

Transgenic Visual-and-Sound Compositions

Artemis Moroni¹

Rafael Bocaletto Maiolla

Jônatas Manzolli

¹Robotics and Computer Vision Division

¹Renato Archer Research Center

Rod. Dom Pedro I, km 143,6

13069-901, Campinas, SP, Brazil

¹E-mail: Artemis.Moroni@cenpra.gov.br

Abstract

Here we present some results obtained in Shapes, an evolutionary environment applied to artistic production in visual domain. In Shapes environment the user can interactively create and evolve visual compositions by using random variations of pre-programmed objects. Pixels of these compositions can be used in another evolutionary environment, ArTVox, to direct a sound sequence creating, in this way, “transgenic” visual-and-sound compositions. The genetic algorithm is described. Both environments, Shapes and ArTVox, were developed by the authors, in Java.

Introduction

Recently, a new generation of computing researchers discovered that by using simulated evolution techniques it is relatively easy to obtain novelty, complex novelty. On the other hand, it is correspondingly difficult to direct the flow that novelty takes. Often, the novelty is not useful. The challenge faced by the designers of evolutionary composition systems is how develop structures and knowledge into the evolutionary loop, to lead the evolutionary search to an interesting region.

Evolutionary techniques have been used for searching large spaces using simulated systems of variation and selection. The loop, in an evolutionary system, is a rather simple one: it *generates, tests and repeats*. Such systems maintain a population of potential solutions; they have a selection process and some “genetic operators”, typically mathematical functions that simulate *crossover* and *mutation*. Basically, a population is generated; the individuals of the population are tested according to certain criteria, and the best are kept. The process is repeated by generating a new population of individuals – or things, or solutions – based on the old ones. This loop continues until the results are satisfactory according to the criteria being used. The effective challenge is to specify what “to generate” and “to test” mean [1, 2].

All evolutionary approaches do, however, share many features. They are all based, like the diagram in Figure 1, on the general framework provided by J. H. Holland’s original genetic algorithm (GA) [3]. In nearly every case, new populations of potential solutions to problems are created, generation after generation, through three main processes: 1. by making sure that better solutions to the problem will prevail over time, more copies of currently better solutions are put into the next generation; 2. by introducing new solutions into the population; that is, a low level of mutation operates on all acts of reproduction, so that some offspring will have characteristics changed randomly; 3. by employing sexual crossover to combine good components between solutions; that is the “genes” of the parents are mixed to

form offspring with aspects of both. With these three processes taking place, the evolutionary loop can efficiently explore many points of the solution space in parallel, and good solutions can often be found very quickly.

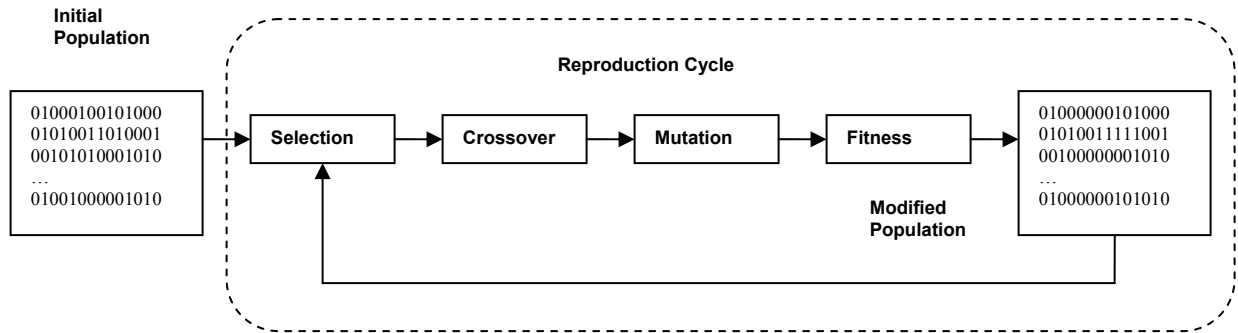


Figure 1: *The simple genetic algorithm.*

Figure 1 shows the diagram of the simple genetic algorithm. The population is initialized with random solutions. Parent solutions are chosen randomly from the “parent population”. Then, crossover and mutation operators are applied, generating a modified population, and the evolutionary cycle begins again. In the next section, we will detail some of the processes involved in this kind of populational search.

