On Evolutionary Computer-Generated Art

PENOUSAL MACHADO
JUAN ROMERO

Many people have a distorted view of Artificial Intelligence (AI). As Astro Teller observes, for most people "Artificial Intelligence is the science of how to get machines to do the things they do in the movies."[1] AI researchers often have to deal with misconceptions created by the doomsday scenarios portrayed in movies such as the 1983 classic *War Games* or the *Terminator* series. These movies provoke fear. Others present a romantic view of AI. For instance, the 1984 movie *Electric Dreams* is about a boy who bought a computer and accidentally spilled a beverage on it. As a consequence, the computer developed AI. Listening to the girl next door practicing the cello, it learned how to play and compose music, fell in love with the girl, tried to kill its owner, and ended up committing suicide.

In reality, since its inception, AI research has given considerable emphasis to logic, reasoning, problem solving, planning, natural language processing, expert systems, chess, et cetera. In several tasks requiring intelligence, state-of-the-art AI systems are now able to attain human competitive results or even to surpass human performance. Yet, in tasks requiring creative reasoning, such as art, design, music, and poetry, computers are far from reaching the accomplishments of humans. As Dissanayake observes, art-making activities are ubiquitous, have evolutionary value, and are a part of human behavior since prehistory.[2] Creativity is often regarded as one of the most remarkable characteristics of the human mind. Not surprising, then, that the search for computational "creativity" should be a central aspect of AI.

During the initial years of AI research, the main source of inspiration was human intelligence. Over the years, researchers have realized that many other sources of inspiration can be used. Establishing analogies with physics phenomena gave rise to novel search methods such as Simulated Annealing,[3] Hopfield Networks,[4] and Elastic Networks.[5] Currently, there is a growing interest in bio-inspired computing, an area of research that comprises techniques such as Evolutionary Computation (EC), Swarm Intelligence, Ant Colony Optimization, and Artificial Life.

Through time, evolution has created a wide variety of species adapted to their environment. Some of these species—for example, humans—exhibit intelligent behavior. Since evolution is the source for natural intelligence, it

SP_TER2_21_156-170.indd 156 1/17/11 11:43 AM

naturally became a source for AI. According to Darwin, evolution is based on two fundamental principles: selection, and reproduction with variation.[6] Selection ensures that fitter individuals are more likely to reproduce. The descendants of these individuals inherit characteristics from the progenitors—which implies that they tend to be fit—but they are not exact copies, which allows evolution. The reinterpretation of Darwin's ideas in light of Mendel's genetics, that is, neo-Darwinism, explains how characteristics are inherited and how and why changes occur. Natural selection occurs at the phenotypic level, while reproduction acts on the genotype.[7] The characteristics of the individuals are not directly inherited. Instead, the genes that codify these characteristics and those that enabled their development are inherited. Variation results from copying errors—that is, mutations—and from the recombination of the genetic material of the progenitors.

The goal of EC research can be synthesized as follows: "How do we turn Darwin's ideas into algorithms?"[8] Nowadays, EC comprises a wide and growing variety of stochastic algorithms. In spite of this variety, they have the same main characteristics, so they can all be seen as instances of a generic evolutionary algorithm (see algorithm 1).

Algorithm 1 Generic Evolutionary Algorithm

P ← generate-initial-population ()
while termination criteria not met **do** *evaluate* (*P*)
P' ← select-individuals (*P*(*t*))
P" ← apply-genetic-operators (*P'*)
P ← create-next-population (*P*,*P"*)
end-while
return simulation result

To illustrate how such an algorithm would work, we'll look at the Four-Color Map problem, which consists in coloring a map using a maximum of four colors in a way that ensures that no country has a neighbor with the same color.[9] Using EC to solve this problem implies a series of analogies. Each individual is a *candidate solution* to the problem. In other words, each individual is a colored map. The fitness of an individual is proportional to the quality of the solution. Thus, a colored map where no neighbors share the same color has maximum fitness, while map-colorings that violate this constraint will be penalized proportionally to the number of times they violate it.

To apply EC, one must find an adequate representation. For instance, one can consider that the genotype of each individual is composed of a chromosome, represented by a string, composed of *n* genes, where *n* is the number of countries in the map. Genes can assume four values, representing the color associated to the corresponding country.

Once the representation is chosen, one must also define adequate genetic operators. When two progenitors reproduce there is a probability, for example, 70 percent, of recombination of their genetic code. When recombination does not occur, the genetic code of the descendants is a copy of the progenitors' code. For the purpose of recombination one could use 1-point crossover: a number (*x*) between 1 and n is randomly selected, and two descendants are generated by copying the first *x* genes from the first parent and the remaining genes from the second (and vice versa for the second individual). After this stage, the mutation operations take place. Each gene of the descendants has a probability of suffering a mutation, for example, 1 percent. Mutation can be implemented as follows: the value of that gene is replaced by a randomly chosen number between 0 and 3.

One must also define a selection scheme. Tournament-based selection is appropriate for this problem, that is, to choose a progenitor, one starts by randomly selecting a given number of individuals (e.g., five) among the population, and the fittest of these individuals will be the progenitor. An individual may be selected more than once for reproduction.

Finally, one must define an initialization method, a replacement scheme, and a termination criterion. For the sake of simplicity one can assume that an initial population of a particular size (e.g., one hundred) individuals is created by randomly choosing values for their genes; non-elitist generational replacement, that is, each population of individuals is entirely replaced by its descendants; and the simulation will stop when the map-coloring with maximum fitness is found.

Considering these choices, the algorithm would proceed as follows. An initial population of one hundred randomly generated individuals is created. Each of the population's individuals is evaluated and its fitness determined. Fifty pairs of individuals are selected as progenitors by using tournament selection; individuals may be selected more than once. These generate one hundred descendants by the application of the genetic operators described earlier. The descendants replace the progenitors, thus becoming the current population. The cycle is repeated, from the evaluation step, until a map where no neighboring countries share the same color is found.

Currently the most prominent EC approaches are Genetic Algorithms (GAs), invented by Holland, and Genetic Programming, popularized by Koza [10] The most significant differences between these approaches concern representation and genetic operators. In a GA, the genome codifies a set of parameters or characteristics necessary to build the phenotype and is typically represented by a string. Thus, the described map-coloring approach is a GA. In GP the genotypes are programs—typically represented by a tree-like structure—the execution of which results in the phenotype.

To a large extent, the appeal of EC is its independence from problem-specific knowledge. Thus, one does not need to know how to solve a problem; the only requirements for applying EC are finding an adequate encoding and a way to assign fitness. Likewise, the ability to use EC in an artistic domain depends on finding an adequate way to represent and evaluate artistic objects.

