

Expressions, Assemblages and Grammars

Penousal Machado
CISUC, Department of Informatics Engineering
University of Coimbra
Pólo II
3030-290 Coimbra, Portugal
E-mail: machado@dei.uc.pt

Abstract

The use of bio-inspired computing offers a wide range of solutions and opportunities, both for scientists, who make an effort to understand and model nature, and for artists, who use nature as a source of inspiration. The use of a metaphor that is relevant for scientists and artists helps to bridge the gap between the scientific and artistic communities, and fosters the collaboration and transfer of knowledge between the two domains. This paper presents an overview of some of the Evolutionary Art projects I have been involved during the past thirteen years.

Introduction

Since its inception Artificial Intelligence (AI) research has given considerable emphasis to logic, reasoning, problem solving, planning, natural language processing, expert systems, chess, etc. Although I come from an AI background I have never been particularly interested in this type of issue. Back in 1984, when I was 14, I saw a movie called “Electric Dreams” that left a lasting impression on me. I had just bought my first computer, a 48K ZX Spectrum; the movie was about a boy who bought a computer and accidentally spilled a beverage on it. As a consequence, the computer developed AI. Listening to the next-door girl 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. I remember thinking that this was the coolest thing ever. To a large extent, I still do.

The current state of the art AI systems are able to attain human competitive results, or even to surpass human performance, in several tasks where intelligence is (or was) regarded as a requirement. Yet, in other domains, namely those requiring creative reasoning, such as art, design, music and poetry, computers are, arguably, far from reaching the accomplishments of humans. Additionally, art-making activities are ubiquitous, and creativity is often regarded as one of the most remarkable characteristics of the human mind. Therefore, I find that the search for computational creativity is 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 may be used. There is a growing interest in bio-inspired computing, an area of research that comprises techniques such as evolutionary computation, swarm intelligence, ant colony optimization, and artificial life. Personally, I find it far more interesting to build a system that possesses a different kind of intelligence than our own, and that can complement ours, rather than trying to imitate human intelligence.

For the above reasons, I have spent a significant portion of the past 14 years developing evolutionary art systems that can help the user create new artifacts or produce them autonomously. This paper gives examples of some of the systems and artworks produced during this period.

Expression-Based Evolutionary Art

The influential works of Richard Dawkins [1], Karl Sims [2] and William Latham [3] led to the emergence of a new research area, usually called Evolutionary Art (EA), which is characterized by the use of nature-inspired computing for artistic purposes. A thorough survey on the application of biological-inspired techniques to visual art can be found in [4].

NEvAr stands for Neuro Evolutionary Art and it is a system that Amílcar Cardoso and I began developing in 1997. Largely inspired by the work of Karl Sims [2], it allows the evolution of populations of images. Like Sims, we employ Genetic Programming (GP). Each genotype is a LISP-like *s-expression*, represented as a tree, and constructed from a lexicon 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 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 *s-expressions*. 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.

The initial versions of NEvAr only allowed interactive evolution, i.e., the user assigns fitness to the images, indirectly determining the survival and mating probabilities of the individuals. The fittest individuals have a higher probability of being selected for the creation of the next population, which is generated through the recombination and mutation of the genetic code of the selected individuals. The genetic operations are performed at the genotype level. Recombination is performed using Koza's crossover operator [5], which exchanges sub-trees between individuals. Five mutation operators are used: sub-tree swap, sub-tree replacement, node insertion, node deletion, and node mutation [6].

The system has a signature – in the sense that it is more prone to generate certain types of images than others – that is intimately related with the used function set, genetic operators, and genotype-phenotype mapping. Nevertheless, in theory, it is possible to generate any image [6]. Assuming that the user is patient, it is possible to evolve from an initial, randomly generated, population to images that fit the preferences of the user. Figure 1 presents examples of images created with NEvAr using interactive evolution.




Figure 1: Examples of images created with NEvAr using interactive evolution.




Figure 2: Images created using complexity estimates.

Over the years, we gradually expanded the system by integrating modules that automate fitness assignment step. In our first attempt, we employed image complexity estimates based on JPEG and quad-tree based fractal image compression [7] to assign fitness. We were pleased with the grayscale images generated using this approach (see Figure 2 for examples) but we were unable to create color images that satisfied us. The inability to solve this issue made us explore a different approach: evolving programs that apply color to grayscale images.

We use GP to evolve programs that take as input a grayscale image and produce as output a Hue channel for the input image. To assign fitness, each individual, i.e. program, is executed over a set of training cases composed of several color images in HSV format. The individuals receive as input the Value channel of each training image and their output is compared with desired Hue channel. Once satisfying coloring programs evolve, they can be applied to any input grayscale image, including images that were not used in their training.