Evolution Applied to Visual Domain

Both biological and simulated evolutions involve the basic concepts of genotype and phenotype, and the processes of expression, selection, and reproduction with variation. The *genotype* is the genetic information that codes for the creation of an individual. The *phenotype* is the individual itself, or the form that results from the developmental rules and the genotype. *Expression* is the process by which the phenotype is generated from the genotype. For example, expression can be a biological developmental process that reads and executes the information from DNA strands, or a set of procedural rules that use a set of genetic parameters to create a simulated structure. Usually, there is a significant amplification of information between the genotype and phenotype [4].

Here, we applied this approach to the generation and evolution of abstract pictures [5]. By *abstract picture* we denote that kind of picture that does not depict objects in the natural world, but instead uses color and form in a non-representational way. For a system to automatically produce “aesthetically pleasing” images, two fundamental components must interact. The first component encompasses the mechanisms to create images. The second component, the hardest in this case, must evaluate the images and choose the next move. Obviously, producing an image is not producing art, as issues of other kind such as aesthetics and expressiveness are involved in the latter. A discussion on aesthetics in this context would imply taking into account artistic variables such as color, composition, proportion, motion, form, shape, to name but a few, and, as a matter of fact, the way they are translated into computational terms. This is a very open question, subject to a flurry of interest and research recently. By this moment, what we have is a new method to generate images by means of genetic algorithms.

Genotype, Phenotype and Expression. In biological systems, genotypes are normally composed of DNA. In simulated evolutions there are many possible representations of genotypes, such as strings of binary digits, sets of procedural parameters, or symbolic expressions. In our development, simulated evolution was applied to the generation and evolution of abstract pictures. Each picture has a chromosome associated to it - its genotype – that could be briefly described as follows:

$$\text{picture} = (\text{object}_1, \text{object}_2, \text{object}_3, \text{object}_4, \text{object}_5, \dots)$$

In this chromosome, each object has its specific attributes, all randomly initialized, corresponding to the fields of the computational geometric function that is used to generate the object. The dimensions of the picture are pre-defined, and the background is treated as another object. The phenotype is the generated picture, and expression is the necessary process, the sequence of computational steps, to show the picture.

Reproduction. *Reproduction* is the process by which new genotypes are generated from an existing genotype or genotypes. For evolution to progress there must be variation or mutations in new genotypes, with some frequency. Mutations are usually probabilistic. *Sexual combination* can allow genetic material of more than one parent to be mixed together in some way to create new genotypes. This permits features to evolve independently and later be combined into an individual genotype. In Figure 2, we show two pictures, which are the parents of the pictures presented in Figure 3. Figure 2.a has only one type of object that we call *Sun*, while Figure 2.b has only *Triangle* objects. The pictures were created in our most recent visual environment, Shapes, a work in progress, developed in Java. This environment has a library of routines for the generation of different types of objects. Such objects have parametric features, which are instantiated with random values when they are generated. There are so many possibilities of combination that it is almost impossible that two identical objects occur.

Figure 3 depicts the offspring obtained from the compositions presented in Figure 2. Note that, if the parent compositions have only objects of type Sun or Triangle, in the offspring we can see compositions with combinations of both objects. These compositions can be evaluated by the user and become the parents of the next generation. They can even be edited in our environment. This would be equivalent to manipulate them genetically before reproduction. In Figure 4 we present some of the objects that are ready to be used in the compositions, with variation. Most of them were inspired by Kandinsky Composition VIII 1923 [6].

