Beyond Interactive Evolution

Expressing Intentions through Fitness Functions

PENOUSAL MACHADO, TIAGO MARTINS, HUGO AMARO AND **PEDRO H. ABREU**

ABSTRACT

Photogrowth is a creativity support tool for the creation of nonphotorealistic renderings of images. The authors discuss its evolution from a generative art application to an interactive evolutionary art tool and finally into a meta-level interactive art system in which users express their artistic intentions through the design of a fitness function. The authors explore the impact of these changes on the sense of authorship, highlighting the range of imagery that can be produced by the system.

In traditional generative art, the system is built by an author who creates an algorithm. Typically this algorithm generates a variety of artworks with similar characteristics. In this scenario, users tend to perform the role of curators, selecting the outputs of the system that they consider valuable. Although interactive generative art is becoming the norm, the degree of control given to the user is relatively limited. The user tends to become a participant in an artistic experience that was designed by the author. As such, the user rarely attains a sense of authorship when interacting with such a system.

The seminal work of Karl Sims [1] led to the emergence of a new art form, evolutionary art, which is characterized by the use of evolutionary computation to evolve populations of artworks. In Sims's work, users guide evolution by indicating their favorite images. This approach to fitness assignment became known as interactive evolutionary computation (IEC). IEC allows users to guide evolution toward regions of the solution space that match their preferences. Unfortunately, it also puts a significant burden on the user, who is "forced" to evaluate thousands of artifacts and compelled to reason based on a local perspective. Additionally, although in theory it is possible to imagine an artifact and use IEC to evolve it, in practice this is impossible. Thus the user must be willing to lose a significant amount of control and embark on a journey of discovery.

Penousal Machado (educator, researcher), CISUC, Department of Informatics Engineering, University of Coimbra, 3030 Coimbra, Portugal. Email: <machado@dei.uc.pt>.

Tiago Martins (researcher), CISUC, Department of Informatics Engineering, University of Coimbra, 3030 Coimbra, Portugal. Email: <tiagofm@dei.uc.pt>.

Hugo Amaro (researcher), CISUC, Department of Informatics Engineering, University of Coimbra, 3030 Coimbra, Portugal. Email: <hamaro@student.dei.uc.pt>.

Pedro H. Abreu (researcher), CISUC, Department of Informatics Engineering, University of Coimbra, 3030 Coimbra, Portugal. Email: <pha@dei.uc.pt>.

Originally presented at the Third International Conference on Evolutionary and Biologically Inspired Music, Sound, Art and Design (EvoMUSART), 23–25 April 2014, Granada, Spain.

See <www.mitpressjournals.org/toc/leon/49/3> for supplemental files associated with this issue.

Considering these issues, we argue that IEC fails to provide a way for users to express their intentions and goals, forcing them to perform low-level assessments of individual artworks. Although users are intensively involved in the iterative cycle and, as such, in the *execution* of artworks, they may lack a sense of authorship of the *idea* that usually seeds the creative process.

To overcome the limitations of IEC, several researchers developed automated fitness assignment schemes, including the use of hardwired fitness functions, machine-learning approaches, the pursuit of novelty and so forth [2]. Automated fitness assignment is an effective way to fight user fatigue. However, to some extent, the solution defeats the purpose, because users are no longer able to express themselves. They act only as curators of the works produced by the automated system.

Our proposal—meta-level interactive evolution—overcomes some of IEC's limitations while giving users the ability to express themselves, hence eliminating one of the main shortcomings of automated approaches. The idea is simple: Take users out of the evolutionary loop and allow them to explicitly specify their goals (Fig. 1). This is accomplished by allowing users to design a fitness function that guides evolution, freeing them from the need to perform individual assessments and allowing them to express aesthetic and artistic goals by specifying the characteristics they desire among the ones considered by the system (Fig. 1).