Exact color reproduction was not a goal. Instead, we were interested in the evolution of aesthetically gratifying colorings. The main difficulty was the development of an appropriate way to compare the colorings produced with the target ones. The detailed explanation of the system and of the fitness assignment scheme can be found in [8]. A similar approach can be applied to evolve programs that produce a Saturation channel. In Figure 3 we present some of the colorings produced by one of the evolved programs trained using some of Wassily Kandinsky's artworks.

Later, in close collaboration with Juan Romero and his team, we developed fitness assignment schemes using Artificial Neural Networks (ANNs) trained with a set of examples, and showed that it is possible to attain human competitive results in Maitland Graves' "Design Judgment Test", a psychological test created to assess how humans respond to several principles of aesthetic order [9].

NEvAr is still an ongoing project. Our latest progress [10] consists of an approach that promotes stylistic change from one evolutionary run to the other. In this case a set of neural networks assigns fitness to the evolved images, guiding the evolutionary engine. These networks are trained by exposing them to two classes of images: a set of artworks of well-known authors; a set of images generated randomly by the system. The goal is twofold: (i) to evolve images that relate to the aesthetic reference provided by the first class, which can be considered an inspiring set; (ii) to evolve images that are novel in relation to the imagery typically produced by the system. Thus, more 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 – i.e. when



Figure 3: Output of one of the evolved coloring programs when applied to one of the training images (leftmost image) and to images that were not training cases.

NEvAr is able to find images that the ANNs fail to classify as being created by NEvAr – these images are added to the second class, the neural networks are re-trained and a new evolutionary run begins. This process is iteratively repeated and, by these means, a permanent search for novelty and deviation from previously explored paths is enforced.

Evolving Assemblages

The research made in the scope of the NEvAr project mainly concerns aesthetic judgment. This imposes a series of constraints. For instance, it is trivial to expand the function set of NEvAr, and by including high-level primitives one can generate aesthetically pleasing images, and explore different types of imagery, significantly faster than with the low-level primitives we are using. However, using simple low-level primitives allows us to get a better grasp of how fitness assignment is guiding evolution. To avoid this type of constraint we got involved in a series of projects that gave us more “artistic” freedom than NEvAr. In this section I will briefly describe *Evolving Assemblages*, a project I developed with Fernando Graça, who was my student at the time.

The goal was to produce large-format reproductions of input images by assembling 3D objects. As previously, we use GP and interactive evolution. Each individual receives as input an image 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 determines the size of each object; the fourth determines the x coordinate where the object will be placed; and, finally, the fifth 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.

The assemblage produced by an individual is functionally dependent on the input image. Graça spent a considerable amount of time evolving assemblages and fine-tuning the system, before he was able to evolve individuals that consistently generated assemblages that we found interesting for different input images. Figure 4 depicts some of the assemblages generated by one of the evolved individuals, while Figure 5 presents a detail of one of the assemblages. A full description of the system can be found in [11].




Figure 4: *Assemblages produced by one of the evolved individuals for different input images.*



Figure 5: *Detail of the leftmost assemblage presented in Figure 4.*

Evolving Visual Grammars

The spark for the project described in this section was realizing, while evolving images in the scope of EA projects, that the populations were often more interesting than the individuals. The individuals of a population tend to share a similar genetic background. As such, in some sense, the variations within the population give hints of the underlying structure of their genotypes, and one can look at a population as a set of individuals of the same species. I always found this variation within a style quite appealing.

The main inspiration was the seminal work of Stiny and Gips [12], who introduced the concept of Shape Grammars and built, among others, shape grammars that capture the architectural “language” of Frank Lloyd Wright’s prairie houses, Mughul Gardens, Palladian Plans, etc. Additionally, they were also able to use these grammars to produce new instances of the same language, e.g. to create new prairie house designs that obey Frank Lloyd Wright’s style to the point of being indistinguishable, even to experts, from his original works [13].

Context Free [14] is a popular open-source application that renders images specified using a simple language entitled Context Free Design Grammar (CFDG) [15]. The use of the CFDG for representation allows the specification of complex families of shapes through a compact set of rules. As Figure 6 illustrates, the use of a simple non-deterministic grammar is enough to generate a family of tree-like shapes.