Crossover and Mutation. Offspring are generated by the use of the crossover operator, which randomly allocates genes from each parent's genotype to each offspring's genotype. Given two parents:

$$\begin{aligned} \text{picture}_1 &= (\text{object}_{11}, \text{object}_{21}, \text{object}_{31}, \text{object}_{41}, \text{object}_{51}) \\ \text{picture}_2 &= (\text{object}_{12}, \text{object}_{22}, \text{object}_{32}, \text{object}_{42}, \text{object}_{52}) \end{aligned}$$

and a random crossover point on position 3, for example, the two offspring generated by the simple GA would be:

$$\begin{aligned} \text{offspring}_1 &= (\text{object}_{11}, \text{object}_{21}, \text{object}_{31}, \text{object}_{42}, \text{object}_{52}) \\ \text{offspring}_2 &= (\text{object}_{12}, \text{object}_{22}, \text{object}_{32}, \text{object}_{41}, \text{object}_{51}) \end{aligned}$$

Crossover is used about 70% of the time to generate offspring; the remaining 30% are simply clones of their parents. Mutation is then applied to offspring, usually with a low probability. In our case, each object has a specific mutation operator. Once the object is chosen, one of its characteristics is also randomly selected and modified. In Figure 5, we can see some mutations applied in the eclipse object. The pictures are presented as they were obtained, but they are not necessarily subsequent. Sometimes the result of the mutation operator is almost imperceptible.

Evaluation. The *evaluation function* is the only means of judging the quality of the evolved solutions. Remember that evolutionary systems maintain a population of potential solutions. In our case, each composition is a solution among all possible compositions with our objects. The interaction between the

evaluation of the solutions and the variation operators (crossover, mutation) to a large part determines the effectiveness of the overall search. Carefully designing suitable evaluation functions requires considerable insight into how the evolutionary algorithm is likely to proceed [7].




Figure 2: *Abstract compositions generated in Shapes.*

But here we find a huge problem. Our function of evaluation is an aesthetic function. Aesthetics is always assumed to be subjective, but aesthetics choices can reflect the opinion of a person, a group of people, or a standard observer that represents some kind of universal aesthetic opinion [8]. How can we measure aesthetics? There are efforts in evolutionary computation that are looking at the question of the aesthetic judgment [9, 10], but like in other evolutionary environments applied to artistic production in the visual domain, here, the aesthetic judgment was left to the user, who observes and evaluates each composition. According to his/her own judgment, the user assigns a score to each composition, in a range of 0 to 10. These scores are then used to determine how many copies of the chromosome of the composition are placed in a temporary area, often termed the “mating pool”, where the crossover and the mutation operators are applied. The higher the fitness of the composition, the more copies are made.

In Figure 6 and 7, we present some visual compositions created in our environment.

Evolution Applied to Sound Domain

Within the last decades, the decreasing cost of technology has fired a revival of interest in the practical association of color-music, or visual and sound domains. Can a chromosome generate a visual *and* sound composition? We have two other previous evolutionary applications, Vox Populi [11, 12], in sound domain, and ArtLab, in visual domain, which later evolved to ArTVox. ArTVox has features for the automatic generation of abstract compositions which can be translated to sound trajectories. We can also use image pixels to generate a sound sequence, in ArTVox. This does not mean that image and sound are strongly correlated, but it is possible.

New problems arise: how can visual and sound features be associated? In his famous book, Kandinsky [13] establishes a parallel between color, form and music. These associations gave rise to two interface features in the ArTVox environment. The first, when active, associates objects with instruments. The second associates color with instruments. The set of relationships applied are those suggested by Kandinsky, but the features of the interface are also open to the user, who can decide his own mappings. After all, according to Kandinsky, “any parallel between color and music can only be relative. Just as a violin can give various shades of tone, so yellow has shades, which can be expressed by various instruments.”

Nowadays, computers have features to create works in both domains. Since “transgenic” organisms are defined as those containing one or more genes that came artificially from other species, carrying attributes from one environment - visual - to the other - sound, we have transgenic compositions!




