SURVIVAL OF THE BEAUTIFUL

PENOUSAL MACHADO
machado@dei.uc.pt

CISUC, Department of Informatics Engineering, University of Coimbra

Abstract
More than twenty-five years have passed since Karl Sims proposed the use of artificial evolution techniques to evolve computer graphics. Since then Evolutionary Computation has been applied to many fields of Art, Music and Design. In this talk we will overview how these techniques have and are being used to create different types of content. We will then focus on an open problem in evolutionary art — fitness assignment — analyzing it from the perspective of the interplay between the evolutionary system and the user. More specifically, we will discuss how Machine Learning and HCI techniques can be combined to create systems that allow the users to express their artistic or aesthetic intentions.

The birth of evolutionary art as an area of research is deeply linked to the work of three pioneers: Richard Dawkins, Karl Sims and William Latham. Back in 1987, at the first Artificial Life conference, Dawkins presented his system, Biomorphs, which demonstrated the power of evolution by letting the user act as a selective breeder, evolving the morphology of virtual creatures. Not long after, Sims (1991) demonstrated how user-guided evolution could be used to evolve abstract imagery, plant like shapes (i.e. Lindenmayer systems) and animations. The work of Latham (1994) popularized evolutionary art by making it available to the general public. Since then many researchers have focused on the application of evolutionary techniques for the generation of computer graphics, and nowadays evolutionary art is a well-established area of research. An analysis of work done throughout these years reveals two main issues: representation and evaluation.

In what concerns representation, we can identify three main types of approach: declarative, parametric and procedural. In declarative representations the genetic code (genotype) directly encodes characteristics of the individual (phenotype). For instance, one could try to evolve images by directly evolving the color of each of their pixels, or the coordinates of the set of lines that compose it (see, e.g., Baker, 1993). As the name indicates, in parametric representations the genotype encodes a set of parameters, which are used by a model to generate the phenotype. Along with Dawkins' Biomorphs the *Electric Sheep* project (see Draves, 2008), which resorts to the evolution of the parameters of a fractal formula, is probably the most notable example of such approach. Finally, in procedural representations the genotype is a program or procedure that when executed generates the phenotype.

The seminal work of Sims (1991) employed a procedural encoding for the evolution of images, which is remarkably similar to Genetic Programming (Koza, 1992) approaches. More often than not, the programs assume a tree-like shape: the internal nodes of the tree are functions (e.g. arithmetic and trigonometric operations) while the leafs are terminals (e.g. variables and random constants).

As pointed out by Machado and Cardoso (2002), and eloquently explained by McCormack (2008), even when the function set is composed of rather simplistic functions, it is possible to demonstrate that these procedural representations have the potential to generate any image. However, as they also point out, practice is an entirely different matter: the type of image these systems tend to generate is intimately linked with the nature of their function sets. Due to its impact on the outcome of the systems, the choice of an adequate representation and the proposal of new representation methods has remained a key topic of evolutionary art throughout the years. Focusing on our most recent efforts in this domain, we highlight the use of a multi-chromosome Genetic Programming approach to evolve assemblages of objects (Graça and Machado, 2015) and graphs to evolve non-deterministic context free design grammars able to create a family of images from a single genotype (Machado, Correia and Assunção, 2015).

While the representation defines what can be generated and, implicitly, the likelihood of generating a given artifact, evaluation determines how the search space is traversed. As previously mentioned, early efforts relied on user-guided evolution, i.e. the user selected which individuals (images) to breed, thus guiding the evolutionary process. While this approach has many merits, it is time consuming, requiring the constant intervention of the user and leading to user fatigue. Furthermore, in these circumstances, users tend to make their choices based on a local and limited perspective of the search space, valuing novelty over quality, which eventually hinders not only the quality of the evolved artifacts but also their nature.