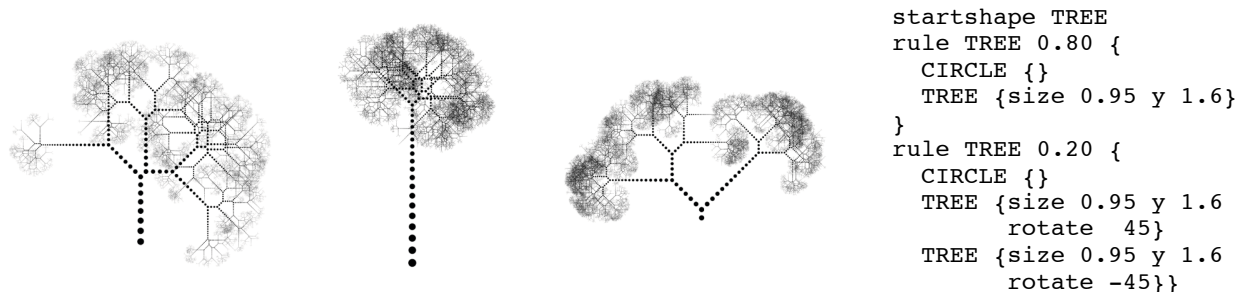


Figure 6: *Instances of images created by the CFDG grammar presented on the right.*

In 2009 we started developing a system for the evolution of visual languages. Together with Henrique Nunes, one of my students, we created an evolutionary engine where each individual is a CFDG and developed appropriate genetic operators, including several mutation operators and graph-based crossover. This approach has several advantages over the previously described ones. Hand-coding an expression-based genotype that produces an interesting image is a hard task, while writing an interesting CFDG program is relatively easy. Evolved expression-based genotypes can easily have thousands of nodes, making them hard to understand and, therefore, edit them meaningfully. Due to the popularity of Context Free, a large number of hand-coded CFDG programs, which can be used to initialize the evolutionary runs, is available.

To test the ability of the system to use hand-coded grammars, we conducted several interactive evolution runs where the first population was made of hand-coded CFDGs (see Figure 7). This allowed us to evolve a wide variety of shapes, such as the ones presented in Fig 8.




Figure 7: Instances of images generated by the six hand-coded grammars of the initial population.




Figure 8: Instances of the images generated by four of the evolved grammars.

To assess the adequacy of the evolutionary engine, and its ability to generate interesting shapes, and families of shapes, without resorting to hand-coded grammars, we performed experiments using several automated fitness assignment schemes as well as interactive evolution [16]. In these runs the initial population was composed of one grammar that draws a single square. In Figure 9 we present some of the images evolved in these runs using automated fitness assignment schemes. Although they may be less appealing than the previous ones, they show that the system is able to promote complexification and that interesting images can be produced in relatively short runs (100 generations) even when starting with the simplest grammar possible.




Figure 9: *Instances of images generated in automated fitness assignment runs. The initial population consisted in a single grammar that generated a black square.*