PHOTOGROWTH

In our tool, Photogrowth, artificial ants with varying life spans, sensory capabilities and behaviors are used to produce a nonphotorealistic rendering (NPR) of an image. Each superimposed line represents the continuous trail of

 doi:10.1162/LEON_a_01103

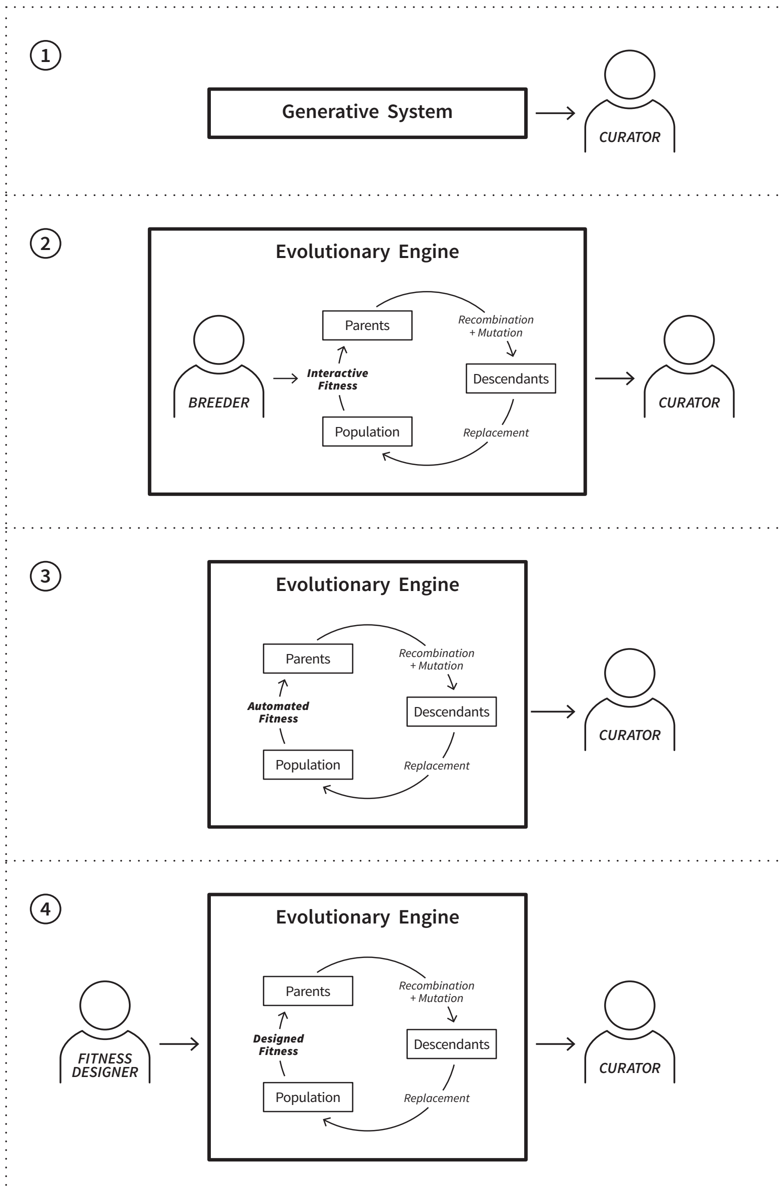

Fig. 1. Roles performed by the user in (1) noninteractive generative art, (2) interactive evolutionary computation, (3) automated evolutionary computation and (4) meta-level interactive evolutionary computation. (© Tiago Martins)

an artificial ant. The intertwinement of trails, together with their variation in width, direction and color, produce richly detailed and expressive artworks from an underlying image. Photogrowth is informed by earlier works on ant-colony art (see [3] for a survey). Photogrowth started as a generative art application, evolved into an IEC application [4] and recently became a meta-level interactive evolution tool [5].

We gather statistics describing the ants' behavior and compute several image metrics. These features are the basis for the construction of fitness functions, which are designed by users. The functions indicate the characteristics to be pursued and avoided, using an intuitive and responsive interface that allows users to perceive the semantics associated with each feature.

