Art and Artificial Life — A Coevolutionary Approach

Gary R. Greenfield
Department of Mathematics and Computer Science
The University of Richmond
Richmond, VA 23173
E-mail: `ggreenfi@richmond.edu`

Abstract

Looking backwards, we recall art of Sommerer and Mignonneau, Sims, and Latham that was inspired by artificial life principles. Assessing current artificial life inspired art, we examine the methods of fitness by aesthetics and user-guided evolution of evolving expressions as practiced by Rooke, Ibrahim, Musgrave, Unemi and the author. Looking forwards, we consider autonomously evolved artistic works using algorithmic aesthetics. We survey what little is known about this topic and proceed to describe our new coevolutionary approach based on hosts and parasites.

Introduction

It would be prohibitive to attempt a comprehensive survey of all the artistic endeavors that have been influenced or inspired by artificial life principles, both for reasons of space and because of the difficulty of referencing many of the works that have been exhibited. But in looking backwards we are struck by two themes: the incorporation of *emergent behaviors* into artistic works, and the exploration of *simulated evolution* for artistic purposes. Emergent behavior forms the cornerstone for many interactive works including, for example, the installations of Sommerer and Mignonneau(39; 40) and Allen(1). Possibly such works trace their origins to the MIT Media Lab *ALIVE* project(24). Spurred on by autonomous robotics, the gaming industry, and the "Furbies" craze, behavior engines and emergent behaviors continue to make their presence felt in the world of fine art. On the other hand, the use of simulated evolution for artistic purposes is perhaps less widely recognized or accepted. We will begin our discussion of this topic by investigating more carefully the origins of simulated evolution in the fine arts.

Michael Tolson, co-founder of the digital effects company *Xaos Tools*, won the prestigious 1993 *Prix Ars Electronica* award for his series of still images titled "Founder's Series." The series was generated with the aid of evolved neural nets. Since Tolson's software was proprietary, details of precisely how this was done are fragmentary. In print, Tolson described his method as applying the genetic algorithm to populations of neural nets to breed intelligent brushes(31). His neural nets were then released onto a specially prepared image where they could sense cues introduced by the artist. By responding to the cues, the image was modified according to the brush *procedures* the neural nets had been bred to implement. As a Siggraph panelist, Tolson showed videotape of the breeding stages of a population of neural nets that were trained to be photosensitive. When patches of pure white were added to an image as cues, the photosensitive neural nets, which were moving at random on the image surface, would streak toward these patches *dragging along the underlying image colors.* Tolson's efforts seem not to have been duplicated[1]. A second example involving neural nets is Lund's *Artificial Painter*(22; 28). It has a much different flavor: Each neural net is coded as a bit string so that the genetic algorithm can be applied, but the computer generated image is obtained by mapping the net's output response at each cell (i.e. pixel) of an environment to a color. Fitness and evolution of the neural nets are accomplished by using the method of Dawkins which we describe next.

In a seminal Artificial Life proceedings paper written by Dawkins(9), the fundamental concept of *user-guided* evolution was introduced. Dawkins implemented, and later marketed commercially, his *Biomorphs* program which allowed users to guide the evolution of a population of two dimensional forms by interactively assigning their fitness values. The forms were recursive Pascal drawing routines whose drawing parameters constituted the form's genotype. Karl Sims combined Dawkin's user-guided evolution technique with an imaging method called evolving expressions for the purpose of providing an artificial life inspired art *medium* in his

[1] Research on intelligent brushes is not easy to characterize. A popular and widely used "intelligent brush" technique, first introduced by Haeberli(17), is based on a drawing program's interactive response to user input through a mouse. Another algorithmically based approach to intelligent brushes is the "computational" brushes recently described by Berzowska(5). Her brushes are Java applets. One intriguing aspect of her work is that her brushes are dynamic and can, for example, erase themselves over time.