In his 1986 book *The Blind Watchmaker*, Richard Dawkins delineates a program that allows the evolution of the morphology of "virtual creatures" or *biomorphs*.[11] More precisely, each biomorph is a drawing, the appearance of which depends on the values of a set of parameters encoded in a string, the genotype. The biomorphs of the current population are displayed on the screen, and the user indicates his/her favorite ones. In other words, the user guides the GA, which circumvents the need to develop a computational fitness function. This influential work led to the emergence of a new research area, evolutionary art.[12] Due to the subjectivity inherent in artistic production, and the subsequent difficulty of creating an algorithm capable of assigning fitness to an artwork, most evolutionary art systems are also guided by the user.

Draves' *Electric Sheep*—the name pays tribute to Philip K. Dick's novel *Do Androids Dream of Electric Sheep?*—is one of the most notable evolutionary art projects.[13] Like Dawkins' biomorphs, *Electric Sheep* features a user-guided parametric evolution model. In this case, the individuals are "fractal flames," a particular kind of fractal, invented by Draves, that belongs to the so-called Iterated Function System category of fractals.[14] The genotypes consist of several hundreds of floating point numbers that encode parameters for the fractal formula, controlling the scattering of the billions of particles that compose each image. *Electric Sheep* is a distributed computing project currently involving more than 350,000 users. Acting as a screensaver, it takes advantage of the computer's idle time to render the individuals, that is, "sheep," that are being evolved collectively. Interested users can vote on their favorite sheep, thus shaping the course of evolution.

In parametric evolution models, the genetic code is a visual language defined by the designer of the system. Therefore, "creating a parametric model implicitly creates a set of possible designs or a solution space."[15] As such, these systems tend to have an identifiable system *signature* that is closely related to the choices made by the human designer. The model should be compact, that is, genotypes should be relatively small; expressive, meaning that it should allow a wide variety of shapes; and robust, in the sense that interesting images should be easy to find.[16] Parametric evolution has been applied to a wide variety of domains including the evolution of cartoon faces,[17] fonts,[18] line drawings,[19] surfaces,[20] and consumer product design.[21]

The seminal work of Sims gave rise to another popular evolutionary art approach: expression-based evolution.[22] In addition to Sims' work, notable examples of this technique include works by Latham, Rooke, and Hart.[23] (Hart's work is on the cover of *TER*.) We shall use NEvAr (Neuro Evolutionary Art) to illustrate this approach. Largely inspired by the work of Sims, it allows the evolution of populations of images, using GP as in Sim's work. Each genotype is a program—in this case, a symbolic expression represented as a tree. These programs are constructed from a lexicon

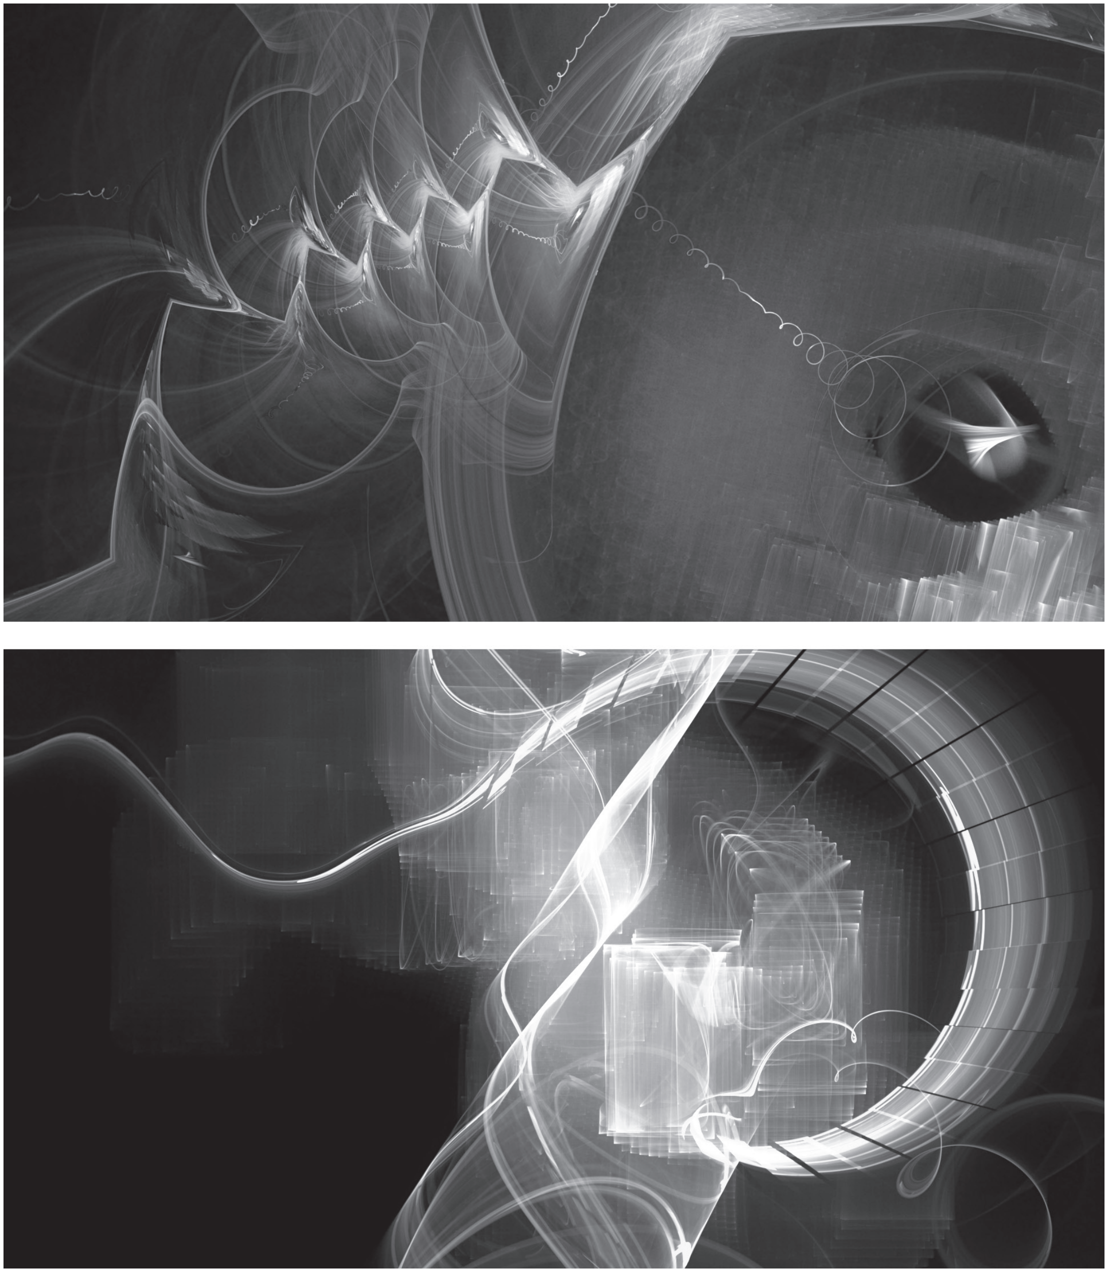

Figure 1. Still images from *Generation 243* by Scott Draves and the Electric Sheep (2009), commissioned by Carnegie Mellon University for the Gates Center of Computer Science

of functions and terminals. The function set is composed mainly of simple functions such as arithmetic, trigonometric, and logic operations. The terminal set is composed of two variables, *x* and *y*, and some random constants. The phenotype (image) is generated by evaluating the genotype for each (*x*, *y*) pair belonging to the image. Thus, the images generated by NEvAr are graphical portrayals of symbolic expressions (see figure 2).