Figure 2 describes the pipeline of our system. Users begin by setting up the evolutionary runs and by designing a fitness function. Then they pass control to the evolutionary engine. Each genotype encodes the parameters of a species of ants, determining how it will react to the input image. Each painting (i.e. phenotype) is produced by simulating the behavior of ants of a given species while they travel on the canvas. When the evolutionary runs are finished, we further empower users by letting them select their favorite images, apply the associated genotypes to different input images and control the details of the final rendering (Fig. 2).

The virtual ants live in a 2D environment initialized with the input image, and they paint on a canvas that is initially empty and used exclusively for depositing ink. The luminance of an area determines the available energy at that point. During simulation, ants gain energy when traveling through bright areas, and this energy is removed from the source image. If the energy of an ant is below a given threshold, the ant dies; if it is above a given threshold, the ant generates offspring. The ants' movement is determined by how they react to light. Each ant has 10 sensory vectors, each with a given direction and length. These sensory "organs" return the luminance value of the area where each vector ends. The direction and length of these sensory vectors, and the way an ant responds to the sensory information, are determined by the ant's genotype: a tuple of 66 floating point numbers that encode the parameters that define the behavior of each ant species. A genetic algorithm (GA) is used to evolve the ant species.

During the simulation of each ant species, a series of behavioral statistics is collected by the program. These characterize aspects such as how often the ants change direction, the distance that they travel, the thickness of their trails and so forth. When the simulation of each ant species ends, some statistics regarding the painting they produced (e.g. its complexity) are also gathered. Video 1 (<http://cdv.dei

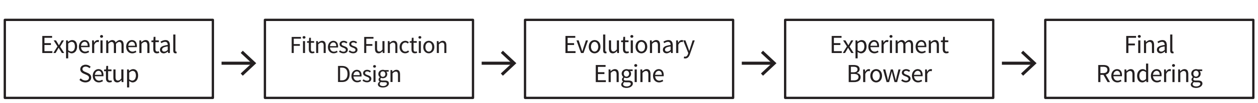

Fig. 2. The pipeline. (© Tiago Martins)

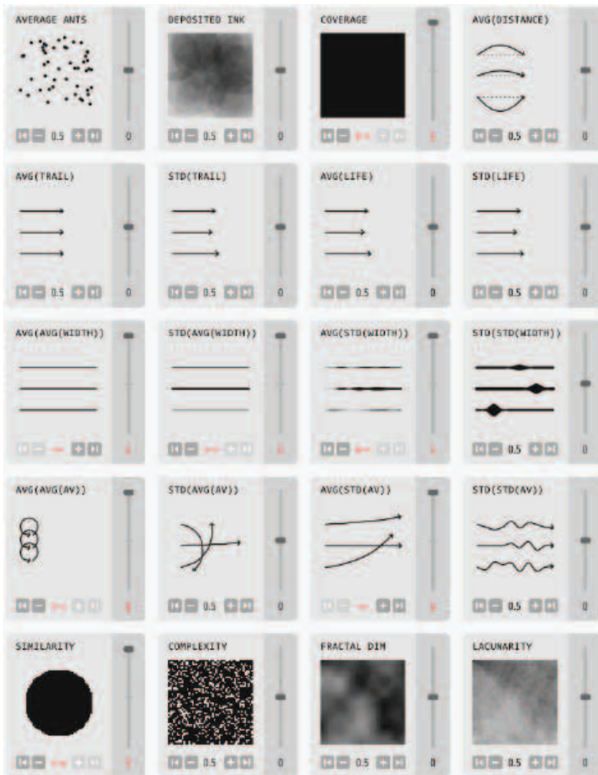

Fig. 3. The interface for fitness function design. A detailed description can be found in Machado et al. [5]. (© Tiago Martins)

.uc.pt/2014/sim.mov>) depicts this simulation, showing the environment and painting canvas (see Machado and Pereira [4] and Machado et al. [5] for further details).

We empower users by giving them the ability to design fitness functions that will guide evolution. For each feature, the user may indicate a weight, which indicates the importance that will be assigned to that feature, and the intention to minimize, maximize or make the feature match a target value. Figure 3 depicts the interface, while video 2 (<http://cdv.dei.uc.pt/2014/int.mov>) illustrates the responsiveness of the icons.