Siggraph paper(34), "Artificial evolution for computer graphics." Subsequently, Sims published several variations on this theme(36; 37). For Sims, an image — the *phenotype* — was generated from a LISP expression — the *genotype*. The user, viewing the population of phenotypes, interactively assigned an aesthetic fitness value to the phenotypes so that mating and mutation of the genotypes of the most fit individuals could take place in accordance with the rules of the underlying artificial genetics. During this period, William Latham, with technical assistance from programmer Steven Todd, began exhibiting computer generated synthetic three dimensional organic sculptures also created using the Dawkins' evolutionary paradigm which requires the user to assign fitness based on aesthetics(41). The difference between the two approaches is that Sim's genotypes were expressions implemented as trees, while Latham's were gene sequences implemented as bit strings. The use of evolving expressions is currently widely recognized because of the fundamental role it plays in Koza's optimization method known as Genetic Programming(21).

This takes us to the present. While Sims and Latham have gone on to form companies for developing and promoting their endeavors (*Genetic Arts Inc.* and *Computer Artworks Ltd.* respectively) others have followed in their footsteps. Since Sims' original design was too computationally demanding to be of general use, Sims' successors have proceeded to refine his methods for designing and implementing image generating systems based on evolving expressions and fitness by aesthetics. Most well known are the works of Rooke(18; 45). However the author was an early convert(12; 13; 14; 15) and Unemi(42) made available a very restrictive X windows version of a Sims' style system. Papka et al built a Sims' style system to evolve three dimensional polygonal isosurfaces as an application to help test the immersive CAVE environment(8). McGuire described a Sims' style three dimensional polygonal modeling system(25). Ibrahim invoked Sims' technique to evolve Renderman shaders(19). Musgrave developed a Sims' style prototype for MetaCreations(26). Bedwell and Ebert investigated the possibility of using Sims' method to evolve implicit surfaces(3). Mount has a web page (http://www.cs.cmu.edu/ jmount/g3.html) which offers a Sims' style system based on quaternion maps. It is a successor to earlier systems designed for the web by Witbrock and Neil-Reilly(47). Additional examples are described by Rowbottom(33). Undoubtedly there are many more examples of which we are not aware.

Looking forwards, what does the future hold? One vision for the future is that virtual reality will be populated by artists (populations of image producing agents) and art critics (populations of image consuming agents) who will decide what is to be viewed by the humans visiting the virtual environment. For this to occur, however, some means must be found to make simulated evolution automatic so that the user need n completely guide the evolution. There are three initial forays that have been made into this area. They deserve our careful attention and will serve as our introduction to simulated aesthetics.

Simulated Aesthetics

In 1994, three CMU graduate students — Baluja, Pommerleau, and Jochem — published the results of their efforts to *fully* automate the Sims' process(2). They first designed and implemented a bare-bones Sims' style image generation system and then logged images from users' sessions in order to create a database of images. The images in the database were then numerically rated for their aesthetic value. Images from the database were resolved to 48×48 pixels, and then training and testing sets of images were drawn with equal representation from low, medium, and high ranked images. A neural net was trained and tested, and then asked to guide the interactive evolution without human assistance. The effort expended, together with the number of different experiments the authors performed, is impressive. According to the authors, the results were "somewhat disappointing" and "mixed and very difficult to quantify." It was concluded that it was difficult for the neural nets to learn or discern any aesthetic principles. The authors also noted that the neural net's exploration of image space was "very limited and largely uninteresting." They pointed out that the greatest potential for using such automated approaches may be to prune away uninteresting images and direct the human (assistant) to more promising ones.

A short time later Rooke undertook efforts of a quite different nature to evolve his own art critics which he called"art commentators" to perform aesthetic evaluations of the images his system generated, and which he could then use to guide the image evolution process(32). Rooke's goal was to present his critics with seeded, as opposed to random, starting populations of genotypes. Thus unlike Baluja et al his critics were not forced to start from scratch. Each critic — itself an expression, but one with the capability of examining selected portions of the image phenotype — assigned an aesthetic fitness value to each image in the population. The training set Rooke used consisted of one hundred evolved images together with Rooke's own aesthetic rankings. As populations of expressions, albeit expressions including image processing functions, "tesselation" functions, and statistical measurements, Rooke's critics could also be evolved. Rooke evolved his critics until they could duplicate his fitness rankings to within an acceptable tolerance. To put his critics to work, Rooke gave them his top ranked images from twenty successive generations of an evolutionary run. After each subsequent generation, the oldest of these images would be removed and the image from the current population with the best aesthetic fitness as judged by the critics would be kept. Thus, after twenty generations the critics were in complete control. Rooke let the critics guide the evolution for three hundred generations. Rooke judged his art critics to have been capable of learning *his* aesthetics, but once again, they seemed incapable of using it to explore new areas of image space. One plausible explanation is that Rooke's critics were being caught in eddies of image space. Rooke suggested that it might be necessary to work side by side with his critics and intervene every so often to put them back on track by reassigning aesthetic fitness rankings to the current image set, and then re-evolving the critics — a human assisted coevolution scenario. Following the 1997 Digital Burgess Conference, Rooke and Steve Grand initiated an on-line discussion[2] about the viability of using artificial life coevolution techniques for aesthetics. One idea that emerged was that one needed a "physics" for aesthetics, by which they meant a theoretical framework from which aesthetic principles could be derived and/or tested.

