Truth to Process – Evolutionary Art and the Aesthetics of Dynamism

Philip Galanter, BA, MFA
Department of Visualization, Texas A&M University, College Station, Texas, USA
e-mail: galanter@viz.tamu.edu

Abstract

After a great deal of initial promise and enthusiasm, evolutionary art seems to have hit a premature and disappointing plateau. This paper proposes a difficult but necessary program for the advancement of evolutionary art.

First two significant problems are discussed. The first is the problem of creating aesthetic fitness functions that would allow evolutionary art systems to execute unattended with industrial sized populations and generations. In this context a quick survey of computational aesthetic evaluation is offered. Then it is suggested that progress in perceptual psychology and neuroaesthetics, coupled with advancements in connectionist computing, may provides new techniques for scoring aesthetic fitness.

The second problem is the sense of "sameness" and lack of innovation exhibited by typical evolutionary art systems. In this regard the related technical problem of genetic representation is noted, various types of genetic representation are reviewed, and a critique is offered relative to evolution in nature. It is noted that evolution in exercises multiple levels of complexification and emergence.

Finally an implication for generative art aesthetics and theory is discussed. It is suggested that current evolutionary art systems and projects are incoherent in so far as they don't focus on the essential virtue of generative art, i.e. a focus on process rather than final form. It is suggested that the general shift from nouns to verbs that is essential to generative art should be pushed to the fore. An aesthetic of truth to process and dynamism is proposed as foundational for evolutionary art.

1. Introduction

Although computational evolutionary art has been an active practice for at least 20 years [1] there is a vague feeling of disappointment with the present state of the art. Industrial applications of evolutionary methods such as genetic algorithms have tracked the general progress seen in other sectors of computation and technology. But evolutionary art seems to have reached a premature plateau. Improvements in the aesthetics of evolutionary art have been slow, and any headway has frequently been by way of narrow techniques that are not generalizable or extensible.

In terms of technology there are two significant challenges, the problem of aesthetic fitness functions and the problem of genetic representation. But there are also related aesthetic and art-theoretical issues that must be faced as well.

2. Fitness Functions for Art

In industrial applications such as automotive design, electronics, routing optimization, investment, chemical process engineering and so on evolutionary techniques have yielded significant value. A common element is that an a priori fitness function can be defined and applied as part of the iterative algorithm. And so fitness scores can be automatically calculated for each candidate solution. [2]

This allows evolutionary systems to run hundreds or thousands of generations with very large populations. And once the system begins it can run for as long as necessary without human intervention.

Artists face a difficult problem in this regard. How does one define and implement an aesthetic fitness function? Various theories proposing a formulaic approach to aesthetic evaluation have come and gone over the years without much success. Some are skeptical as to whether aesthetics can ever be systematized for human understanding let alone computer implementation.

The typical response is to keep the artist "in the loop." The artist manually assigns fitness scores as each new generation is created. This approach is usually called "interactive evolutionary computing" or IEC.

Lewis has cataloged a large number of evolutionary art projects with nearly 200 citations. Most of those systems by far are interactive and use some form of case-by-case human judgment to provide fitness scores. [3] This is sometimes called "the fitness bottleneck" because the artist, usually by orders of magnitude, cannot evaluate populations as quickly as the computer can generate them. An additional problem is that the artist's judgment will become inconsistent due to fatigue and boredom. The result is that compared to industrial applications evolutionary art systems typically suffer from small populations and few generations.