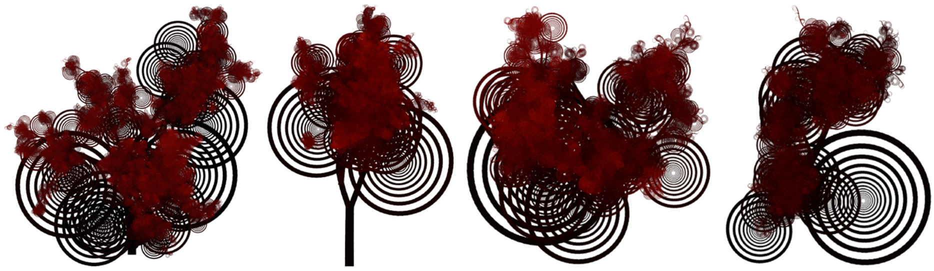

Fig. 1a, 1b, 1c, 1d
Samples of the set of images generated by an evolved non-deterministic context free design grammar, illustrating how a non-deterministic grammar may generate a wide variety of images.

The attempts to automate fitness assignment can be divided in two main groups: the use of hardwired fitness functions and the use of Machine Learning.

In the first case, the authors try to encode some sort of aesthetic criteria that may guide evolution through a function or program. The main difficulty is, not surprisingly, that it has been proven extremely hard to formally define and capture such kind of criteria. In most cases, if not all, it is trivial to show using counter-examples that the conditions considered by the authors are neither sufficient nor necessary to capture a general notion of aesthetics. Nevertheless, many examples exist that illustrate how some aesthetic principles may be explored and exploited in this context. For instance: Machado and Cardoso (2002) use complexity estimates to assign fitness; Greenfield (2003) proposes a multi-objective optimization approach to evolve images that satisfy several criteria; Ross et al. (2006) promote the evolution of images that show a "natural" distribution of color gradients; Romero et al. (2012) demonstrate how complexity measure can be used in aesthetic appreciation tasks, later showing how they relate to humans' perception of complexity (Machado et al., 2015); Reed (2013), revisiting Birkhoff's work, uses aesthetic measures to evolve vase designs.

Baluja et al. (1994) were the first to apply Machine Learning techniques in the context of evolutionary art. Their approach was based on artificial neural networks, which were trained using examples of images generated through user-guided evolution. Unfortunately, as the authors recognized, the results were disappointing. Romero et al. (2003) put forward the idea of combining a general purpose evolutionary art system with an image classifier trained to recognize faces, or other types of objects, to evolve such type of image. Ten years passed until the actual implementation of the idea by Correia et al. (2013), who were able to evolve recognizable faces, flowers, leafs, lips and other sorts of image using an expression based general purpose evolutionary system. In a later work (Machado et al., 2015), several classifiers are combined to evolve ambiguous images.

These works highlight the power of Machine Learning, but also its current limitations. As Baluja et al. (1994) already indicated, the evolutionary engine tends to find ways of exploiting the limitations of the neural networks, fooling them. Therefore, the convergence to images that are classified by the network as faces but that do not resemble faces to the human eye is quite frequent. This shortcoming can be explored for artistic and scientific purposes, as Correia et al. (2016) demonstrate by evolving images that are not classified as faces by the neural network they employ, although humans easily identify them as faces. In a different line of research, Machado et al. (2008) present a system that promotes the competition between the neural network classifiers and the evolutionary system, which results in a continuous pursuit of novelty, style change and re-invention.

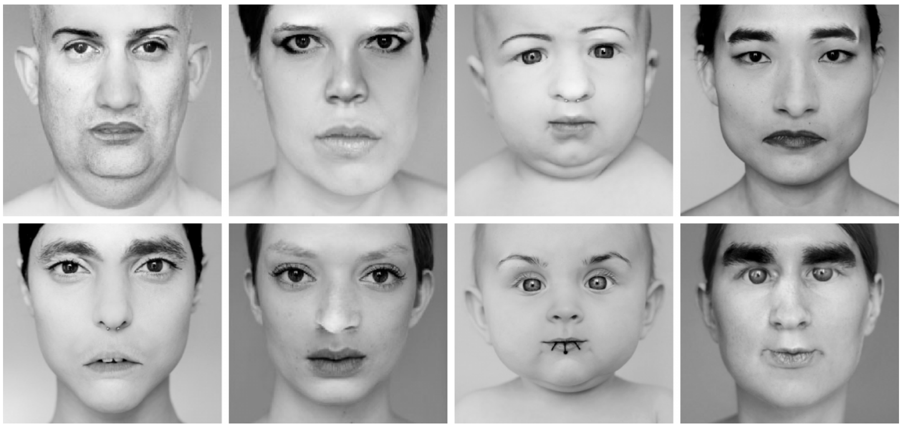

Fig. 2
Examples of images that are not recognized as faces by a neural network.