A tangential development which we find significant for understanding the aesthetics of visual images is the recent effort of Belpaeme(4) to evolve expressions whose internal nodes consist solely of image processing primitives in an attempt to discover new and useful digital filters. The aesthetic or fitness value used for the experiment was how successful the filter was at *distinguishing* between the various images of a test set. One intriguing outcome from Belpaeme's experiments was how small the evolved filters turned out to be. One might think that there was a hidden bias towards computational efficiency incorporated into the fitness metric. The explanation Belpaeme offered was that the chaining of image processing functions caused a significant loss of image information content.

Such prior work helps motivate why we think the automation problem for evolving images is a difficult one, and why we see coevolution based on the Sims' method as a significant challenge for the future. Before taking up this challenge, we must review some of the developments of artificial life coevolution research. The first artificial life coevolution simulation was published by Hillis(16), who applied coevolutionary techniques to an optimization problem involving sorting. In the visual arena, Sims gave a stunning example of artificial *learning* based on coevolution(35; 38). Using directed graphs for genotypes, Sims constructed virtual "creatures" to compete in virtual contests of "capture the flag." Sims made mesmerizing videos of the evolved behavior of his creatures. Since then, by linking creature evolution with environments supporting artificial physics, other impressive behaviors have been evolved, including artificial walking and swimming(43).

There are several obstacles to overcome in trying to adapt coevolutionary research to evolving images based on fitness by aesthetics. First, Hillis' sorters are easy to assign a fitness too. Clear optima are recognizable. The sorters either sort or they don't. Similarly his coevolving population of parasites are either successful at invading the sorters by finding examples of difficult lists to sort or they are not. Second, the coevolutionary behavior generators of Sims and his successors seem to depend on competition between individuals within a single population for a resource, or success at completing a task. Further complicating matters is the fact that recently the fundamental nature of coevolution and its underlying principles have begun to be reconsidered. Cliff and Miller point to the difficulty of recognizing and measuring the so-called Red Queen effect that results as two coevolvers change each other's fitness landscapes(7), while Ficici and Pollack question the sustaining power of the coevolutionary arms race by analyzing mediocre stable states and the prevalence of cycling(10).

Now we are able to frame the fundamental problem we are trying to solve: Given a Hillis coevolutionary framework, where now one species consists of host images and another species consists of image parasites, in order for the parasites to prey upon the host images based on aesthetics, how will the parasites be judged? In other words, what will the *algorithmic* assessment of aesthetic fitness be? Nake in commenting on early attempts by Max Bense and Abraham Moles to use the Shannon concept of information as the guiding principle for an analysis of the aesthetic processes concludes:

> Although some exciting insight into the nature of aesthetic processes was gained this way, the attempt failed miserably. Nothing really remains today of their theory that would arouse any interest for other than historical reasons(27).

In a recent issue of the Journal of Consciousness Studies, Ramachandran and Hirstein sparked considerable debate by offering a series of computational aesthetic principles, the primary one being *exaggeration*(29). The framework for exaggeration that they propose is strikingly similar to artificial life sexual selection experiments of Werner and Todd(46). In considering the challenge of evolving populations of images and populations of aesthetic observers, one is heartened by the words of one of the visionaries of artificial life, Thomas S. Ray, who writes,

> We do not know yet, if we can ever expect evolution in the digital medium to express a level of creativity comparable to what we have seen in the organic medium. However, it is likely that evolution can only reach its full creative potential, in any medium,

[2]http://www.biota.org/conf97/reviews.html.