To overcome these fatigue and bottleneck problems Drave's Electric Sheep system accepts evaluation information from the users of the thousands of computers displaying the results as a screen saver. [4] Earlier Sim's Galapagos system used the amount of time an observer would spend looking at a given image as a fitness function. While such systems overcome the fitness bottleneck to some extent, the long term aesthetic utility of this approach is open to question. As both demonstrated and satirized by Komar and Melamid, aesthetics derived from polling leads to the unremarkable and mediocre. [5]

2.1 Computational Aesthetic Evaluation

When it comes to IEC one is reminded of the Mechanical Turk. In the 18th century this was described as a machine that could play chess. [6] In reality the Mechanical

Turk was more like stage magic than computation. Despite making a great show of revealing the see-thru machinery supposedly involved, in reality a human operator was hidden inside the cabinet. This operator made all the decisions and won or lost the game. And in an interactive evolutionary computing system there is a hidden artist making all the decisions and winning or losing the aesthetic game.

From one point of view the challenge of computational aesthetic evaluation (CAE) is to create aesthetic fitness functions that can solve the problems presented by IEC. If human evaluation could be eliminated evolutionary art could benefit from the large populations and multiplicity of generations enjoyed by industrial applications.

So far though, computational aesthetic evaluation remains a grand challenge-level problem. For example, attempts to partially solve this problem have been made by Sanders and Gero [7] and Greenfield [8] using agent-based systems; Jaskowski et al [9], Fornari et al [10], and Ciesielski et al [11] using error measurement relative to an exemplar; McDermott et al [12] similarly using perceptual measures, spectral analysis, and low-level comparison to sound targets; and Khalifa and Foster [13] using music theory-based rules.

Various numeric measures as aesthetic indicators have been explored such as Zipf's law (Manaris et al [14]), fractal dimension (Mori et al [15] and Taylor [16]), and various complexity measures (Birkhoff [17], Machado and Cardoso [18]).

Attempts to use connectionist models such as neural networks in computational aesthetic evaluation include Machado et al [19], Phon-Amnuaisuk et al [20], and Gedeon [21].

Some researchers have tried to deal with the fatigue problem by only using human evaluation for a subset of the population. The scores are then leveraged across the entire population based on similarities in the population. Takagi offers a broad overview of attempts to fuse evolutionary computing and interactive evolution computing. He includes both art systems as well as other applications up to 2001 with over 250 citations. [22] More recent reports on hybrid systems for evolutionary art include Yuan and Gong [23], Pallez, Machado et al [24], and Machwe and Parmee [25].

2.2 Connectionism and the Psychology and Neurology of Aesthetics

Despite the above efforts computational aesthetic evaluation remains an unsolved problem. This shouldn't be terribly surprising because we don't know much about how human aesthetic evaluation works either. And this may be a hint as to where future progress will be found.

Significant research is underway by experimental psychologists as reported in journals such as Psychology of Aesthetics, Creativity, and the Arts, Empirical Studies of the Arts, and the Journal of Consciousness Studies. A picture of how human aesthetics works is being assembled piece by piece from the point of view of experimental psychology.

A new addition to such research is the nascent field of neuroaesthetics. Neuroaesthetics is the scientific study of the neurological bases for the creation, experience, and contemplation of works of art. An example of such work is a recent chapter by Martindale where he proposes a general neural network model of aesthetic perception. In support of this model he cites the compatible results of 25 empirical studies of aesthetic perception. [26]

It's important to note that Martindale's model is theoretical and has not been implemented as a computational neural network. But this kind of work may give others ideas and incentive for doing so. As an example of connectionist computing inspired by neurological theory, Hawkins has introduced a new connectionist design he calls "hierarchical temporal memory" (HTM) based on a theory of the neocortex. [27]

The human brain has about 100 billion (10^{11}) neurons each with about 10 thousand (10^4) synaptic connections. In addition there are about 900 billion glial cells in the brain. At one time these cells were thought to be relatively passive. Current thinking, however, is that glial cells also actively process information. Given there are upwards of 10^{15} connections it seems unlikely digital technology will be able to duplicate the mechanics of the brain any time soon.

In this regard it's worth noting that some evolutionary psychologists have hypothesized that our aesthetic capabilities arise from adaptations related to mate selection. In addition, animals with much simpler neurology seem to select mates based on a simple kind of aesthetic evaluation. It may be that computation orders of magnitude less than that exhibited by the human brain will be up to the task.

