NEvAr – The Assessment of an Evolutionary Art Tool

Penousal Machado*; Amílcar Cardoso**
* Instituto Superior de Engenharia de Coimbra, CISUC; Qta. da Nora, 3030 Coimbra, Portugal;
** CISUC, Dep. Eng. Informática, Uni. Coimbra, Polo II; 3030 Coimbra, Portugal;
machado@dei.uc.pt; amilcar@dei.uc.pt

Abstract

The use of Evolutionary Computation approaches to create images has reached a great popularity, leading to the appearance of a new art form – Evolutionary Art – and to the proliferation of Evolutionary Art Tools. In this paper we present and make an assessment of one of these tools: NEvAr. We also systematise and describe the work methodology currently used to generate images. When working with NEvAr we focus on the reuse of useful individual, which we store in an image database. The size of this database, and the importance of its role, led us to the development of automatic seeding procedures, which we also describe.

1 Introduction

In the past few years, a new AI area has begun to emerge, usually named Creative Reasoning. Several aspects contributed to the growth of interest in the study of computational creativity: artificial creative systems are potentially effective in a wide range of artistic, architectural and engineering domains where conventional problem solving is unlikely to produce useful solutions; their study and development may contribute to the overall understanding of the mechanisms behind human creativity; in some ways, the study of creativity can be viewed as the next natural step in AI, considering that we already can build systems capable of solving tasks requiring intelligence, can we build systems that are able to solve tasks that require creativity?

Models of the human creative process (e.g. Dewey (1910), Guilford (1968), Wallas (1926) and De Bono (1986)) may constitute an important source of inspiration to the development of artificial creative systems. Human creativity, however, isn't the only source of inspiration available. When looking at nature, we can see all living species in a permanent struggle for life. Long term survival is connected with the capability of adapting to environmental changes. The survival of the fittest individuals and the recombination of their genetic material is the key element of the adaptation process. The recombination of "good pieces" of different individuals can give rise to new and better ones. Furthermore, the slight modification of individuals' genetic code, can also increase the quality of the individuals.

Over the time, natural selection was capable of producing an incredible amount (and variety) of solutions, *species*, to a common problem, *survival*. Thus, there is no doubt that the evolutionary process is a way of producing innovative solutions (Goldberg, 1998). Whether these solutions can or cannot be considered creative, is a different question, and the answer depends on the way we define creativity. Therefore, and since no uncontroversial globally accepted definition exists, we can consider this to be an open question. However, our current standing is that these solutions can be considered creative.

In the past few years, two Evolutionary Computation (EC) approaches (Genetic Algorithms (GA) and Genetic Programming (GP)) have been used as a mean to implement computational creativity, resulting in the appearance of a set of new applications in areas such as music and image generation, architecture and design.

GA are the most common EC approach in the musical field, some examples are the works of Horowitz (1994), Ralley (1995), Biles (1994), Jacob (1995). However, and in spite of the numerous applications, Wiggins et al (1999), which have studied the performance of this type of systems, defend that these approaches are not ideal for the simulation of human musical thought. In the field of image generation, GP is the most used approach. Examples of works in this field are: Dawkins (1987), Sims (1991), Todd (1993), Rooke (1996), which resort to GP to evolve images, and Baker (1993), where GP is used to evolve human faces. GP has also been successfully applied in the fields of design (Bentley, 1999; Graf, 1996) and animation (Sims, 1991; Angeline, 1996; Ventrella, 1999).

Due to the difficulty of creating an evaluation function in domains such as image or music generation, most of the above mentioned systems use Interactive Evolution (IE). In IE systems the user evaluates the individuals, thus guiding evolution. In the musical field we can already find several systems that resort to automatic evaluation (e.g. Horner et al (1991), McIntyre (1994), Spector (1994, 1995), Hodgson (1996,1990,1999), Papadopoulos et al (1998)). In image generation, the picture is quite different: as far as we know there has been only one attempt to automate fitness assignment, the work of Baluja et al (1994). However, the results produced by this system, which uses neural networks to evaluate images, were disappointing.