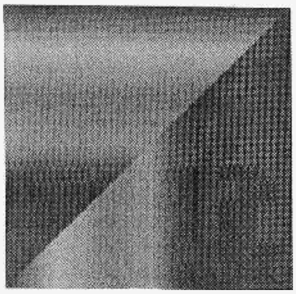

Figure 1: The phenotype from the genotype of a binary basis function (after Maeda) defined on the unit square. Its postfix expression is `V0 V1 B19`.

> when it is free to operate entirely by natural selection, in the context of an ecological community of co-evolving replicators(30).

Inspiring words. In the sequel, we shall take up the coevolutionary challenge based on aesthetic fitness.

Images from Expressions

In this section we will describe how we generate phenotypes (images) from genotypes (expresssions). A genotype is an expression tree E written in postfix form. The leaves of the tree are chosen from a set consisting of constants with values ranging from 0.000 to 0.999 in increments of 0.001 together with the variables $V0$ and $V1$. The internal nodes are chosen from sets of unary and binary primitives or *basis* functions. A unary primitive is a function from the unit interval to itself, and a binary primitive is a function from the unit square to the unit interval. A left to right stack evaluation procedure assigns to each point $(V0, V1)$ of the unit square a value $E(V0, V1)$ in the unit interval which is then mapped to a color. For convenience, we shall resolve phenotypes at a resolution of 100×100. The nodes of the genotype, sometimes referred to as alleles or nucleotides, possess *arity* — technically, the number of arguments each basis function requires — zero for terminals, one for unary basis funcions, two for binary basis functions. A sample binary basis function, which we adapted from Maeda(23), is shown in Fig. 1.

Parasites for Images

Our motivation is as follows: A visually "interesting" image is one that causes our filtering apparatus — our eyes — to generate anomalies for our brain to process. Recent findings of Vogt (44) offer evidence in favor of this hypothesis. We want our images to evolve in such a way that our filtering apparatus will be affected by them. Thus we attach simple digital filters to fixed locations on the image, convolve local portions of the image with the filters, and then compare the convolved image with the original image. We are seeking images for which the convolved image is significantly different from the original. The filter is parasitic upon the image, attempting to blend with the image, while the host image attempts to repel the parasite by making it visible as a blemish. We provide the necessary details.

Given a 100×100 host image with values $h_{i,j}$ in the interval $[0, 1]$, at location L we extract a 10×10 patch $p_{i,j}$ with $1 \leq i, j \leq 10$. A parasite is represented as a 3×3 matrix of integers $(f_{i,j})$ with $0 \leq i, j \leq 2$ whose values are restricted to lie in the interval $[-P_{\max}, P_{\max}]$. The neighborhood of the patch is the 12×12 region of the image consisting of the original patch surrounded by a one pixel wide border. When we pass the filter over the neighborhood we obtain a convolved patch $v_{i,j}$ defined as

$$v_{i,j} = \frac{\sum_{r=0}^{2} \sum_{c=0}^{2} p_{i+(r-1),j+(c-1)} f_{i+(r-1),j+(c-1)}}{S}, \quad (1)$$

where

$$S = 1 + |\sum_{i,j} f_{i,j}|. \quad (2)$$

To make precise the comparison between the original patch and the evolved patch we assign a fitness to the host image via

$$h_{\text{fitness}} = \sum_{i,j} \delta_{i,j}\ , \quad (3)$$

where

$$\delta_{i,j} = \begin{cases} 1 & \text{if } |v_{i,j} - p_{i,j}| > \varepsilon \\ 0 & \text{otherwise}\ , \end{cases} \quad (4)$$

and ε is the host's *exposure* threshold. Since the patch is 10×10 we define the fitness for the parasite to be

$$p_{\text{fitness}} = 100 - h_{\text{fitness}}. \quad (5)$$

When *multiple* parasites are attached to a host, the host's fitness is the average fitness taken over all parasites.