EXPERIMENTATION

We are interested primarily in answering the following set of questions:

1. Are users able to express intentions through the design of fitness functions?
2. Are the fitness functions able to guide the evolutionary algorithm toward regions of the search space consistent with those intentions?
3. Are users able to evolve imagery that they are unable to create by other means?
4. Does the tool confer a sense of authorship on users?
5. Does the tool offer users a new range of artistic possibilities?

We gave access to the tool to 16 users. After a short explanation of the tool and workflow, they designed several fitness functions, conducted multiple evolutionary runs and curated the outcomes. When the sessions were over, they filled out a questionnaire, which was followed by short interviews. The main conclusion that we draw from this experiment is that the tool's results depend deeply on users' commitment: in other words, on their willingness to learn how to use the tool.

To answer the first question: While all users were able to express goals using the interface, some of these goals were rather simplistic (e.g. to create renderings that match the original image using straight lines). Only some of the users—those who made the effort to understand the consequences of their choices and the philosophy of the tool—were able to convey complex design goals (e.g. to create renderings that approximate the original using a combination of thin and thick lines of constant width, producing long spirals of varying curvature and covering the entire canvas).

Concerning the second question, the evolutionary algorithm is always able to find local or global optima for the designed fitness functions. However, the optima are not necessarily consistent with the user's original intentions. On the positive side, the system frequently finds varied and surprising ways to match the user's intentions; this results from the stochastic nature of the algorithm, from the multiple objectives encoded in each fitness function and from the expressive power of the painting algorithm, which allows a large diversity of outcomes.

The answer to the third question is clear: It is unfeasible to recreate the results attained in our experiments by using meta-level interactive evolution, by direct manipulation of the code or by interactive evolution. Even simple goals become extremely hard to accomplish by direct manipulation of the parameters governing the ants' behavior. Using the generative version would necessitate setting 66 separate parameters, whose impact on the outcome is hard to predict by hand. The fitness functions use up to 20 features (usually users employ a significantly lower number, however) that are directly associated with specific image properties and hence are "closer" to the space in which users form their intentions. Interactive evolution allows the discovery of interesting images, but these rarely match a specific intention.

The answer to the fourth question is probably the most subjective, because the sense of authorship depends heavily on the user. Having said that, the most experienced users, the ones who put the most effort into the design of the fitness functions, tended to be the ones with a higher sense of authorship.

The fifth question can also be answered positively. The system produces a wide variety of images that would be difficult to execute using conventional drawing tools. More important, some of these images are difficult to conceive. Additionally, as previously mentioned, the system often surprises users (even the most experienced ones) with a wide variety of unexpected outcomes that match user intentions, which can be seen as an advantage or a drawback.

To convey the type of results that can be attained by a committed user, we focused on the interaction of one of the 16 users: a graphic designer who was not familiar with either

Fig. 4. Examples of phenotypes resulting from each of the 15 fitness functions. (© Tiago Martins)

