 (1990) 461-476
16. Birkhoff, G.D.: *Aesthetic Measure*. Harvard University Press, Cambridge, MA (1933)
17. Ramachandran, V.S.: *The Emerging Mind*. Reith Lectures; 2003. BBC in association with Profile Books, London (2003)
18. Doczi, G.: *The Power of Limits: Proportional Harmonies in Nature, Art and Architecture*. Shambhala (Distributed by Routledge & Kegan Paul), London (1981)
19. Nake, F.: How Far Away Are We from the First Masterpiece of Computer Art? In Brunstein, K., Raubold, E. (eds.) *13th World Congress 94*, Elsevier Science, B.V., North-Holland (1994) 406-413
20. Zelevansky, L.: *Beyond Geometry: Experiments in Form, 1940s-70s*. MIT Press, Cambridge, MA (2004)
21. Partridge, D., Rowe, J.: *Computers and Creativity*. Intellect Books, Oxford, England (1994)
22. Boden, M.A.: What Is Creativity? In Boden, M.A. (ed.) *Dimensions of Creativity*, MIT Press, Cambridge, MA (1994) 75-117
23. Saunders, R., Gero, J.S.: Artificial Creativity: A Synthetic Approach to the Study of Creative Behaviour. In Gero, J.S. (ed.) *Proceedings of the Fifth Conference on Computational and Cognitive Models of Creative Design*, Key Centre of Design Computing and Cognition, Sydney (2001)
24. Cariani, P.: Emergence and Artificial Life. In Langton, C.G., Taylor, C., Farmer, D., Rasmussen, S. (eds.) *Artificial Life II*, SFI Studies in the Sciences of Complexity, Addison-Wesley, Redwood City, CA (1991) 775-797
25. Whitelaw, M.: *Metacreation: Art and Artificial Life*. MIT Press, Cambridge, MA (2004)