Artificial Genetics for Hosts and Parasites

We use the standard genetic operators for the host expression genotypes: The host crossover operator exchanges subtrees between two host genotypes, and the host point mutation operator causes every node of the host genotype to have a small probability of being substituted for using a different basis function selected from among the set of basis functions of the same arity. We are not aware of any artificial genetics having been previously implemented for (3×3) image filters. Because we felt that filters/parasites should be viewed as exceedingly primitive organisms we did not use any mating operators. Instead reproduction was accomplished by cloning the parasite and subjecting the clone, with small probability, to a small number of transcription operators (e.g., exchange of two rows or columns, shifts of a row or column, or exchange of two entries) before passing it to the parasite point mutation operator which, with small probability, allows each entry of the array to be perturbed.

Coevolution of Evolved Images

Our coevolutionary scenario is now straight forward to describe. During initialization we fix a number of locations for parasites to attach to. We generate a random host population and attach a randomly generated parasite to each of the fixed locations on each host. The parasite populations are managed according to the location they are specific to, analogous to the way a species of fish might have wholly different parasites for specific internal organs. At each time step, fitness updates are calculated and the least fit hosts are removed from the population. Random matings between the survivors are used for replacements. Similarly, for each *location*, the least fit parasites are removed and their replacements are determined by cloning and mutating the most fit survivors from that location's population. A new host inherits the parasites that were attached to the host it is replacing[3]. Since a host's parasites only act by filtering a small patch on the host, and since the phenotype does *not* have to be generated in visual form, the coevolution implementation is fast. Of course to monitor the coevolution we must *cull* the host population and examine the phenotypes. Typically we cull one or two hosts with the highest fitness every two hundred time steps.

Some Results

Since our goal is to obtain visually interesting images, we found it necessary to impose one additional constraint on the genetics of our system. Before describing it, we remark that by having individual parasite populations irritate the hosts locally while the hosts can only react globally *viz., the basis functions used as nodes in the genotype are globally defined*, a tension is set up between these local irritations and the global response. Unfortunately, there are two obvious and uninteresting ways for hosts to fight the local invasion. The first is to evolve thin vertical bands in the phenotype that have many contrasting values so that a parasite can not adjust to the resulting global discontinuities, and the second is to create tiny islands of such discontinuities resembling a cloud of droplets. Examples of such degenerate images are given in Fig. 2. As is so often the case when using the genetic algorithm, our hosts quickly found this flaw in our physics. We countered by restricting the percentage of *unary* primitives that could appear in a genotype. The purpose of using such primitves is to allow visual "smoothing." Since our binary primitives offer much more better visual contrast, forcing our hosts to incorporate sufficiently many of them into their genotypes helped sustain the diversity we needed in the population's gene pool.

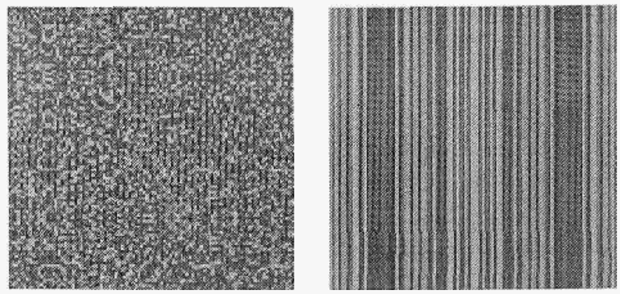

Figure 2: Degenerate images, presumably local optima, that were coevolved by exploiting a flaw in the aesthetic measure of fitness. These cloud or striped hosts present a rugged fitness landscape to the parasites because of the many local discontinuities.

We have tested our coevolutionary simulation using up to thirty hosts and by assigning up to five locations per host giving rise to populations of up to one hundred and fifty parasites. During one representative simulation run using these parameters and lasting 1000 time steps, 10,030 hosts and 79,505 parasites were considered though only ten were culled. No human could examine this many host images. Given our upper bound on the size of the host genotype (up to 100 nodes) and our restriction on the number of unary primitives allowed in the host genotype, it is remarkable that only 24 times during this run was a mating attempt unsuccessful in the ten tries alloted for achieving a valid host crossover.