Although the automation of fitness assignment poses many relevant scientific challenges and questions, full automation has a cost: the users are no longer able to express themselves through such systems. In recent years we have focused on overcoming this problem. The core idea is to allow the users to become designers of fitness functions, allowing them to express their intentions by using a responsive interface, which implicitly defines fitness. Unlike fully automated systems, our approach engages the users making them a decisive part of the system and giving them a sense of authorship, while freeing them from the need to evaluate images individually, as it happens in traditional user-guided evolutionary systems.

Photogrowth is the first example of this approach (Machado et al., 2014). The system uses a parametric evolution approach to evolve species of artificial ants that produce non-photorealistic renderings of input images. The users are responsible for setting up the evolutionary runs and designing a fitness function through a graphical user interface. This allows them to indicate features pertaining the behavior of the ants during simulation, and features that pertain the images the ants generate. When the evolutionary runs are concluded, the users are also able to select their favorite images, apply the associated genotypes to different input images, and control the details of the final rendering.

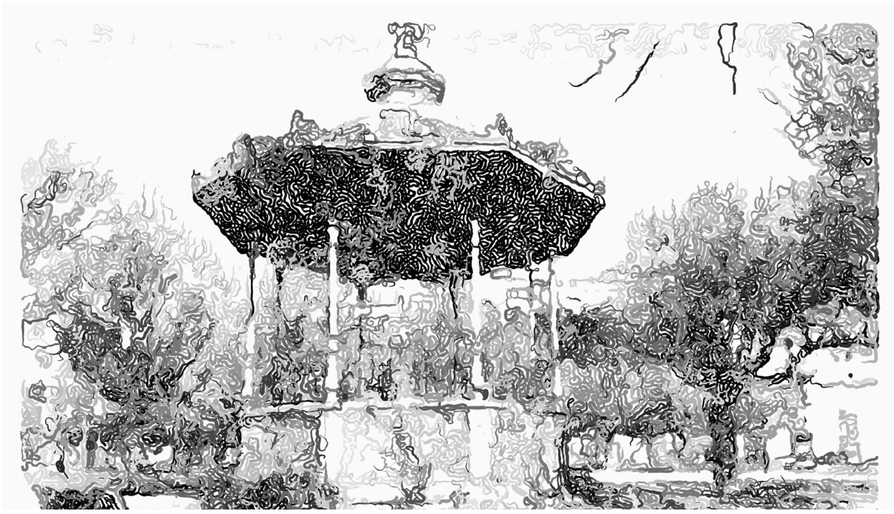

Fig. 3
Non-photorealistic rendering produced by *Photogrowth* via *insta.ants*. For additional information see cdv.dei.uc.pt/insta-ants/.

Given the current popularity of Machine Learning and the consequent wide availability of systems and tools, we believe that one of the key challenges that lies ahead, in art and in general, is to develop provably beneficial artificial intelligence systems, empowering artists and audiences, and expanding the realms of artistic creation.

REFERENCES

Karl Sims. 1991. Artificial evolution for computer graphics. *In Proceedings of the 18th annual conference on Computer graphics and interactive techniques* (SIGGRAPH '91). ACM, New York, NY, USA, 319-328. DOI=http://dx.doi.org/10.1145/122718.122752

Richard Dawkins. 1986. *The Blind Watchmaker*. New York: W. W. Norton & Company, Inc.

Stephen Todd and William Latham. 1994. *Evolutionary Art and Computers*. Academic Press, Inc., Orlando, FL, USA.

Ellie Baker. 1993. Evolving Line Drawings. *In Proceedings of the 5th International Conference on Genetic Algorithms*, Stephanie Forrest (Ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 627-.

Scott Draves. 2008. Evolution and Collective Intelligence of the Electric Sheep. In: Romero J., Machado P. (eds) *The Art of Artificial Evolution*. Natural Computing Series. Springer, Berlin, Heidelberg.