So perhaps breakthroughs in both the neurology of aesthetics and connectionist computing will synergistically lead to breakthroughs in computational aesthetic evaluation.

3. Aesthetic Sameness and Genetic Representation

The second major technical problem is that of genetic representation. This problem exhibits two related symptoms in evolutionary art. Frequently work from a given evolutionary systems displays a certain sameness or cast that many find disappointing. And as a corollary a given evolutionary system typically will not shift paradigms or otherwise innovate in significant ways.

To be sure such systems can assist in the discovery of surprises, but crossing genres rarely happens if at all. Innovation at the level of a Picasso transitioning from his Blue Period through the Rose, African, and then Cubist periods is not something we see given the current state of evolutionary art technology. Note that this is <u>not</u> an observation about a lack of quality, although there is that as well, but rather a statement about the inability of evolutionary art systems to significantly break aesthetic paradigms at <u>any</u> level of quality.

In nature complexity at higher scales is an emergent property of self-organization at lower scales. This process has been called by some “complexification.” [28] The problem of aesthetic sameness and lack of innovation is due, at least in part, to the lack of complexification capacity found in most evolutionary art systems.

Artists typically have a vague to specific notion of what the desired result is, and then designs an evolutionary system to explore that aesthetic space. In doing so the design of the genetic representation is of meta-significance because it will constrain the space of *all possible* evolutionary paths.

There are at least four kinds of genetic representation and each allows for a greater or lesser degree of complexification.

Fixed parameters offer the simplest kind of genetic representation. For example, in a system for creating drawings of insects there would be a gene for head size, another for body color, another for leg length, and so on. Such a system will always draw six legs, and so it will never draw a spider. A fixed parametric evolutionary system is highly constrained.

Extensible parameters offer a slightly more complicated representation. In an insect drawing system an arbitrary number of leg genes would be allowed, and so spiders and even centipedes could be drawn. But without wing or fin genes the system will never draw birds or fish. And so an extensible parametric system is still tightly constrained.

Direct mechanical representations are more complicated, but allow for a significant increase in complexification. A version for drawing insects would include genes that construct a (typically virtual) drawing machine as well as genes that instruct the drawing machine as to what marks to make. In theory this kind of genetic representation can allow most any picture to be drawn. And mutations in the genes that construct the drawing machine might, for example, turn pencil-like marks into brushed ink-like marks. But the complexification such a system affords is still limited because the genetically constructed machines only exhibit a single level of emergence.

Reproductive mechanical representations are those most like DNA as found in nature. In such a system genes can construct machines that can go on to construct other machines and so on. And indeed such machines may create copies of themselves, i.e. reproduce. Such a system can induce an arbitrary number of levels of emergence across multiple scales. This kind of genetic representation has the potential to provide complexification as found in nature.

A quick review of complexification in nature is useful at this point. (See figure 3.) According to current theories, shortly after the big bang matter self-organized as it settled into lower energy states resulting first in subatomic particles and then atoms. Hydrogen and helium atoms, in turn, diversified and gained mass through fusion within stars. Self-organization then transitioned from the atomic level to the chemical level with the formation of molecules. Some molecules increased their numbers more rapidly than others due to the appearance of yet other molecules acting as catalysts. This process intensified when feedback loops of reactions that produced their own catalysts appeared. Some of these reactions complexified into networks of self-reinforcing catalytic reactions called “autocatalytic sets.” [29]