Does cycling ever occur during our coevolution? Yes, at least to some extent. Hybrids and variants of the degenerate images like the kind shown in Fig 2 do appear and re-appear during the course of coevolution. However, we have visual evidence that, even though we start with random populations, subsequent evolutionary trajectories escape from these uninteresting degeneracies. Our explanation is that since parasites in some sense "chase" their hosts over the fitness landscape, or to put it another way, hosts ward off parasites by fleeing from them, under evolutionary pressure hosts flee parasites following *different* trajectories in image space, hence

[3] Here our biological motivation is weak. While it is easy to argue that clones from the most fit parasites would be the ones most successful at attaching themselves to a location where a host has just successfully repelled a parasite, it would take a stranger scenario to justify re-attaching existing parasites to newly bred hosts. Presumably space limitations would need to dictate that a newly bred host must be "deposited" at a spot being vacated by a host that is to be removed.

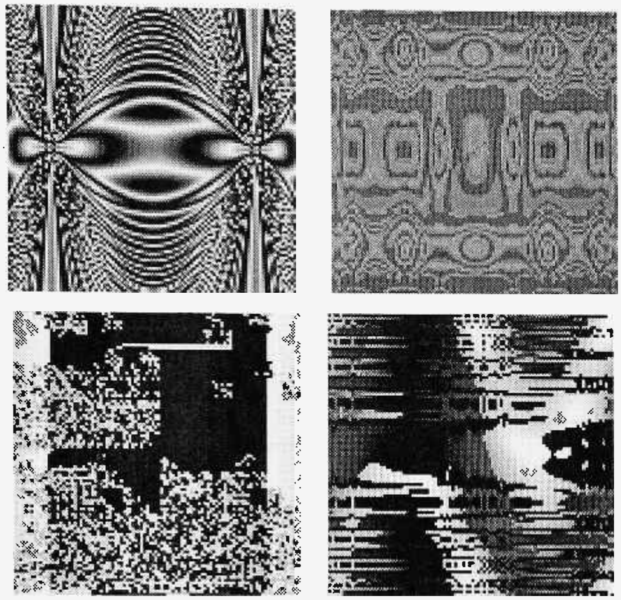

Figure 3: Coevolved images from four different runs. All were evolved starting with small random populations. Noteworthy is the ability of the system to produce diverse imagery by exploring different evolutionary trajectories in image space.

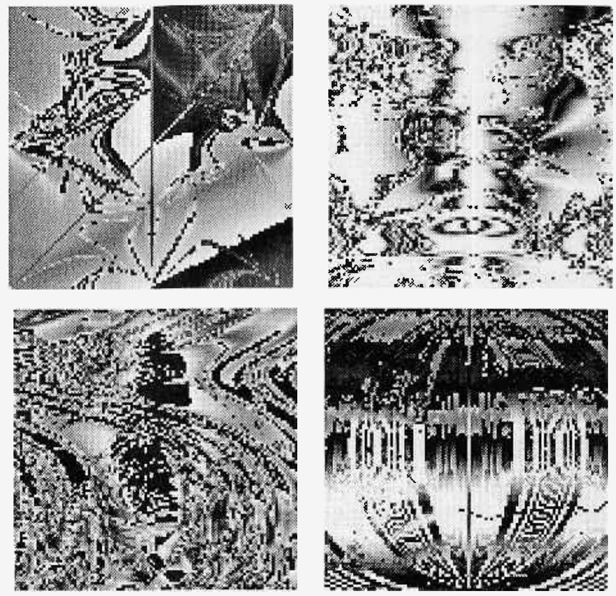

Figure 4: Coevolved images from four runs using larger genome sizes but with the simulation run for only 1500 generations. Host population size is thirty with three parasites per host. The most fit images are culled for inspection every 200 generations.

newly emerging fitter hosts will tend to be quite different than fitter hosts from earlier epochs. (See Fig 3.)

Future Work

Additional benchmarking of the coevolutionary simulation's capabilities is needed. To make reasonable comparisons with images evolved by humans and images evolved from other systems of the type we have described will require considerable additional effort, which is further complicated by the practice of seeding runs from archived gene banks and by the wide disparity of basis function sets that artists and researchers use. Because we did not use seeding, images such as those shown in Fig. 3 are probably best thought of as organisms from the "primordial ooze." More indicative of the image *complexity* that can arise using coevolution is revealed by the representative images shown in Fig. fig4 obtained from runs which permitted hosts to use larger genomes, but restricted the number of generations the simulation was allowed to run for. Work is currently in progress on extending our model so that it is coevolutionary in the sense of Hillis, incorporating a network topology of demes which at intervals share hosts and parasites with neighboring demes(6).