In order to produce color images, NEvAr resorts to a special kind of terminal that returns a different value depending on the color channel being processed. Recombination is performed using the standard GP crossover operator, which exchanges sub-trees between individuals.[24] Five mutation operators are used: sub-tree swap, sub-tree replacement, node insertion, node deletion, and node mutation.[25]

Like most, if not all, evolutionary art systems, NEvAr has a signature—in the sense that it is more prone to generate certain types of images than others—that is closely related with the function set, genetic operators, and genotype-phenotype mapping used. Nevertheless, as demonstrated by Machado and Cardoso, it is theoretically possible to generate any image with NEvAr, and the same is true for several other expression-based evolutionary art systems.[26] This means that "Every great (and not so great) work of visual art is in there, past, present and future, as are images of political assassinations, nude celebrities (even ones that have never posed nude), serial killers, animals, plants, landscapes, buildings, every possible angle and perspective of our planet at every possible scale and all the other planets, stars, galaxies in the universe, both real and imaginary. Pictures of next week's winning lottery ticket, and of you holding that winning ticket."[27]

In practice, the images tend to be abstract and have a computer-generated appearance. Nevertheless, a patient and disciplined user is able to guide evolution from an initial, randomly generated population, to populations of images that meet the users' preferences.

Recently, Graça and Machado have developed "Evolving Assemblages," a system that evolves large-format reproductions of input images by assembling 3-D objects, using GP and interactive evolution.[28] In this case, each individual is a program that receives an image as input and generates an assemblage as output. To accomplish this, each genotype comprises five trees. Based on the input image and for each of its pixels, the first tree selects, from a list of available objects, which type of object will be placed on the canvas; the second tree determines the rotation applied to each object; the third one determines the size of each object; the fourth one determines the *x* coordinate where the object will be placed; and, finally, the fifth one determines the *y* coordinate. Once the assemblage is calculated, the color of objects is determined: each object assumes the color of the pixel of the input image where its center is placed.

GP approaches have also been used in domains such as the evolution of painterly renderings,[29] line-based drawings,[30] plant-like shapes and other 3-D objects,[31] l-systems (see McCormack for a survey[32]), filters,[33] animations,[34] and architectural plans.[35]

Outside the field of the visual arts, evolutionary approaches based on GAs or GP have been applied to sound synthesis, improvisation, harmonization and composition (see Miranda and Biles for a survey[36]); poetry generation[37]; choreography;[38] and many other fields. (See figures 3, 4, and 5.)

Although interactive evolution techniques have been used to produce a wide variety of artistic artifacts, this presents several problems. The most pressing problem is the user fatigue caused by the need to evaluate a large number of individuals.[40] Other problems include the subjectivity of the task, the lack of consistency in the users' evaluations, and the bias toward novelty. All these factors have a negative effect on the evolutionary process. Moreover, interactive evolution techniques skirt an important AI objective: building a computational system capable of performing aesthetic judgments, even if limited ones.

In general terms, there are three main approaches for the automation of the fitness assignment step of evolutionary art systems: using handwritten fitness functions; using machine learning techniques and using co-evolutionary approaches.

Machado and Cardoso took inspiration from the works of Arnheim as well as from research that points toward a preference for simple representations of the world, and a tendency to perceive it in terms of regular, symmetric