The core subject of this paper is the assessment of NEvAr as a tool. In Section 2 we make a brief overview of the previous work in Evolutionary Art Tools, focusing on the most prominent systems. In Section 3, we introduce NEvAr and describe the used evolutionary model. Section 4 concerns the assessment of NEvAr and the description of the work methodology currently employed. This methodology gives emphasis to the reutilization of good individuals, which are stored in a database. The difficulties of managing an increasingly large database led to the study of seeding procedures, which will be described in Section 5. Finally, in the 6th section we make some overall remarks and draw some conclusions.

2 State of the Art

In the past few years, the use of IE to the generation of images has achieved a great popularity. The source of inspiration of most of these applications can be found in Richard Dawkins book “The Blind Watchmaker”, in which the author suggests the use of a GA to evolve the morphology of virtual organisms, *biomorphs*. In these systems, the evolution is guided by the user accordingly to hers/his aesthetic criteria. This inspired the works of K. Sims (91) and W. Latham et al (92), which can be considered as the first applications of IE in the field of the visual arts, and are usually considered as the most influential works in this area. The success of these approaches has led to the emergence of a new art form, “Evolutionary Art” (EA), and also to the proliferation of IE applications in this field, usually called Evolutionary Art Tools.

In spite of the increasing number of this type of applications, few are the ones that can be compared favourably with the above mentioned works. The vast majority of these applications adds nothing new to these works, and are, frequently, inferior both in terms of potential and results. Moreover, few are the ones that have been thoroughly tested, i.e. in most cases there was no attempt to use them to create art. Therefore, it seems safe to say that the classification of these applications as Evolutionary Art Tools is misguiding, and that few are the applications that deserve this name. In this restricted set, we can include the works of: K. Sims (91), W. Latham and S. Todd (92), S. Rooke (96), Vetrella (99). The description of the characteristics of these systems and the analysis of their potential is clearly beyond the scope of this paper. These systems share many features, most notably: they resort to GP, use IE and have been successful in the generation of visual artworks.

3 NEvAr

NEvAr (Neuro Evolutionary Art) is an evolutionary art tool, inspired in the works of K. Sims (1991) and R. Dawkins (1987). It allows the evolution of populations of images from an initial one, and resorts to IE. In this section, we will make a brief description of the evolutionary model used in NEvAr. NEvAr is in many ways similar to the application developed by K. Sims (91), namely in what concerns the representation of the individuals and the used genetic operators. Therefore we won’t make a description of these aspects.

For the current purpose, it is enough to say that in NEvAr the individuals are represented by trees. The genotype of an individual is a symbolic expression, which is constructed from a lexicon of functions and terminals. In NEvAr, we use a function set composed mainly by simple functions such as arithmetic, trigonometric and logic operations. The interpretation of a genotype results on a phenotype, i.e. an image. All the genetic manipulations (e.g. crossover, mutation) are performed at the genotype level. In Figure 1, we present two images generated with NEvAr and in Figure 2 some images generated by the mutation and crossover of the genetic code of these images.

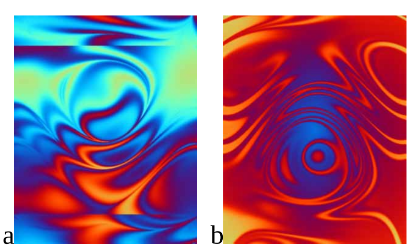
a b

Figure: 1 –Two images created with NEvAr.

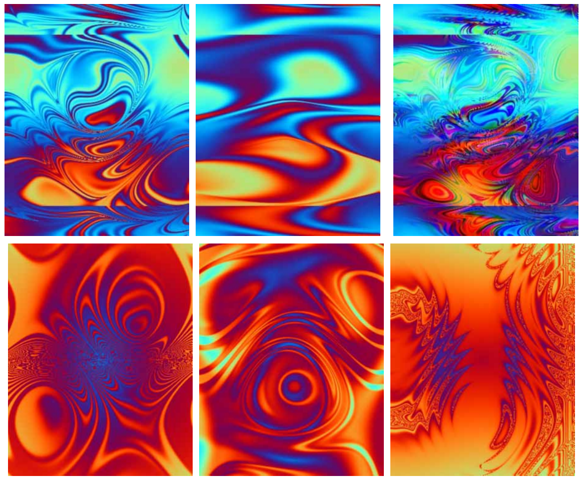

Figure 2: The top row images were created through the mutation of image *a* of Figure 1. The bottom row images result from the crossover of images *a* and *b* of Figure 1.