References

Allen, R., Emergence, *Siggraph Conference Abstracts and Applications, Computer Graphics Annual Conference Series*, 1999, ACM Publications, New York, NY, 182.

Baluja, S., D. Pomerleau, and T. Jochem, Towards automated artificial evolution for computer-generated images, *Connection Science*, **6** (1994), 325–354.

Bedwell, E., and D. Ebert, Artificial evolution for implicit surfaces, *Conference Abstracts and Applications, Computer Graphics Annual Conference Series, 1998*, Association for Computing Machinery, New York, NY, 1998, 261.

Belpaeme, T., Evolving visual feature detectors, *Advances in Artificial Life, Proceedings of the Fifth European Conference*, 1999, D. Floreano et al (eds.), Lecture Notes in Artificial Intelligence, Vol. 1674, Springer Verlag, 266–270.

Berzowska, J., Computational expressionism: a model for drawing with computation, *Siggraph Conference Abstracts and Applications, Computer Graphics Annual Conference Series*, 1999, ACM Publications, New York, NY, 210.

Clement, D., Coevolution of evolved expressions, *Honor's Thesis*, University of Richmond, Richmond, VA, 2000.

Cliff, D., and G. Miller, Tracking the Red Queen: measurements of adaptive progress in co-evolutionary simulations, *Advances in Artificial Life: Third European Conference on Artificial Life, Granada Spain, June 1995, Proceedings*, Lecture Notes in Artificial Intelligence 929, Springer-Verlag, Berlin, 1995, 200–218.

Das, S., T. Franguiadakis, M. Papka, T. DeFanti, and D. Sandin, A genetic programming application in virtual reality, *IEEE World Congress on Computational Intelligence, Proceedings of the First IEEE Conference on Evolutionary Computation June 1994*, IEEE Press, 1994, 480–484.

Dawkins, R., The evolution of evolvability, *Artificial Life*, C. Langton (ed.), Addison-Wesley, Reading, MA, 1989, 201–220.

Ficici, S., and J. Pollack, Challenges in co-evolutionary learning; arms-race dynamics, open-endedness, and mediocre stable states, *Artificial Life VI*, C. Adami et al (eds.), MIT Press, Cambridge, MA, 1998, 238–247.

Graf, J., and W. Banzhaf, Interactive evolution of images, *Genetic Programming IV: Proceedings of the Fourth Annual Conference on Evolutionary Programming*, J. McDonnell et al (eds.), MIT Press, 1995, 53–65.

Greenfield, G., An algorithmic palette tool, University of Richmond Technical Report TR-94-02, 1994.

Greenfield, G., Evolving expressions and art by choice, *Leoonardo*, Vol 33. No. 2, 2000, to appear.

Greenfield, G., New directions for evolving expressions, *1998 Bridges Conference Proceedings*, R. Sarhangi (ed.), Gilliland Publishing, Arkansas City, KS, 1998, 29–36.

Greenfield, G., On understanding the search problem for image spaces, *1999 Bridges Conference Proceedings*, R. Sarhangi (ed.), Gilliland Publishing, White Plains, MD, 1999, 41–54.

Hillis, D., Co-evolving parasites improves simulated evolution as an optimization procedure, *Artificial Life II*, C. Langton et al (eds.), Addison-Wesley, Reading, MA, 1991, 313–324.

Haeberli, P., Paint by numbers : abstract image representations, *Computer Graphics (Proceedings SIGGRAPH '90)*, 1990, 207–214.

Hitchcock, N., Painting pictures through time, *Computer Artist*, December/January 1996, 9–9.

Ibrahim, A., GenShade, *Ph.D. Dissertation*, Texas A&M University, 1998.

Keith, M., and M. Martin, Genetic programming in C++: implementation issues, *Advances in Genetic Programming*, K. Kinnear, Jr. (ed.), MIT Press, Cambridge, MA, 1994, 285–310.