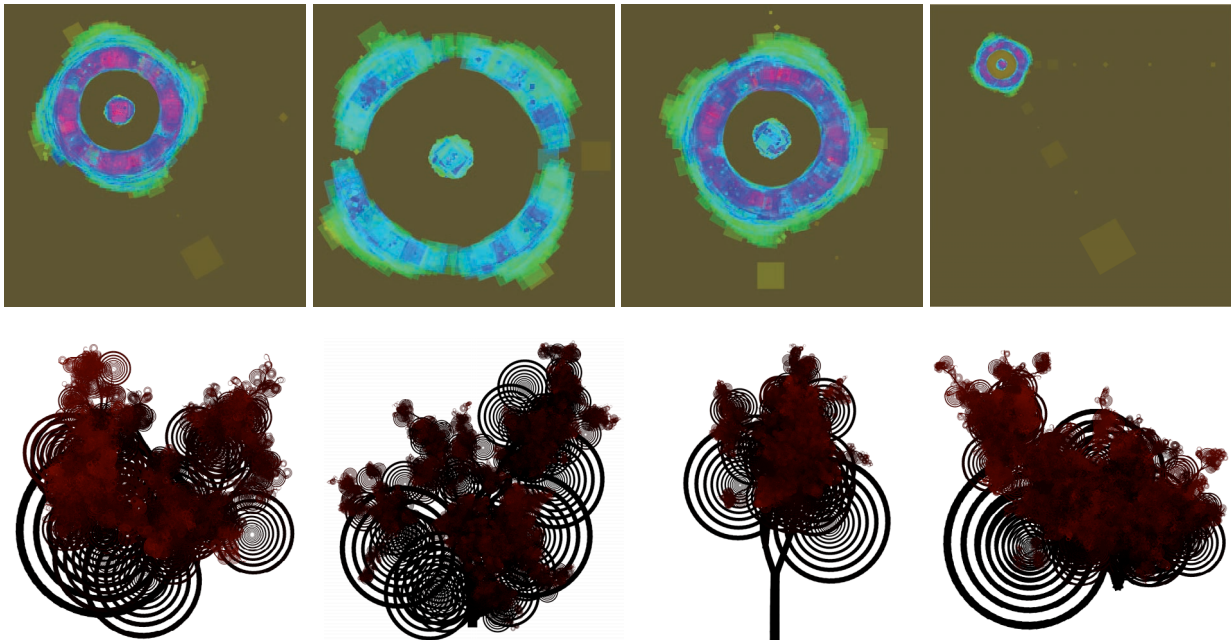


Figure 10: *Instances of the shapes generated by two of the evolved grammars.*

Our main goal was the evolution of visual languages. In Figure 10 we present two of the many visual languages developed in the course of the experiments. Currently we are working with designers in the creation of dynamic visual identities, i.e., instead of establishing the visual identity by defining a specific image, the visual identity is defined by the use of a set of closely related images over time. A well-known example of such approach is the visual identity of “Casa da Música” created by Stefan Sagmeister [17]. In a nutshell, our idea consists in using our system to evolve visual identities: a grammar becomes the visual identity producing instances of that visual identity on demand.

Conclusions

Looking at the evolutionary art projects I have been involved in during the past years, I find that the results attained are secondary when compared with the joy their development gave me. I had the opportunity to work in something I am passionate about and to work with fascinating students, researchers, artists and designers, willing to take risks and think outside of the box. The evolutionary art projects also played an important role in the classroom, allowing me to motivate students by showing them that Artificial Intelligence techniques can be used to create what they consider to be “cool stuff”.

Acknowledgements.

I like to thank the anonymous reviewers for their thorough and insightful comments. I would also like to thank Juan Romero, Amílcar Cardoso, Fernando Graça and Henrique Nunes who have played a key role in some of the projects mentioned here.

References

- [1] R. Dawkins, *The Blind Watchmaker*, Penguin Books, 1986.
- [2] K. Sims, Artificial Evolution for Computer Graphics, *ACM Computer Graphics*, 25(4): 319-328. 1991.
- [3] S.Todd and W. Latham, *Evolutionary Art and Computers*, Academic Press, 1992.
- [4] M. Lewis, Evolutionary Visual Art and Design, in *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, Springer, pp. 3-37, 2007.
- [5] J. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Complex Adaptive Systems. The MIT Press. 1992.
- [6] P. Machado and A. Cardoso, All the Truth About NEvAr, *Applied Intelligence*, Special Issue on Creative Systems, 16(2):101–19. 2002.
- [7] Yuval Fisher (Ed.), *Fractal Image Compression: Theory and Application to Digital Images*, Springer Verlag, New York, 1995
- [8] P. Machado, A. Dias and A. Cardoso, Learning to Colour Greyscale Images, *The Interdisciplinary Journal of Artificial Intelligence and the Simulation of Behaviour – AISB Journal*, 1(2): 209-219, 2002.
- [9] P. Machado, J. Romero, A. Santos, A. Cardoso, A. Pazos, On the Development of Evolutionary Artificial Artists. *Computers & Graphics* 31(6): 818-826. 2007.
- [10] P. Machado, J. Romero, B. Manaris, Experiments in Computational Aesthetics, *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, Springer, pp. 381-415, 2007.
- [11] P. Machado and F. Graca, Evolutionary Pointillist Modules: Evolving Assemblages of 3D Objects. *EvoWorkshops 2008*: 453-462. 2008.
- [12] G. Stiny and J. Gips, Shape grammars and the generative specification of paintings and sculpture, *Information Processing 71*, Amsterdam, North Holland Publishing Co. pp.1460-1465. 1971.
- [13] M.A. Boden, *The Creative Mind: Myths and Mechanisms* Basic Books, New York, 1990.
- [14] J. Horigan and M. Lentzner: *Context Free*, <http://www.contextfreeart.org/>, accessed in Jan. 2010.
- [15] C. Coyne., *Context Free Design Grammar*, <http://www.chriscoyne.com/cfdg/>, accessed in Jan. 2010.
- [16] P. Machado , H. Nunes, J. Romero, *Graph-Based Evolution of Visual Languages*, EvoApplications, 2010
- [17] S. Sagmeister, *The power of time off*. TED Talks, 2009.