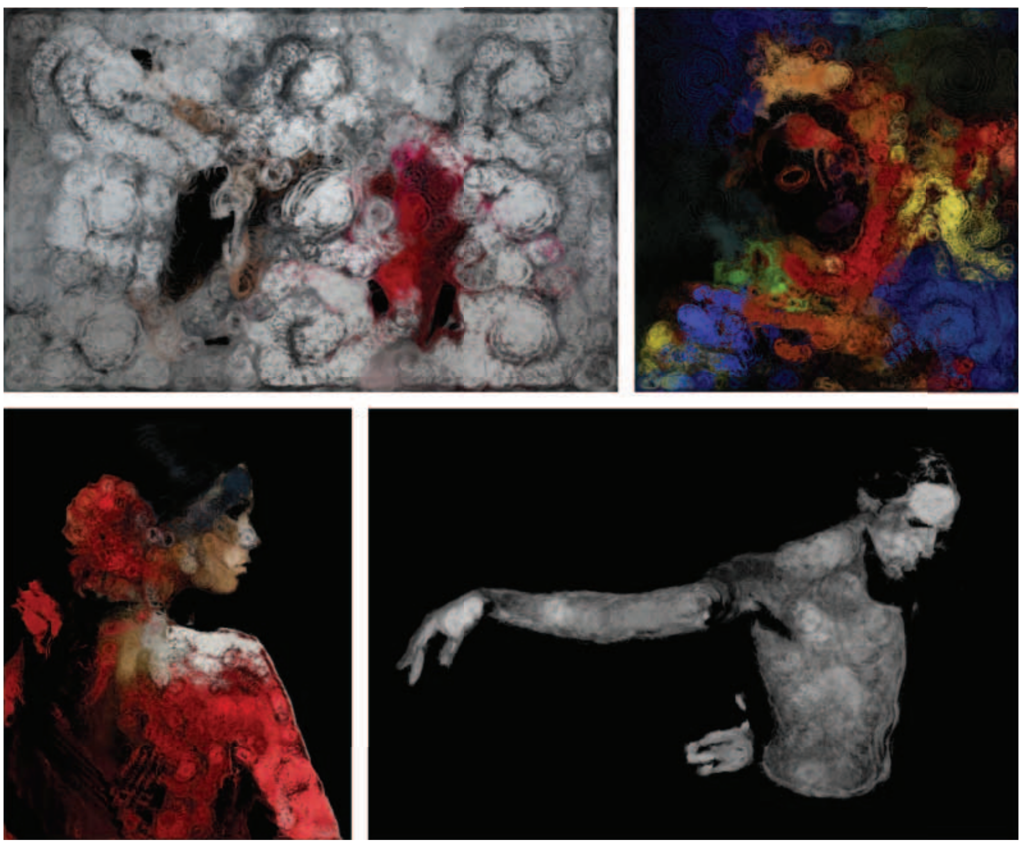

Fig. 5. The same genotype applied to different images. (© Tiago Martins)

the inner workings of the system or the interface. Figure 4 summarizes the results of 15 different experiments and corresponding fitness functions by depicting the fittest individual in each experiment. The results highlight the diversity of the outcomes. Although we present only one image per fitness function, the system creates hundreds of images that correspond to local or global optima for each function (Fig. 4).

Over time, users can compile ant species that match their preferences and intentions and then use their species with novel source images. In Fig. 5, we show the results of applying the genotype corresponding to the rightmost image in the bottom row of Fig. 4 to different input images. These results indicate that although the ant species are sensitive to the environment (i.e. the input image), the characteristics of the painting (e.g. the curviness of its lines) are inherent to the specific ant species. Therefore, applying the same ant species to different input images tends to result in ant paintings with similar aesthetic qualities (Fig. 5).

The final rendering interface gives users an additional degree of control, allowing them to fine-tune rendering options and explore alternative rendering modes. Figure 6 illustrates how various combinations of parameters affect the visual outcome. Since the details of the rendering are difficult to perceive in a small format, video 3 (<http://cdv.dei.uc.pt/2014/ren.mov>) illustrates the final rendering process (Fig. 6).

CONCLUSIONS AND FUTURE WORK

We have described a novel paradigm, meta-level interactive evolution, as well as its application for the generation of non-photorealistic image renderings. This approach empowers users by giving them the ability to design fitness functions and freeing them from the need to perform individual assessments. Our experimental results indicate that committed users are able to evolve images that are consistent with their aesthetic and artistic intentions.

For each of the designed fitness functions, the algorithm creates a wide diversity of outcomes that address those intentions in surprising ways. Moreover, experienced users gain a sense of authorship of the *idea* that seeds the generative process. Finally, we argue that the proposed tool creates a new range of artistic possibilities for the user, in the sense that it produces images that would be hard to execute or even imagine. In the future, we intend to explore novelty search mechanisms and clustering techniques to summarize the results of the evolutionary runs, simplifying the curation stage of the process, and to further increase the diversity of results consistent with users' intentions.