As the images of Figure 2 show, the mutation and crossover operations can produce interesting and unexpected results. The high plasticity, which is inherent to the used representation and, to some extent, to GP approaches, allows the generation of radically different, yet fit, phenotypes from similar genotypes.

3.1 The Model

In NEvAr, the assignment of fitness is made by the user and, as such, she/he has a key role. The interaction of human and computer poses several problems (e.g. limited population size, limited runs). The fact that NEvAr is an interactive tool has the advantage that a skilled user can guide the evolutionary process in an extremely efficient way. She/he can predict which images are compatible, detect when the evolutionary process is stuck in a local optimum, etc. In other words, the user can change its evaluation criteria according to the context in which the evaluation is taking place.

In the design of NEvAr's model, we took under consideration these idiosyncrasies. In Figure 3, we show the model of NEvAr. From here on, we will call *experiment* to the set of all populations, from the initial to the last, of a particular GP run. NEvAr implements a parallel evolutionary algorithm, in the sense that we can have several different and independent *experiments* running at the same time. It is also asynchronous, which means that we can have an *experiment* that is in population 0 and another one that is in population 100. Additionally, we can transfer individuals between *experiments* (migration).

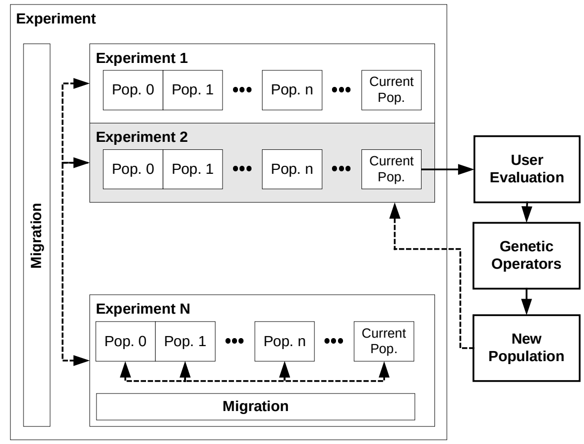

Figure 3: The evolutionary model of NEvAr. The active *experiment* is depicted in grey.

We will illustrate the use of this model through an example. Suppose that the user creates two different *experiments*, *a* and *b*, the initial population of *a* is randomly generated and has size N, and the initial population of *b* has size 0. The user focuses his efforts in *experiment a* and evaluates the individuals of successive populations generated by NEvAr. When the user finds images that she/he likes, she/he adds these images to the current population (in this case the population 0) of *experiment b*. If at a given point the user feels that the evolutionary process would benefit if the next population was generated by the combination of the individuals of the current population with individuals previously transferred to population *b*, she/he adds those individuals to the current population and the evolutionary process continues.

At a certain point, the user decides to focus on *experiment b*, and orders the generation of a new population from the current one (population 0), which is composed, exclusively, by individuals transferred from *a*. Thus, the initial population of *experiment b* is not random, but exclusively composed by fit individuals that were originally generated in other *experiments*. In fact, *experiment b* can be seen as a database of images, which may be used to initialise future *experiments*. We may generalise this approach by organising a gallery of images.

NEvAr also allows the migration within *experiments*. This feature is important due to the limited size of each population, since it allows the revival of images from previous populations. It is also possible to go back to a previous population and change the evaluation of the individuals, which allows the exploration of different evolutionary paths.

4 Working with NEvAr – The artistic point of view

NEvAr is an Evolutionary Art Tool, therefore the main goal is the production of artworks. Its analysis must be performed with this in mind. Like any other tool, NEvAr requires a learning period. To explore all the potential of a tool, the user must know it in detail and develop or learn an appropriate work methodology. The results, and user satisfaction, depend not only on the tool but also on its mastering.