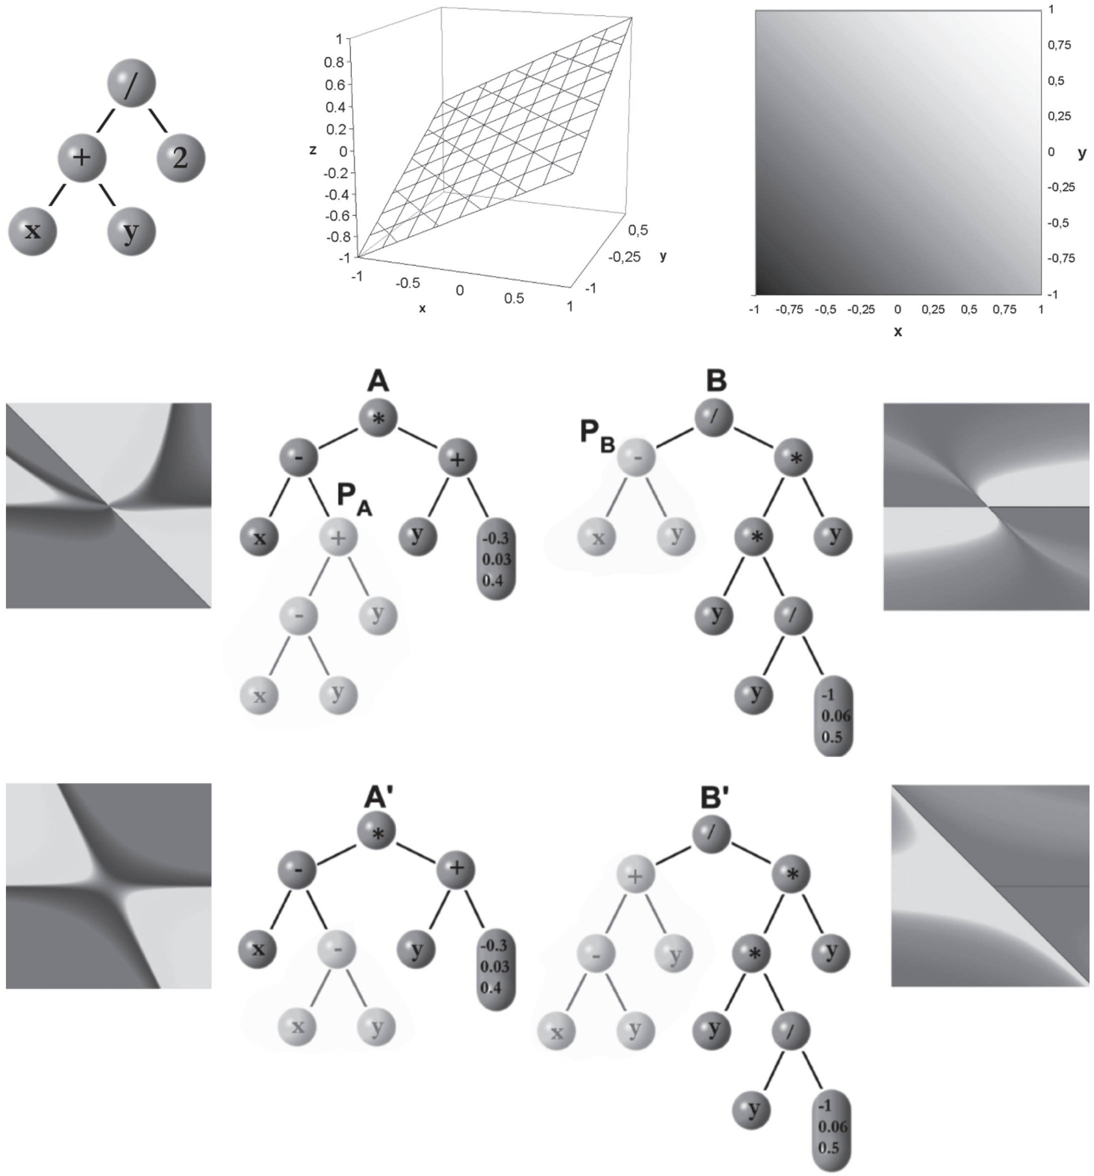

Figure 2. *In the top row*: function $f(x, y) = (x + y) / 2$ represented as a tree (the format of the genotype); as a 3-D chart and as an image produced by assigning to each pixel a luminance proportional to the height of the corresponding position of the 3-D chart. *In the second row*: two individuals and their corresponding genotypes. *In the third row*: the descendants produced by performing a crossover operator at points P_A and P_B and swapping the corresponding sub-trees (depicted in lighter gray).

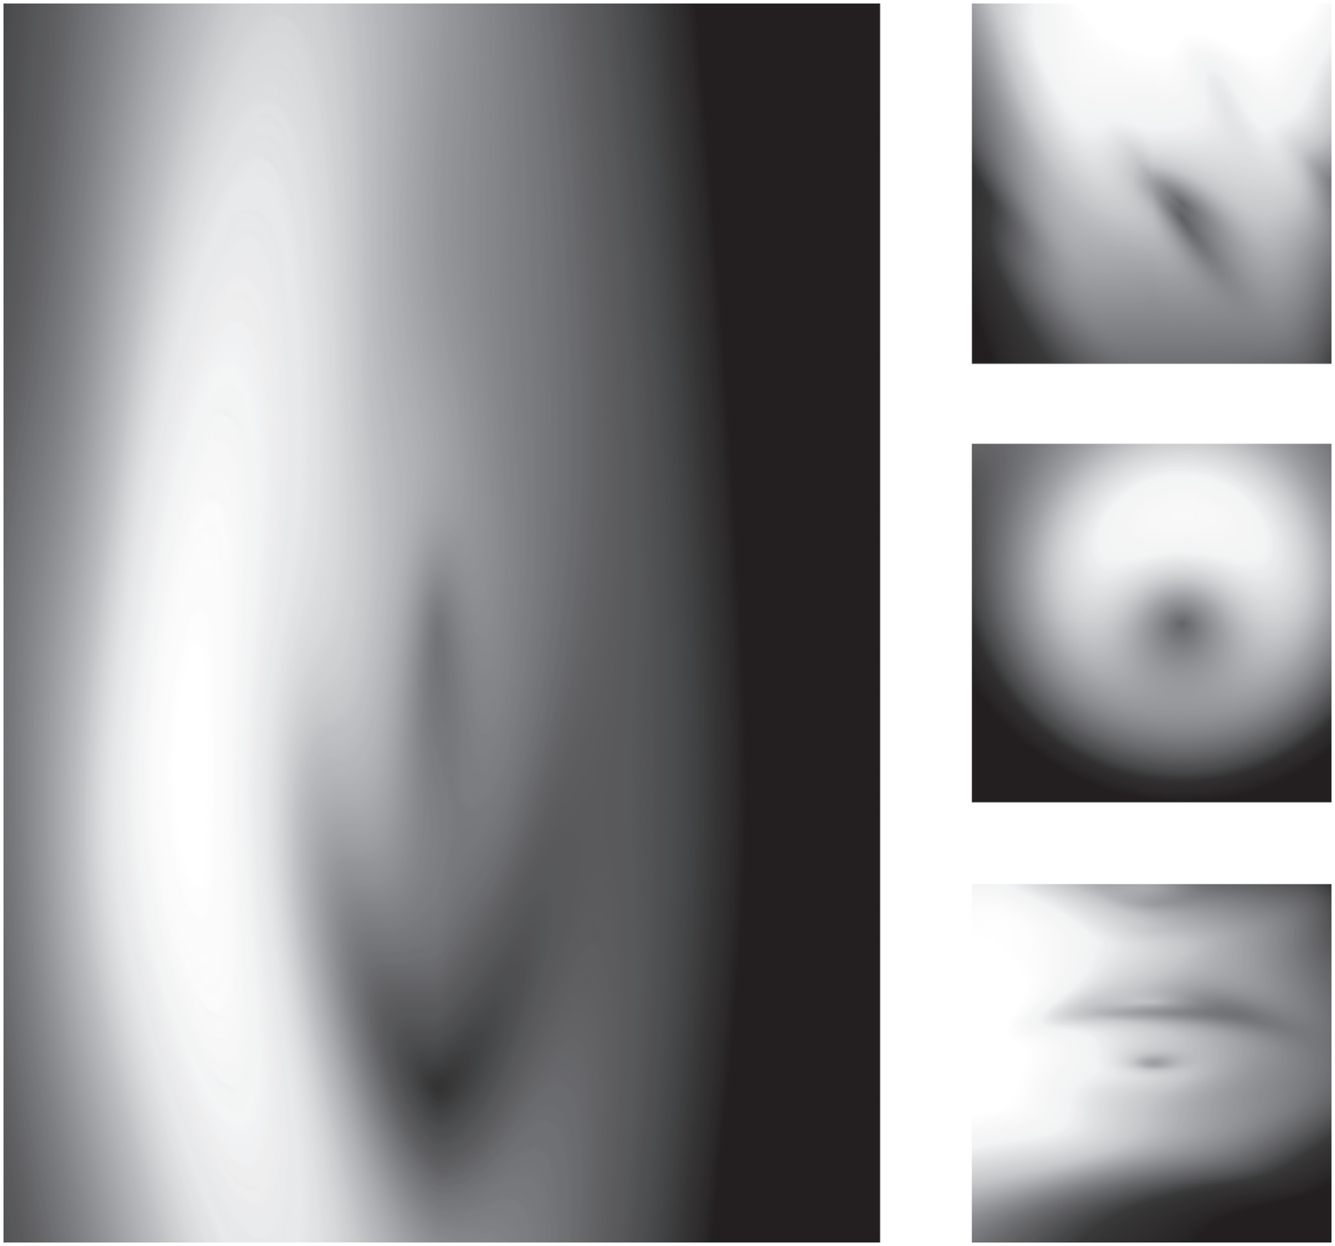

Figure 3. Images from the series *bodies of sin* created with NEvAr using interactive evolution. These images are, arguably, the first set of pictorial images created using expression-based evolution and were first displayed at EvoMUSART'2005, Lausanne, Switzerland.