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

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

Penousal Machado and Amílcar Cardoso. 2002. All the Truth About NEvAr. *Applied Intelligence* 16, 2 (January 2002), 101-118. DOI: https://doi.org/10.1023/A:1013662402341

Jon McCormack. 2008. Facing the Future: Evolutionary Possibilities for Human-Machine Creativity. In: Romero J., Machado P. (eds) *The Art of Artificial Evolution*. Natural Computing Series. Springer, Berlin, Heidelberg

Fernando Graça and Penousal Machado. 2015. Evolving assemblages. *International Journal of Arts and Technology* (IJART). Vol. 8, iss. 2, pp. 167-184.

Penousal Machado, João Correia and Filipe Assunção. 2015 Graph-Based Evolutionary Art. In: Gandomi A., Alavi A., Ryan C. (eds) *Handbook of Genetic Programming Applications*. Springer, Cham.

Gary Greeenfield. 2003. Evolving aesthetic images using multiobjective optimization. *Congress on Evolutionary Computation*, CEC '03. IEEE Press.

Juan Romero, Penousal Machado, Adrian Carballal and Antonino Santos. 2012. Using complexity estimates in aesthetic image classification. *Journal of Mathematics and the Arts*, vol. 6, iss. 2-3, pp. 125-136.

Brian J. Ross, William Ralph and Hai Zong. 2006. Evolutionary Image Synthesis Using a Model of Aesthetics. *In Proceedings of the 2006 IEEE Congress on Evolutionary Computation*. IEEE Press.

Juan Romero, Penousal Machado, Adrian Carballal and Antonino Santos. 2012. Using complexity estimates in aesthetic image classification. *Journal of Mathematics and the Arts*, vol. 6, iss. 2-3, pp. 125-136.

Penousal Machado, Juan Romero, Marcos Nadal, Antonino Santos, João Correia, and AAdrian. Carballal. 2015. Computerized measures of visual complexity. *Acta Psychologica*. Vol. 160, iss. 1, pp. 43-57.

Kate Reed. 2013. Aesthetic Measures for Evolutionary Vase Design. In: Machado P., McDermott J., Carballal A. (eds) *Evolutionary and Biologically Inspired Music, Sound, Art and Design*. EvoMUSART 2013. Lecture Notes in Computer Science, vol 7834. Springer, Berlin, Heidelberg.

Shumeet Baluja, Dean Pomerleau and Todd Jochem. 1994. Simulating a Users Preferences: Toward Automated Artificial Evolution for Computer Generated Images. *Connection Science*. Vol. 6 , Iss. 2-3.

Juan Romero, Penousal Machado, Antonio Santos, and Amilcar Cardoso. 2003. On the development of critics in evolutionary computation artists. *In Proceedings of the 2003 international conference on Applications of evolutionary computing* (EvoWorkshops'03), Stefano Cagnoni, Colin G. Johnson, Elena Marchiori, David W. Corne, and Jens Gotlieb (Eds.). Springer-Verlag, Berlin, Heidelberg, 559-569.

João Correia, Penousal Machado, Juan Romero, and Adrian Carballal. 2013. Evolving Figurative Images Using Expression-Based Evolutionary Art. *In Proceedings of the fourth International Conference on Computational Creativity* (ICCC), pp. 24-31.

Penousal Machado, Adriano Vinhas, João Correia, and Anikó Ekárt. 2015. Evolving ambiguous images. *AI Matters* 2, 1 (October 2015), 7-8. DOI: https://doi.org/10.1145/2813536.2813539

João Correia, Tiago Martins, Pedro Martins, and Penousal Machado. 2016. X-Faces: The eXploit Is Out There. *In Proceedings of the Seventh International Conference on Computational Creativity* (ICCC 2016), pp. 164-182.

Penousal Machado, Juan Romero, and Bill Manaris. 2008. 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*, pp. 381-415.

Penousal Machado, Tiago Martins, Hugo Amaro and Pedro H. Abreu. 2014. An Interface for Fitness Function Design. *In Proceedings of the Third International Conference on Evolutionary and Biologically Inspired Music, Sound, Art and Design*, pp 23-25.