Additionally, the evaluation of the results (images) can only be made by the user that generated them. The evaluation of an art tool can only be made by the artist using it. The key aspect is that the artist must review her/himself in the produced artworks. Thus, the fact that a tool can generate "interesting" images is irrelevant from the artistic point of view. What is really important is that the produced artworks convey the artistic ideas of the artist. In other words the artist must be able to express her/himself through the use of the tool.

The images generated with NEvAr during the early stages of experimentation were clearly disappointing. This failure didn't result from the lack of power of the tool, but from our lack of expertise in its use. Next we will present the work methodology that we currently use to generate images with NEvAr.

4.1 The Process

The creation of an artwork encompasses several stages, such as: genesis of the idea, elaboration of sketches, exploration of the idea, refinement, and artwork execution. The methodology that we propose can be considered, in some way, analogous. It is composed by four main stages: Discovery, Exploration, Selection and Refinement.

These stages can be described, concisely, as follows: the stage of Discovery consists on finding a promising evolutionary path, which, typically, corresponds to evolving a promising set of images from an initial random population (genesis of the idea); in the second stage, Exploration, the "ideas" evolved on the previous stage are used to generate images of high aesthetic value (exploration of the ideas); the Selection stage involves choosing the best produced images; the selected images, when necessary, will be subjected to a process of Refinement, whose goal is the alteration of small details or the correction of imperfections (final execution of the artwork).

Our empirical experience allows us to classify the Discovery stage as the most crucial of the process, and, together with the Exploration stage, the one in which the faculties of the user are more important.

Discovery corresponds to the genesis of the idea, therefore it is inappropriate to approach this stage with pre-conceived ideas regarding the final aspect of the artwork. In other words, it is impossible in practice (yet tempting) to think on an image and use NEvAr to evolve it. This is probably the most important aspect to retain, because it contrasts with what is usually expected in a tool, i.e. that it allow the implementation of an idea. This aspect can be viewed as a weakness, but it is also the distinguishing feature and strength of NEvAr (and other evolutionary art tools). A conventional art tool only plays an important role in the artistic process in stages subsequent to the generation of the idea. Furthermore, the idea frequently determines which tool will be used in its execution, since some are more adequate than others. NEvAr, however, plays a key role in the generation of the idea. Its influence is noticeable through all the artistic process and in its main creative stage. In NEvAr, the artist is no longer responsible for the creation of the idea, she/he is responsible for the recognition of promising concepts. More precisely, the idea results from an evolutionary process, and is created by the artist and the tool, in a (hopefully) symbiotic interaction.

In the Exploration stage the initial idea is already set and we are dealing with images of high aesthetic value. Through the recombination of these images, we explore a space of forms which is smaller than the one explored in the discovery stage, and is therefore more thoroughly searched. The Exploration stage can prolong itself conducting the artist to a point which, at least apparently, has nothing to do with the original one. Like in the Discovery stage, the expertise of the user is determinant to the success of this stage. With the accumulation of experience, the user learns how to distinguish between promising paths and ones that lead nowhere, to predict which combinations of images produce best results, how to manipulate crossover and mutation rates in order to produce best results, etc.

The Selection stage can be divided in two different ones, one that is concurrent with the evolutionary process, and one that is posterior. During the stage of Exploration, the best images (according to the user criteria) are added to a different *experiment*, that works as a gallery. As stated before, NEvAr stores all populations, which allows the review of the evolutionary process and the addition to the gallery of images that were previously neglected. This revision is highly recommended, and a substantial amount of time should separate the generation of the images and its review in order to allow the necessary distance between generation and criticism.

The Refinement process usually occurs separately from the *experiment* that generated the image. The common procedure is to initialise a new *experiment* with the image that we want to refine (i.e. the initial population of this *experiment* will be composed by the image and, in some cases, similar ones). The generation of new populations, from this initial one, allows the exploration of a search space in the vicinity of the image that we want to refine. In Figure 4 we present some images created with NEvAr.