Figure 4. Assemblages evolved by Graça and Machado in the scope of the *Evolving Assemblages* project, which won the 2010 Genetic and Evolutionary Computation Conference Evolutionary Art Competition.

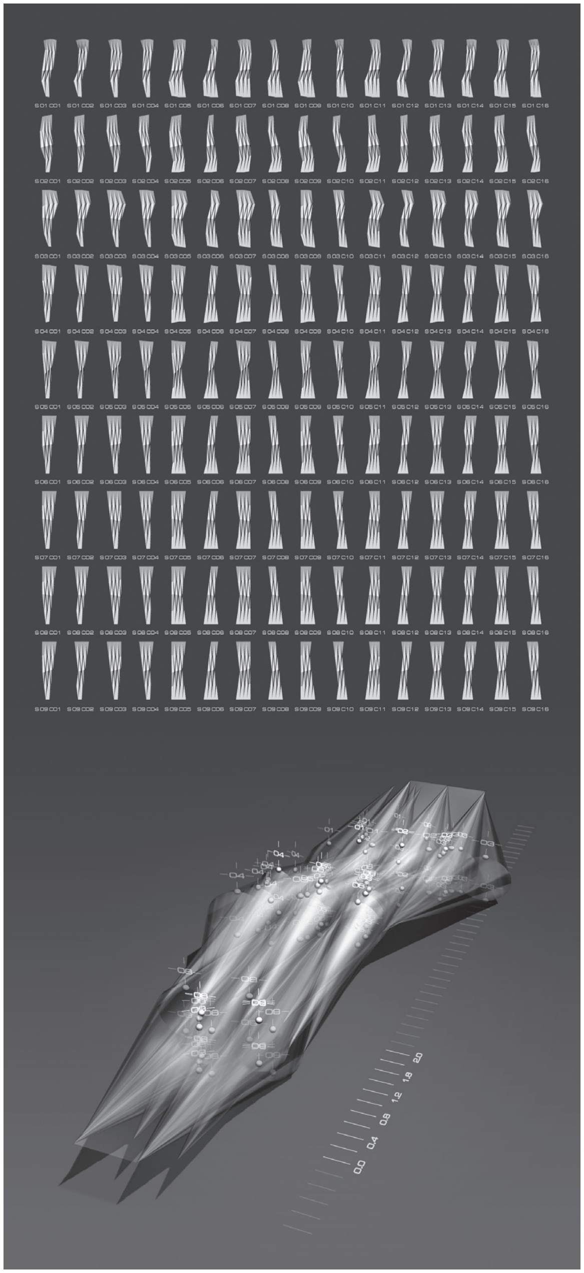

Figure 5. Morphogenetic design experiment—*AA Strawberry Bar*, 2003, Achim Menges, using Genr8.[39]

and constant shapes.[41] As a consequence, they have explored the working hypothesis that aesthetic value has a link with the sensorial and intellectual pleasure experienced in finding a compact percept (internal representation) of a complex visual stimulus. The identification of symmetry, repetition, rhythm, balance, et cetera, can be a way of reducing the complexity of the percept, which may explain these aesthetic principles and the ability of the brain to recognize them in an "effortless" way.

Following this set of ideas, in "Computing Aesthetics," we propose an aesthetic theory: those images that produce a complex visual stimulus and yet result in a compact percept—that is, a compact internal representation—tend to be valued.[42] For instance, fractal images are usually complex, and highly detailed; yet they can be compactly described by a simple mathematical formula. The self-similarity of these images can make them easier for our brain to process, allowing one to build a compact percept, which would explain, according to this theory, why fractal images tend to be aesthetically interesting.

To test this theory, we used JPEG- and quad-tree–based fractal image compression to estimate the complexity of the visual stimulus and the percept.[43] In "All the Truth," we used a variation of the proposed formula to assign the fitness, thus making NEvAr autonomous.[44] In recent years, several evolutionary art systems that use complexity estimates to assign fitness have been proposed. Neufeld, Ross, and Ralph present a genetic programming engine generating non-photorealistic filters by means of a fitness function based on Ralph's bell curve distribution of color gradient.[45] The model was implemented by doing an empirical evaluation of hundreds of artworks. Their paper contains examples of some of the non-photorealistic filters created. In one of his works, Greenfield uses geometric measurements induced by the color organization of the images.[46] The algorithm reduces the images to a small number of regions of the same color. Fitness is assigned by performing a weighted sum of three geometric assessments of these regions: the sum of their areas, of their boundary lengths, and of the adjacencies among them. In a later work, Greenfield used a multiple-objective optimization approach to fitness assignment.[47] The goal was to maintain several "species" within the same population of images. He defines multiple static fitness functions. Each of these functions uses two out of three previously described assessments. Since the assessments are incompatible, this fosters the specialization of the individuals and the appearance of species.

The first attempt to fully automate an evolutionary art system appears in the work of Baluja, Pomerlau, and Todd.[48] They begin by using interactive evolution to build a set of images evaluated by users. In a later stage, they use machine-learning techniques, namely, Artificial Neural Networks (ANNs). ANNs are inspired in the structure and functional aspects of the brain. They are composed by a set of interconnected (artificial) neurons, and have been used successfully in a wide variety of tasks, including time series prediction, robotic controllers, and face and character recognition. The appeal of ANNs is their ability to learn from a set of examples. Baluja and colleagues use the set of images evaluated by the user to train an ANN, which is meant to learn the preferences of the users. Although the authors classify the results as "somewhat disappointing," this work is an important step toward the automation of fitness assignment.

Saunders and Gero use a Self-Organizing Map ANN to assign fitness to the images produced by an expression-based evolutionary system.[49] Self-Organizing Map ANNs do not require a training step; they automatically organize the input data—in this case, images—into clusters according to their similarity. (It is important to notice that this computational similarity may be significantly different from similarity as perceived by humans.) The goal of their system is to study the emergence of novelty. As such, fitness depends on the dissimilarity of an image to the existing clusters of images. In general terms, the experimental results indicate that their evolutionary algorithm tends to produce increasingly complex images.

The work of Saunders and Gero also has a co-evolutionary inspired component. They use an agent-based framework, where each artificial agent has its own expression-based evolutionary system and Self-Organizing Map. When an agent finds an image that has the "right" degree of novelty, it shares the genotype with other agents. If a receiving agent also finds the image adequately novel, the genotype is included in the population of its evolutionary system and the receiving agent issues a credit to the agent that discovered the image. Thus, although there is not a direct arms race among agents, agents influence each other by communicating the genotypes of the artworks they produce to other agents.