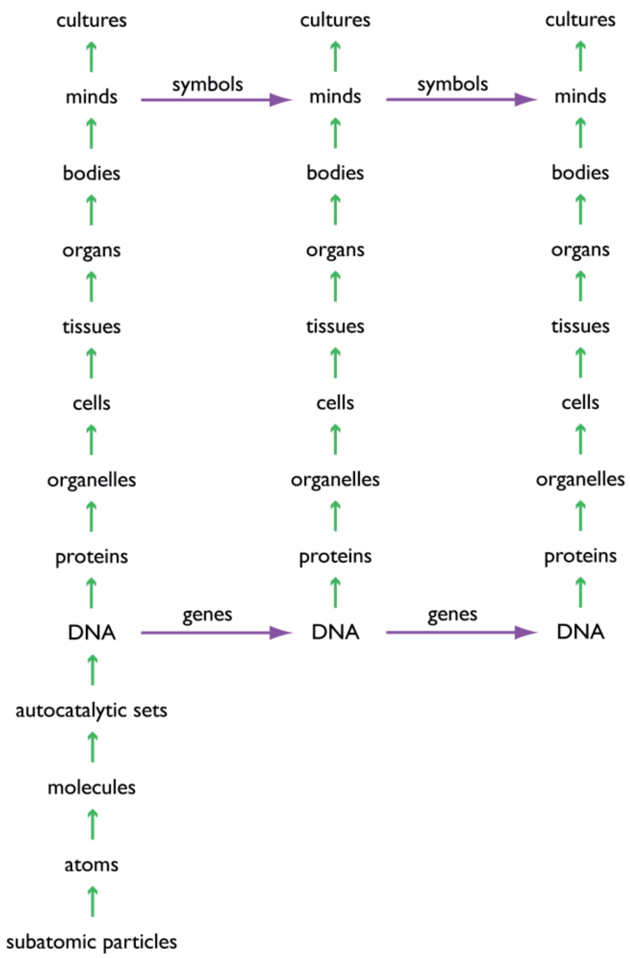

Figure 3

It is hypothesized that out of this sea of autocatalysis arose a primordial form of DNA that had the unique ability to reproduce itself in a way similar to autocatalysis despite the fact that its chemical structure wasn’t fixed. As is now well-known DNA is, in part, made up of sequences of nucleotides that can vary in quantity and order. This created the basis for an extremely robust system of complexification spanning many scales and requiring a number of levels of emergence.

And so there is an upward increase of complexity as DNA creates proteins, proteins organize to create organelles, organelles organize to create cells, cells organize to create tissues, tissues organize to create organs, and organs organize resulting in creatures.

At the level of DNA there is the unique ability to reproduce genes and then kick off a nearly identical sequence of emergence. This results in new living individuals inheriting, with perhaps a little variation, the complexity slowly built up over millions of years.

(It is beyond the scope of this paper to discuss, but it’s not until minds capable of symbolic understanding emerge that another mode of inheritance from individual to individual becomes possible. This has suggested to some that evolutionary art systems should generate emergent individuals capable of further information exchange at a higher level).

Like DNA, an evolutionary art system using reproductive mechanical genetic representations would allow for greater complexity because it would allow multiple levels of emergence across multiple scales. And perhaps then evolutionary art systems could then exhibit the wished for variety and innovation we find in nature.

Unfortunately most evolutionary art systems currently utilize fixed or extensible parametric genetic representations, and a few arguably use direct mechanical representations. The creation of evolutionary art systems that use reproductive mechanical representations remains an unsolved problem as a matter of technology. It also presents an interesting art theoretical problem.

4. Dynamism and Truth to Process

As should be clear from the previous section, evolution in nature is a bottom up process. There are multiple levels of emergence at increasing scales. And this has happened over millions of years due to iterative and immediate competition and survival and not long term planning.

In short evolution is not teleological. It does not set out to create an advanced creature and then project downward to determine what is needed at lower levels.
For all their complexity creatures in the natural world are unplanned and unanticipated emergent properties resulting from many levels of bottom up complexification.

This is the opposite of how evolutionary art is typically created. Generally the artist has some idea of approximately what the end product should be, and then he designs a system and genetic representation that will lead to the anticipated result. This usually involves at most a single level of emergence. Simply as a practical matter this dooms the work to the sameness and systemic lack of innovation noted earlier.