Figure 3: *The offspring obtained from the compositions in Figure 4.*




Figure 4: *These are some of the objects that are ready to be used in the visual compositions.*



Figure 5: *The mutation applied to the eclipse object.*





Figure 6: *Above, the parents.*

Conclusions


Programs using evolutionary algorithms can generate unexpected structures in different domains: visual, graphic, or musical arts, or in chemistry, engineering, or robotics. The main question surrounding visual and sound compositions does not concern the number of elements, but rather the balanced use of a relatively small family of elements. It seems that if in the past technological problems arose in building engines that worked in the visual and sound domain, the current issue is of a more conceptual nature; a new aesthetical era begins.




(a)




(b)




(c)




(d)




(e)



(f)



(g)



(h)

Figure 7: *And now, the descendants.*

Acknowledgements

We would like to thank Leonardo Laface de Almeida and Daniel Gurian Domingues, who worked in the development of JaVox and ArTVox environments. We would like to thank PIBIC/CNPq program and CenPRA, for making this research possible. This research work is part of the AURAL project, supported by FAPESP process 05/56186-9.

References

- [1] A. Moroni, J. Manzolli, F. J. Von Zuben, R. Gudwin, *Vox Populi: An Interactive Evolutionary System for Algorithmic Music Composition*, Leonardo Music Journal, 10, pp. 49-54, 2000.
- [2] P. M. Todd, G. M. Werner, *Frankensteinian Methods for Evolutionary Music Composition* in Griffith, N. & Todd, P. M. (eds) *Musical Networks: Parallel Distributed Perception and Performance*, Cambridge: The MIT Press, pp. 313-339, 1999.
- [3] J. H. Holland, *Genetic Algorithms*, Scientific American, July, 1992.
- [4] K. Sims, *Evolving Three-Dimensional Morphology and Behaviour*, in Bentley, P. (ed.), *Evolutionary Design by Computers*, San Francisco: Morgan Kaufmann, pp. 297-321, 1999.
- [5] A. Moroni, F. J. Von Zuben, J. Manzolli, *ArTbitration: Human-Machine Interaction in Artistic Domains*, Leonardo, 35(2), pp. 185 – 188, 2002.
- [6] W. Kandinsky, *Concerning the Spiritual in Art*. Dover Publications, 1977.
- [7] Z. Michalewicz, D. B. Fogel, *How to Solve It: Modern Heuristics*. New York: Springer, 1998.
- [8] F. Hoenig, *Defining Computational Aesthetics*, Computational Aesthetics 2005, Eurographics Workshop on Computational Aesthetics in Graphics, Visualization and Imaging, pp. 13-16, 2005.
- [9] S. Baluja, D. Pomerleau, T. Jochem, *Towards Automated Artificial Evolution for Computer-generated Images* in Griffith, N. & Todd, P. M. (eds) *Musical Networks: Parallel Distributed Perception and Performance*, The MIT Press, 341 – 370, 1999.
- [10] P. Machado, J. Romero, B. Manaris, A. Santos, & A. Cardoso. *Power to the critics - a framework for the development of artificial art critics* in IJCAI'2003 Workshop on Creative Systems, Acapulco, Mexico, 2003.
- [11] A. Moroni, J. Manzolli, F. J. Von Zuben, R. Gudwin, *Vox Populi: Evolutionary Computation for Music Evolution*, in Bentley, P. and Corne, D. (eds.) *Creative Evolutionary Systems*, San Francisco, USA: Morgan Kaufmann, pp. 205 - 221, 2002.
- [12] A. Moroni, J. Manzolli, F. J. Von Zuben, *ArTbitrating JaVox: Evolution Applied to Visual and Sound Composition*. in Brunet, P., Correia, N., Baranoski, G. (eds.) *Ibero-American Symposium in Computer Graphics 2006*, Santiago de Compostela, Eurographics Chapter Proceedings, Eurographics Association, pp. 97 – 108, 2006.
- [13] W. Kandinsky, *Concerning the Spiritual in Art*. Dover Publications, 1977.