Rooke was the first one who made an attempt to use a co-evolutionary approach in the context of evolutionary art, but he never published the results of these experiments. Gary Greenfield states that "Rooke's idea was to try to co-evolve a population of art critics, which he called 'image commentators,' to perform the aesthetic evaluations of the images his Sims' inspired system generated."[50]

Greenfield implemented an interesting solution to the automation of the fitness assignment. He co-evolves a population of images, using an expression-based evolutionary system, and a population of convolution image filters. For determining fitness, the filters are applied to the images, possibly changing them. The fitness of an image is proportional to the amount of change introduced by the filters. The fitness of a filter is inversely proportional to the changes it introduces in the population of images. Thus, in simple terms, each filter acts as a "parasite" and its survival depends on the ability to pass unnoticed. On the other hand, the fitness of an image depends on its ability to identify the parasites, by making them visible. This co-evolutionary setup tends to converge toward states where white-noise images are predominant. Nevertheless, a given type of imagery is, typically, unable to dominate the population for too long because of the evolutionary pressure caused by the parasites. This constant tension generates an interesting evolutionary dynamic.

The latest development of NEvAr is also inspired by co-evolution. Like Baluja and colleagues, we employ ANNs. However, there is an important difference: by employing a set of metrics—for example, complexity estimates—we measure several characteristics of the images; the ANNs base their judgments on the characteristics of the images. Thus, in simple terms, the ANNs never see the images; they only have access to their characteristics. By these means, the training of the ANNs is performed using information of a higher level of abstraction than the images' pixels.

The goals of this setup are twofold: first, the creation of images without human intervention. The only information provided to the system is a set of images—in this case five thousand paintings made by renowned artists—that acts as an aesthetic reference. Second, the creation of a system that systematically explores new styles. For this purpose, we employ a method, inspired by co-evolution that promotes stylistic change from one evolutionary run to the next.

The ANNs are trained by showing them two classes of images: a set of paintings by well-known authors and a set of images generated randomly by NEvAr. Once this is done, an evolutionary run is initiated and the trained ANNs are used to assign fitness to the images evolved by NEvAr. The evolved images that the ANNs fail to identify as being produced by NEvAr have maximum fitness. The goal is twofold: first, to evolve images that relate to the aesthetic reference provided by the first class, which can be considered to be an inspiring set; second, to evolve images that are novel in relation to the imagery typically produced by the system. Thus, rather than trying to replicate a given style, the goal is to break away from the traditional style of the system. Once novel imagery is found—that is, when NEvAr is able to find images that the classification system fails to classify as being created by NEvAr—these images are added to the second class, the classifier is retrained and a new evolutionary run begins. This process is iteratively repeated and the method fosters a permanent search for novelty and deviation from previously explored paths.

In our last experiment, the system performed twelve consecutive evolutionary runs.[51] During these runs, the arms race between the ANNs and the evolutionary system was always balanced: the evolutionary engine was always able to find images that were misclassified by the ANNs. However, after these were added to the set of training images, the ANNs were always able to discriminate between the images created by evolution and the set of paintings (with a success rate above 98 percent), thus fostering a stylistic change in the next evolutionary run. Additionally, we conducted some tests using the ANNs from different iterations to classify

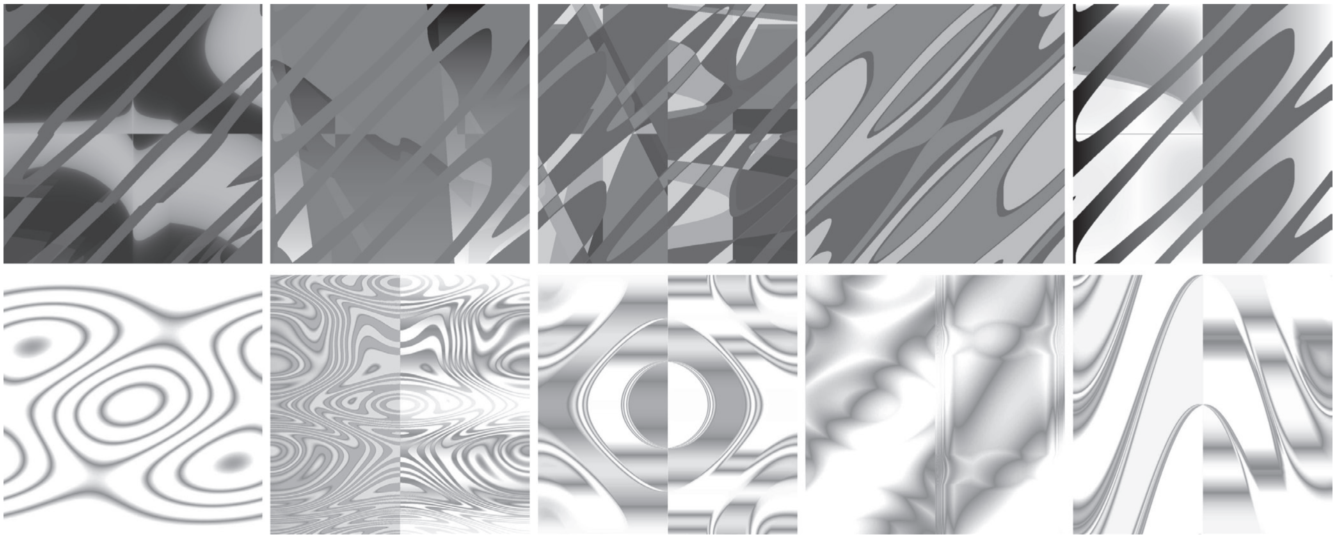

Figure 6. Examples of images created using NEvAr's approach to stylistic change. The images in the top row characterize the type of imagery being produced during the first evolutionary run of the process; the ones in the bottom row depict the style of the eleventh evolutionary run of the experiment.

external imagery that did not belong to the training sets. The experimental results show a gradual improvement of performance as the number of runs increases. This result indicates that the competition between systems improved not only the performance of the ANNs but also their generalization abilities. The evaluation of the aesthetic merit of the evolved images is subjective. Nevertheless, it is reasonable to state that in each run evolution converged to a consistent style of imagery, and that the evolved styles are all significantly different from each other. Hence, the main goals of the approach were attained.

The assessment of the aesthetic merit of the evolved images is subjective; therefore we provide examples of images evolved during the first and twelfth evolutionary runs. The CD accompanying the handbook *The Art of Artificial Evolution* includes the thirty thousand images evolved in the course of this experiment.[52] Additionally, it is safe to say that in each run evolution converged to a consistent style of imagery, and that the evolved styles are all significantly different from each other.