Perhaps this is acceptable and even desired for certain kinds of generative art and design. For example, if one is designing tables its not useful if the system suddenly starts offering lampshades, however surprising and innovative that may be.

But in the realm of fine art where innovation and conceptual focus are the hallmarks of truly great work, one has to ask whether the typical approach is not only practically insufficient but also art-theoretically incoherent.

In previous writing I’ve discussed generative art, and especially complexity-based generative art, as a move from art objects to art processes. What is essential to generative art is not the final object. There are many non-generative ways to create objects too. The defining aspect of generative art is the way the artist cedes control to an autonomous system. [30, 31]

Previous art movements have promoted the notion of “truth to materials.” In the context of formalism it was thought that the most powerful aesthetic would be the presentation of the essential nature of the medium. That would deliver the purest distillation of significant form. Applied to architecture this meant that concrete was presented as concrete, and steel beams were shown as steel beams. For Clement Greenberg paintings as simulated windows into illusory space presented a weak formal aesthetic. Only when the canvas was literally considered as a flat support for paint presented as paint could painting harness its true form and essential power.

In taking a top down rather than bottom up approach, current evolutionary art has turned the process of evolution upside down. And the teleology that doesn’t exist in nature has been introduced in art that is supposed to be inspired by nature. This is part of theoretical incoherence referred to previously.

What is essential to generative art is not any particular material but rather the harnessing of process. At this point in art history a powerful aesthetic for generative art could be called “truth to process.”

Evolutionary art created in the context of truth to process should be created from the bottom up. It should start with reproductive mechanical genetic representations. Gene expression should not directly create the object, but rather should kick off a process of complexification that crosses multiple scales and levels of emergence. This is currently a yet to be implemented lesson from nature.

Most importantly generative art in the spirit of “truth to process” should not obsess on formal issues surrounding the final object. Formal aspects of the final object are important only in so far as they lead back to the processes that created them. Evolutionary art, and generative art in general, should give the audience a sense of dynamism and offer the generative system itself. In the move from nouns to verbs generative art should embrace dynamism, the aesthetic of creation as an activity, and truth to process as intrinsically beautiful.

References

[1] Todd, S. and W. Latham, *Evolutionary art and computers*. 1992, London ; San Diego: Academic Press. 224 p., 32 p. of plates.

[2] Fogel, L.J., *Intelligence through simulated evolution : forty years of evolutionary programming*. Wiley series on intelligent systems. 1999, New York: Wiley. xiii, 162.

[3] Lewis, M., *Evolutionary Visual Art and Design*, in *The art of artificial evolution : a handbook on evolutionary art and music*, J. Romero and P. Machado, Editors. 2008, Springer: Berlin. p. 3-37.

[4] Draves, S., *The electric sheep screen-saver: A case study in aesthetic evolution.* Applications of Evolutionary Computing, Proceedings, 2005. **3449**: p. 458-467.

[5] Ross, A., *Poll stars.* Artforum International, 1995. **33**(5).

[6] Aldiss, B., *The mechanical Turk - The true story of the chess-playing machine that changed the world.* Tls-the Times Literary Supplement, 2002(5170): p. 33-33.

[7] Saunders, R. and J.S. Gero, *Curious agents and situated design evaluations.* Ai Edam-Artificial Intelligence for Engineering Design Analysis and Manufacturing, 2004. **18**(2): p. 153-161.

[8] Greenfield, G., *Robot paintings evolved using simulated robots.* Applications of Evolutionary Computing, Proceedings, 2006. **3907**: p. 611-621.

[9] Jaskowski, W., *Learning and recognition of hand-drawn shapes using generative genetic programming.* Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2007. **4448 LNCS**: p. 281-290.

[10] Fornari, J., *Creating soundscapes using evolutionary spatial control.* Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2007. **4448 LNCS**: p. 517-526.