Fig. 6. The same genotype rendered with different final rendering options. (© Tiago Martins)

Acknowledgments

This research is partially funded by the Portuguese Foundation for Science and Technology, project PTDC/EIA–EIA/115667/2009 and grant SFRH/BD/105506/2014, and by project ConCreTe. Project ConCreTe acknowledges the financial support of the Future and Emerging Technologies (FET) program within the Seventh Framework Programme for Research of the European Commission under FET grant number 611733.

References and Notes

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

2 See e.g. P. Machado et al., "On the Development of Evolutionary Artificial Artists," *Computers & Graphics* 31, no. 6 (2007) 818–26; S. Baluja, D. Pomerlau and J. Todd, "Towards Automated Artificial Evolution for Computer-Generated Images," *Connection Science* 6, no. 2 (1994) 325–54; and T. Kowaliw, A. Dorin and J. McCormack, "An Empirical Exploration of a Definition of Creative Novelty for Generative Art," in K.B. Korb, M. Randall and T. Hendtlass, eds., *Artificial Life: Borrowing from Biology*, *Lecture Notes in Computer Science*, vol. 5865 (Springer, 2009) pp. 1–10.

3 G. Greenfield and P. Machado, "Ant and Ant Colony Inspired A-Life Visual Art," *Artificial Life* (2015) (forthcoming). Gary Greenfield and Penousal Machado, "Ant- and Ant-Colony-Inspired ALife Visual Art," *Artificial Life* 21, No. 3 (August 2015) 293–306. <DOI=http://dx.doi.org/10.1162/ARTL_a_00170>.

4 P. Machado and L. Pereira, "Photogrowth: Non-Photorealistic Renderings through Ant Paintings," *Genetic and Evolutionary Computation Conference* (Philadelphia: ACM, 2012) pp. 233–40.

5 P. Machado et al., "An Interface for Fitness Function Design," in J. Romero, J. McDermott and J. Correia, eds., *Evolutionary and Biologically Inspired Music, Sound, Art and Design—Third International Conference*, EvoMUSART 2014, Granada, Spain, April 23–25, 2014. *Proceedings: Lecture Notes in Computer Science*, vol. 8601 (Springer, 2014).

PENOUSAL MACHADO *is a teacher at the University of Coimbra, the coordinator of the Cognitive and Media Systems group and the scientific director of the Computational Design and Visualization Lab. and of the Centre of Informatics and Systems of the University of Coimbra. His research interests include nature-inspired computation, artificial intelligence, computational creativity, computational art and design. He is the author of more than 100 refereed journal and conference papers in these areas, chair of several scientific events, and member of the Programme Committee of numerous conferences. He is also the recipient of scientific awards, including the prestigious award for Excellence and Merit in Artificial Intelligence granted by the Portuguese Association for Artificial Intelligence. Recently, his work was featured in Wired UK and included in the "Talk to me" exhibition of the New York Museum of Modern Art.*

TIAGO MARTINS *is a cross-media designer by training who smoothly mutated into an algorist while interacting with computer scientists during his studies. He is particularly interested in the convergence of multiple disciplines such as graphic design, art and computation. He creates his own tools to develop experimental artifacts and design solutions through the exploration of the expressiveness of computational, parametric, generative and emergent processes. He is a researcher of the Computational Design and Visualization Lab. of the Centre of Informatics and Systems of the University of Coimbra, where he is enrolled in the Doctoral Program of Information Science and Technology, working on his thesis project "Author Tools in Design."*

HUGO AMARO *is a researcher at the Computational Design and Visualization Lab. of the University of Coimbra, with previous work experience in multiple areas, including evolutionary computation, machine learning, classification, data analysis, web applications and databases. He is interested in evolutionary computation, artificial intelligence and computer-generated graphics.*

PEDRO H. ABREU *is an assistant professor at the Department of Informatics Engineering, University of Coimbra, Portugal. He received his Informatics Engineering Degree from the University of Porto in 2006 and a Ph.D. in Soccer Teams Modeling from the University of Porto in 2011. His interests include (but are not limited to) medical informatics and personal healthcare systems applied to oncological diseases. He is the author of more than 40 publications in international conferences and journals.*