Evolutionary art research is reaching maturity, and part of this process is the growing awareness of the various social, artistic, and scientific challenges the area faces. From a social standpoint, building tools that foster the creativity of the user is probably the most prominent goal. From an artistic perspective, the greatest challenge is the acceptance of the evolutionary approach as a significant art practice. Promoting the participation and close involvement of artists in the design of these systems is crucial for meeting this challenge. From a scientific standpoint, the most important challenges include the automation of fitness assignment, the creation of systems that develop their own aesthetics, the integration and interaction of these systems in a cultural environment, and the pursuit of new forms of human–machine interaction.

Evolutionary art results in dynamic models whose behavior is neither totally defined nor predictable by the model's creator. In fact, since the only requirement is having a way to assign fitness, EC allows the discovery of solutions to problems for which very little knowledge is available. This allows their application to creative tasks, such as art, that we are—and might always be—far from fully understanding.

BIBLIOGRAPHY

Arnheim, Rudolf. *Towards a Psychology of Art/Entropy and Art: An Essay on Disorder and Order*. Berkeley: University of California Press, 1966.

Appel, Kenneth, and Wolfgang Haken. "The Solution of the Four-Color Map Problem." *Scientific American* 237 (1977): 108–21.

Baker, Ellie. "Evolving Line Drawings." Technical Report TR-21-93, Harvard University Center for Research in Computing Technology, 1993.

Baluja, Shumeet, Dean Pomerlau, and Jochem Todd. "Towards Automated Artificial Evolution for Computer-Generated Images." *Connection Science* 6 (1994): 325–54.

Collomosse, John. "Evolutionary Search for the Artistic Rendering of Photographs." In *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, edited by Juan Romero and Penousal Machado, 39–62. Berlin: Springer, 2007.

Darwin, Charles. *On the Origin of Species*. London: John Murray, 1859.

Dawkins, Richard. *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe without Design*. New York: W.W. Norton, 1987.

Dennett, Daniel. *Darwin's Dangerous Idea*. London: Penguin, 1995.

Dissanayake, Ellen. *What Is Art For?* Seattle: University of Washington Press, 1987.

Draves, Scott. *Electric Sheep*, http://www.electricsheep.org.

———. "Evolution and Collective Intelligence of the Electric Sheep." In *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, edited by Juan Romero and Penousal Machado, 63–78. Berlin: Springer, 2007.

Dubbin, Greg, and Stanley Kenneth. "Learning to Dance through Interactive Evolution." In *Applications of Evolutionary Computation, EvoApplications 2010: EvoCOMNET, EvoENVIRONMENT, EvoFIN, EvoMUSART, and EvoTRANSLOG, Istanbul, Turkey, April 7–9, 2010, Proceedings, Part II*, edited by Cecilia Di Chio et al., 331–40. Berlin: Springer, 2010.

Durbin, Richard, and David Willshaw. "An Analogue Approach to the Travelling Salesman Problem Using an Elastic Net Method." *Nature* 326 (1987): 689–91.

Field, David, Anthony Hayes, and Robert F. Hess. "The Roles of Polarity and Symmetry in the Perceptual Grouping of Contour Fragments." *Spatial Vision* 1 (2000): 51–66.

Fisher Yuval. *Fractal Image Compression: Theory and Application*. New York: Springer, 1995.

Greenfield, Gary. "Co-Evolutionary Methods in Evolutionary Art." In *The Art of Artificial Evolution: A Handbook on Evolutionary

Art and Music, edited by Juan Romero and Penousal Machado, 357–80. Berlin: Springer, 2007.

———. "Color Dependent Computational Aesthetics for Evolving Expressions." In *Bridges: Mathematical Connections in Art, Music, and Science; Conference Proceedings*, edited by Reza Sarhangi, 9–16. Winfield, KS: Central Plains Book Manufacturing, 2002.

———. "Evolving Aesthetic Images Using Multiobjective Optimization." In *Evolutionary Computation, CEC 2003 Proceedings, Vol. 3*, edited by Bob McKay et al., 1903–09. Canberra, Australia: IEEE Press, 2003.

———. "Robot Paintings Evolved Using Simulated Robots." In *Applications of Evolutionary Computing, EvoWorkshops 2006: EvoBIO, EvoCOMNET, EvoHOT, EvoIASP, EvoINTERACTION, EvoMUSART, and EvoSTOC, Budapest, Hungary, April 10–12, 2006, Proceedings*, edited by Franz Rothlauf et al., 611–21. Berlin: Springer, 2006.

Hart, David August. "Toward Greater Artistic Control for Interactive Evolution of Images and Animation." In *Applications of Evolutionary Computing, EvoWorkshops 2007: EvoCoMnet, EvoFIN, EvoIASP, EvoINTERACTION, EvoMUSART, EvoSTOC and EvoTransLog Proceedings*, edited by Giacobini et al., 527–36. Berlin: Springer, 2007.

Hemberg, Martin, Achim Menges, Katrin Jonas, Michel da Costa Gonçalves, and Steve Fuchs. "Genr8: Architects Experience with an Emergent Design Tool." In *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, edited by Juan Romero and Penousal Machado, 167–88. Berlin: Springer, 2007.

Hervás, Raquel, Jason Robinson, and Pablo Gervás. "Evolutionary Assistance in Alliteration and Allelic Drivel." In *Applications of Evolutinary Computing, EvoWorkshops 2007: EvoCoMnet, EvoFIN, EvoIASP,EvoINTERACTION, EvoMUSART, EvoSTOC and EvoTransLog Proceedings*, edited by Giacobini et al., 537–46. Berlin: Springer, 2007.

Holland, John H. *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press, 1975

———. "Building Blocks, Cohort Genetic Algorithms, and Hyperplane-Defined Functions." *Evolutionary Computation* 8 (2000): 373–91.

Hopfield, John. "Neurons with Graded Response Have Collective Computational Properties Like Those of Two-State Neurons." *Proceedings of the National Academy of Sciences of the United States of America* 81 (1984): 3088–92.

Jacob, Christian, and Gerald Hushlak. "Evolutionary and Swarm Design in Science, Art, and Music." In *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, edited by Juan Romero and Penousal Machado, 145–66. Berlin: Springer, 2007.

Kirkpatrick, Scott, D. Gelatt Jr., and Mario P. Vecchi. "Optimization by Simulated Annealing." *Science* 220 (1983): 671–80.

Koza, John. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge: MIT Press, 1992.

Lewis, Matthew. "Evolutionary Visual Art and Design." In *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, edited by Juan Romero and Penousal Machado, 3–37. Berlin: Springer, 2007.

Machado, Penousal, and Amilcar Cardoso. "Computing Aesthetics." *XIVth Brazilian Symposium on Artificial Intelligence SBIA'98 (Porto Alegre, Brazil)*, edited by Flávio Moreira de Oliveira, 219–29. Berlin: Springer, 1998.