[11] Ciesielski, V., *Evolution of animated photomosaics.* Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2007. **4448 LNCS**: p. 498-507.

[12] McDermott, J., N.J.L. Griffith, and M. O'Neill, *Toward user-directed evolution of sound synthesis parameters.* Applications of Evolutionary Computing, Proceedings, 2005. **3449**: p. 517-526.

[13] Khalifa, Y. and R. Foster, *A two-stage autonomous evolutionary music composer.* Applications of Evolutionary Computing, Proceedings, 2006. **3907**: p. 717-721.

[14] Manaris, B., et al., *Developing fitness functions for pleasant music: Zipf's law*

and interactive evolution systems. Applications of Evolutionary Computing, Proceedings, 2005. **3449**: p. 498-507.

[15] Mori, T., Y. Endou, and A. Nakayama, *Fractal analysis and aesthetic evaluation of geometrically overlapping patterns.* Textile Research Journal, 1996. **66**(9): p. 581-586.

[16] Taylor, R.P., *Chaos, Fractals, Nature: a new look at Jackson Pollock*. 2006, Eugene, USA: Fractals Research.

[17] Birkhoff, G.D., *Aesthetic measure*. 1933, Cambridge, Mass.,: Harvard University Press. xii, 1 p., 2 l., 3-225, 1 p.

[18] Machado, P., *Computing aesthetics.* Lecture Notes in Artificial Intelligence (Subseries of Lecture Notes in Computer Science), 1998. **1515**.

[19] Machado, P., J. Romero, and B. Manaris, *Experiments in Computational Aesthetics*, in *The art of artificial evolution : a handbook on evolutionary art and music*, J. Romero and P. Machado, Editors. 2008, Springer: Berlin. p. 417-451.

[20] Phon-Amnuaisuk, S., *Evolving music generation with SOM-fitness genetic programming.* Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2007. **4448 LNCS**: p. 557-566.

[21] Gedeon, T. and D. s, *Neural network for modeling esthetic selection.* Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2008. **4985 LNCS**(PART 2): p. 666-674.

[22] Takagi, H., *Interactive evolutionary computation: Fusion of the capabilities of EC optimization and human evaluation.* Proceedings of the Ieee, 2001. **89**(9): p. 1275-1296.

[23] Yuan, J., *Large population size IGAs with individuals' fitness not assigned by user.* Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2008. **5227 LNAI**: p. 267-274.

[24] Machado, P., et al., *Partially Interactive Evolutionary Artists.* New Generation Computing, 2005. **23**(2): p. 143-155.

[25] Machwe, A.T., *Towards an interactive, generative design system: Integrating a 'build and evolve' approach with machine learning for complex freeform design.* Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2007. **4448 LNCS**: p. 449-458.

[26] Martindale, C., *A Neural-Network Theory of Beauty*, in *Evolutionary and neurocognitive approaches to aesthetics, creativity, and the arts*, C. Martindale, P. Locher, and V.M. Petrov, Editors. 2007, Baywood Pub.: Amityville, N.Y. p. 181-194.

[27] Hawkins, J. and S. Blakeslee, *On intelligence*. 1st ed. 2004, New York: Times Books. 261 p.

[28] Casti, J.L., *Complexification : explaining a paradoxical world through the science of surprise*. 1st ed. 1994, New York, NY: HarperCollins. xiv, 320 p.

[29] Kauffman, S.A., *At home in the universe : the search for laws of self-organization and complexity*. 1995, New York: Oxford University Press. viii, 321.

[30] Galanter, P., *What is Generative Art? Complexity theory as a context for art theory*, in *International Conference on Generative Art*. 2003, Generative Design Lab, Milan Polytechnic: Milan, Italy.

[31] Galanter, P., *Complexism and the role of evolutionary art*, in *The art of artificial evolution : a handbook on evolutionary art and music*, J. Romero and P. Machado, Editors. 2008, Springer: Berlin. p. 311-332.