Koza, J., *Genetic Programming III : Darwinian Invention and Problem Solving*, Morgan Kaufmann, San Francisco, CA, 1999.

Lund, H., L. Pagliarini, O. Miglino, Artificial painter, *Abstract Book of the Third European Conference on Artificial Life (Ecal '95)*, 1995.

Maeda, J., *Design by Numbers*, MIT Press, Cambridge, MA, 1999.

Maes, P., ALIVE : Artificial Life Interactive Video Environment, in *Computer Graphics Visual Proceedings, Annual Conference Series, ACM Siggraph 1993*, ACM, New York, NY, 1993, 189–190.

McGuire, F., The origins of sculpture: evolutionary 3D design, *IEEE Computer Graphics & Applications*, January, 1993, 9–12.

Musgrave, K., http://www.wizardnet.com/musgrave/mutatis.html.

Nake, F., Art in the time of the artificial, *Leonardo*, Vol. 31, No. 3, 1998, 163–164.

Pagliarini, L., H. Lund, O. Miglino, and D. Parisi, Artificial life: a new way to build educational and therapeutic games, *Artificial Life V*, C. Langton and K. Shimohara (eds.), MIT Press, Cambridge, MA, 1997, 152–156.

Ramachandran, V., and W. Hirstein, The science of art: a neurological theory of aesthetic experience, *Journal of Consciousness Studies*, Vol. 6, No. 6-7, June/July 1999, 15–52.

Ray, T., Evolution as artist, in *Art@Science*, C. Sommerer & L. Mignonneau (eds.), Springer-Verlag/Wien, 1998, 81–91.

Robertson, B., Computer artist Michael Tolson, *Computer Artist*, August/September 1993, 20–23.

Rooke, S., personal communication.

Rowbottom, A., Evolutionary art and form, in *Evolutionary Design by Computers* (ed. P. Bentley), Morgan Kaufmann Publishers, San Francisco, CA, 1999, 261–277.

Sims, K., Artificial evolution for computer graphics, *Computer Graphics*, **25** (1991), 319–328.

Sims, K., Evolving 3D morphology and behavior by competition, *Artificial Life IV*, R. Books and P. Maes (ed.), MIT Press, Cambridge, MA, 1994, 40–48,

Sims, K., Interactive evolution of dynamical systems, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, MIT Press, 1991, 171–178.

Sims, K., Interactive evolution of equations for procedural models, *The Visual Computer*, **9** (1993), 466–476.

Sims, K., Virtual creatures, *Computer Graphics, Annual Conference Proceedings, 1994*, ACM SIGGRAPH, 1994, 15–23.

Sommerer, C., and A. Mignonneau, Art as a living system, in *Art @ Science*, C. Sommerer & A. Mignonneau (eds), Springer-Verlag/Wien, 1998, 148–161.

Sommerer, C., and A. Mignonneau, Art as a living system : interactive computer art works, *Leonardo*, **32** No. 3 (1999), 165–173.

Todd, S., and W. Latham, *Evolutionary Art and Computers*, Academic Press, San Diego, CA 1992.

Unemi, T., *sbart* - Interactive Art Using Artificial Evolution, User's Manual, 1994.

Ventrella, J., Attractiveness vs. efficiency - how mate preference affects locomotion in the evolution of artificial swimming organisms, *Artificial Life VI*, C. Adami et al (eds.), MIT Press, Cambridge, MA, 1998, 178–186.

Vogt, S., Looking at paintings: patterns of eye movement in artistically naive and sophisticated subjects, Extended Abstract, *Leonardo*, Vol. 32, No. 4, 1999, 325.

Voss, D., Sex is best, *WIRED*, December, 1995, 156–157.

Werner, G., and P. Todd, Too many love songs: sexual selection and the evolution of communication, in *Fourth European Conference on Artificial Life*, P. Husbands & I. Harvey (eds.), The MIT Press, Cambridge, MA, 1997, 434–443.

Witbrock, M., and S. Neil-Reilly, Evolving genetic art, in *Evolutionary Design by Computers*, P. Bentley (ed.), Morgan Kaufmann Publishers, San Francisco, CA, 1999, 251–259.