———. "All the Truth about NEvAr." *Applied Intelligence*, Special Issue on Creative Systems 16 (2002): 101–19

Machado, Penousal, and Fernando Graça. "Evolutionary Pointillist Modules: Evolving Assemblages of 3D Objects." In *Applications of Evolutionary Computing, EvoWorkshops 2008: EvoCOMNET, EvoFIN, EvoHOT, EvoIASP, EvoMUSART, EvoNUM, EvoSTOC, and EvoTransLog, Naples, Italy, March 26–28, 2008 Proceedings*, edited by Mario Giacobini et al., 453–62. Berlin: Springer, 2008.

Machado, Penousal, Juan Romero, and Bill Manaris. "Experiments in Computational Aesthetics: An Iterative Approach to Stylistic Change in Evolutionary Art." In *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, edited by Juan Romero and Penousal Machado, 381–415. Berlin: Springer, 2007.

McCormack, Jon. "Aesthetic Evolution of L-Systems Revisited." *Applications of Evolutionary Computing, EvoWorkshops2004: EvoBIO, EvoCOMNET, EvoHOT, EvoIASP, EvoMUSART, EvoSTOC (Coimbra, Portugal)*, edited by Günther Raidl et al., 475–86. Berlin: Springer, 2004.

———. "Facing the Future: Evolutionary Possibilities for Human-Machine Creativity." In *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, edited by Juan Romero and Penousal Machado, 417–51. Berlin: Springer, 2007.

Miranda, Eduardo Reck, and John Al Biles, eds. *Evolutionary Computer Music*. London: Springer, 2007.

Neufeld, Craig, Brian Ross, and William Ralph. "The Evolution of Artistic Filters." In T*he Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, edited by Juan Romero and Penousal Machado, 335–56. Berlin: Springer, 2007.

Nichio, Kenichi. "Fuzzy Fitness Assignment in an Interactive Genetic Algorithm for a Cartoon Face Search." In *Genetic Algorithms and Fuzzy Logic Systems: Soft Computing Perspective*, edited by Elie Sancheze, Takanori Shibata, and Lotfi Zadeh, 175–92. River Edge, NJ: World Scientific Publishing Co., 1997.

Pontecorvo, Michael S., and Neil Elzenga. "Exploring Designer/Consumer Dialog in Evolutionary Design Systems." In *AISB Symposium on Creative Evolutionary Systems*, edited by Peter Bentley and David Corne, 115–24. Brighton: Society for the Study of Artificial Intelligence and Simulation of Behaviour, 1999.

Romero, Juan and Penousal Machado, eds. T*he Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*. Berlin: Springer, 2007.

Sims, Karl. "Artificial Evolution for Computer Graphics." *ACM Computer Graphics* 25 (1991): 319–28.

Saunders, Rob, and John S. Gero. "The Digital Clockwork Muse: A Computational Model of Aesthetic Evolution" In *AISB'01 Symposium on Artificial Intelligence and Creativity in Arts and Science*, edited by Geraint Wiggins, 12–21. Brighton: Society for the Study of Artificial Intelligence and Simulation of Behaviour, 2001.

Takagi, Hideyuki. "Interactive Evolutionary Computation: Fusion of the Capacities of EC Optimization and Human Evaluation" *Proceedings of the IEEE* 9, no. 89 (2001): 1275–96.

Teller, Astro. "Smart Machines, and Why We Fear Them." *New York Times*, March 21, 1998.

Todd, Steven, and William Latham. *Evolutionary Art and Computers*. Winchester, UK: Academic Press, 1992.

Tyler, Christopher. *Human Symmetry Perception and Its Computational Analysis*. Mahwah, NJ: Lawrence Erlbaum Associates, 2002.

World, Linda. "Aesthetic Selection: The Evolutionary Art of Steven Rooke." *IEEE Computer Graphics and Applications* 1 (1996):16.

NOTES

1. Teller, "Smart Machines."
2. Dissanayake, *What Is Art For?*
3. Kirkpatrick, Gelatt, and Vecchi, "Optimization."
4. Hopfield, "Neurons."
5. Durbin and Willshaw, "Analogue Approach."
6. Darwin, *Origin of Species*.
7. Dennett, *Darwin's Dangerous Idea*.
8. Holland, "Building Blocks," 377.
9. The four-color theorem was first conjectured the theorem in 1852 by F. Guthrie and was proven in 1977 by Appel and Haken ("Solution"), who constructed a computer-assisted proof.
10. Holland, *Adaptation*; Koza, *Genetic Programming*.
11. Dawkins, *Blind Watchmaker.*
12. For a comprehensive survey of evolutionary and biological concepts in the visual arts, see Lewis, "Evolutionary Visual Art."
13. Draves, *Electric Sheep*.
14. See Draves, "Evolution."
15. Lewis, "Evolutionary Visual Art," 5.
16. Draves, "Evolution."
17. Lewis, "Evolutionary Visual Art"; Nichio, "Fuzzy Fitness Assignment."
18. Lewis, "Evolutionary Visual Art."
19. Baker, "Evolving Line Drawings."
20. Hemberg et al., "Genr8."
21. Pontecorvo and Elzenga, "Exploring Designer/Consumer Dialog."
22. Sims, "Artificial Evolution."
23. Todd and Latham, "*Evolutionary Art*"; World, "Aesthetic Selection"; Hart, "Toward Greater Artistic Control."
24. Koza, *Genetic Programming*.
25. Machado and Cardoso, "All the Truth."
26. Machado and Cardoso, "All the Truth."
27. McCormack, "Facing the Future," 425.
28. Machado and Graça, "Evolutionary Pointillist Modules."
29. Collomosse, "Evolutionary Search."
30. Greenfield, "Robot Paintings Evolved."
31. Jacob and Hushlak, "Evolutionary and Swarm Design."
32. McCormack, "Aesthetic Evolution."
33. Neufeld and Ross, "Evolution."
34. Hart, "Toward Greater Artistic Control."
35. Hemberg et al., "Genr8."
36. Miranda and Biles, eds., *Evolutionary Computer Music.*
37. Hervás, Robinson, and Gervás, "Evolutionary Assistance."
38. Dubbin and Kenneth, "Learning to Dance."
39. Hemberg et al., "Genr8."
40. Takagi, "Interactive Evolutionary Computation"
41. Machado and Cardoso, "Computing Aesthetics"; Arnheim, *Towards a Psychology*; Field, Hayes, and Hess, "Roles of Polarity"; Tyler, *Human Symmetry Perception*.
42. Machado and Cardoso, "Computing Aesthetics."
43. Fisher, *Fractal Image Compression*.
44. Machado and Cardoso, "All the Truth."
45. Neufeld, Ross, and Ralph, "Evolution."

46. Greenfield, "Color Dependent Computational Aesthetics."
47. Greenfield, "Evolving Aesthetic Images."
48. Baluja, Pomerlau, and Todd, "Automated Artificial Evolution."
49. Saunders and Gero, "Digital Clockwork Muse."
50. Greenfield, "Co-Evolutionary Methods," 361.
51. Machado, Romero, and Manaris, "Experiments."
52. Romero and Machado, *Artificial Evolution*.