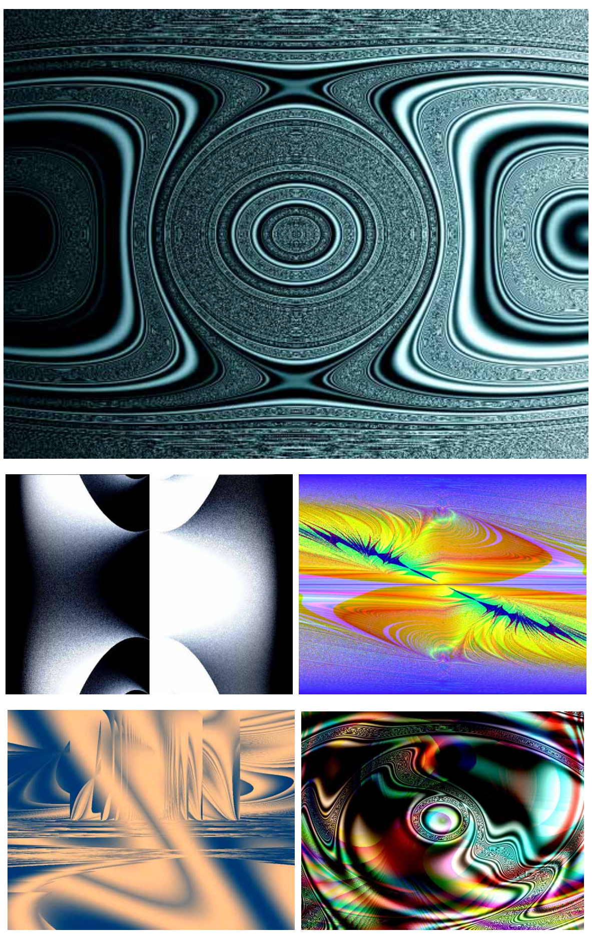

Figure 4: Some examples of images created with NEvAr. Additional images can be found in the CD-ROM accompanying P. Bentley (1999).

4.2 Image Database

One of the weaknesses of EC approaches lies on the fact that they do not have long-term memory (although we can view multiploidy as a limited memory mechanism). The use of several *experiments* allows the accumulation of individuals, which can be used in later experiences. Thus, we can create a Knowledge Base (KB) of images.

The KB has been used mainly in two situations: To initialise new *experiments* and to add individuals to the current population of an *experiment*. The goal of the first form of use is to shorten, or even avoid, the initial stages of the evolutionary process (Discovery and Exploration). The addition of previously generated individuals to the current population usually follows an opportunistic reasoning. There are several situations in which this may be useful, for instance, to avoid a local optimum, or when we find an image whose combination with a previously created one is previewed as promising.

The KB is playing an increasingly important role in the process of image generation, and is currently a priceless feature of the system. The size of the KB is also increasing rapidly and, consequently, the KB is becoming harder to manage and use. This led us to the development of automatic seeding procedures, which we will describe in the following section.

5 Seeding

The development of automatic seeding procedures is part of our ongoing research. In this section we describe our current approaches, which should be considered preliminary.

Our idea is inspired in Case Based Reasoning, and can be described as follows: the user chooses an image, and the seeding procedure selects, from the database, similar ones, to initialise the GP *experiment*. To implement this idea, we need to develop a similarity metric, i.e. a way to compare images. Unfortunately, this task is not trivial.

In our first attempt we used the root mean square error (*rmse*) among two images, which is usually applied to evaluate the error involved in image compression, as similarity metric. Minimum error implied maximal similarity. The similarity between two images, *a* and *b*, was given by the following formula:

$$rmse\ sim_{a,b} = \frac{100}{1+\sqrt{rmse_{a,b}}} \quad (1)$$

The experimental results showed the inappropriateness of this approach. This failure can be easily explained, the goal is to find images that are similar to the eye and not "mathematically" similar images. To illustrate the shortcomings of *rmse* based similarity, we resort to an example: consider two images composed by alternate vertical black and white stripes of one pixel width, one starting with a black stripe and the other with a white one; these images will be almost indistinguishable to the eye, however, the *rmse* among them will be maximal.

Our current idea is to compare images according to their properties. It is a well known fact that image complexity affects the performance of image compression methods, i.e. it is easier to represent compactly a simple image than a complex one. Moreover, some compression methods work better with some types of images than with others.

JPG compression was designed for the representation of natural images. In this type of images, colour transition is usually smooth. Although adequate for these images, the performance of *JPG* severely degrades when dealing with images possessing abrupt colour transitions (e.g. a black and white text image). Fractal Image Compression takes advantage of the self-similarities present on the images and will, therefore, perform better when these similarities are high.

Our previous experience with image compression methods led us to believe that we could use the quality of the compression to develop a similarity metric. For the scope of this paper, we will define compression quality as:

$$\frac{compression\ ratio}{rmse}, \qquad (2)$$

and compression complexity as the inverse.

We use two different compression methods: *jpg* and fractal based. The fractal image compression algorithm makes a quad-tree partitioning of the image. By changing the maximum depth of the tree, we can specify, indirectly, the limits for the error involved in the compression. During compression, the colour information is discarded, the images are converted to greyscale and then compressed.

Let's define *IC* as the compression complexity resulting from the use of the *JPG* method; *PC* as the compression complexity resulting from the application of the fractal based approach. We use two different maximum tree depths, N and N-1, therefore we have PC_1 and PC_2.

To compare two images, *a* and *b*, we start by calculating *IC*, PC_1 and PC_2, for each of them. The similarity between images *a* and *b* is given by the following formula:

$$sim_{a,b} = \frac{1}{1+\sqrt{\left|IC_a - IC_b\right| + \left|PC1_a - PC1_b\right| + \left|PC2_a - PC2_b\right|}} \qquad (3)$$

In Figure 5, we present a subset the images belonging to the database. In Table 1, we present the *IC*, PC_1 and PC_2, measures for each of them as well as the similarity of these individuals with images 9 and 14 of the population.

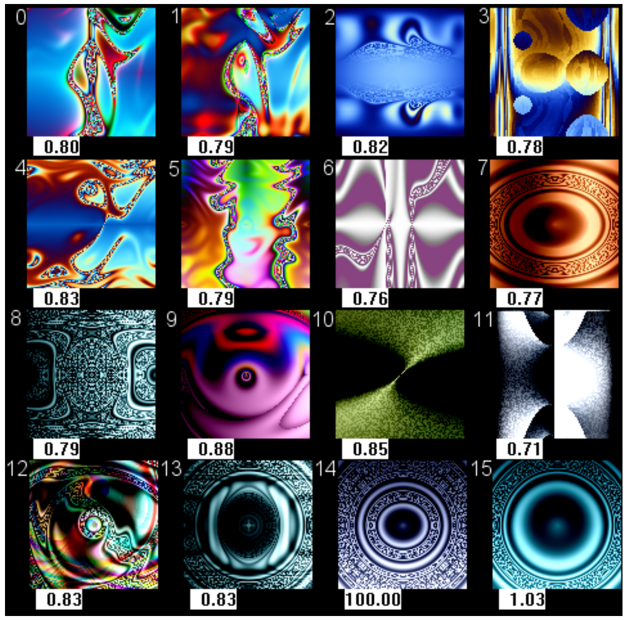

Figure 5: The numbers bellow the images indicate the *rmse* similarity to image 14 (see formula 1). According to this metric, the closest image is image 15, which is good, and the second closest is image 9, which is bad.

Ordering the individuals according to their similarity to image 14, yields the following list: {14, 8, 13, 12, 15, 4, 5, 7, 0, 11, 10, 6, 3, 1, 2, 9}. This ordering seems to be appropriate, the major deficiency being that individual 7 is considered less similar than individuals 4 and 5. The comparison to image 9 gives the ordered list: {9, 2, 1, 3, 6, 10, 11, 0, 7, 5, 4, 15, 12, 13, 14, 8}, which also appears to be approximately correct. Image 9 is characterised by its fluid and organic forms, and so are individuals 2, 3, 6, 0, and, although in a lesser degree, individuals 10 and 11.

Table 1: The *IC*, PC_1 and PC_2 measures for each of the images presented in Figure 5, and the similarity among these images and individuals 14 and 9 of the same figure.

Image	CI	CP1	CP2	Similarity to 14	Similarity to 9
0	5.053	19.228	5.957	10.397	17.500
1	4.455	10.503	4.646	9.790	22.703
2	2.926	5.518	2.403	9.365	37.261
3	4.085	11.256	5.957	9.879	21.529
4	6.401	21.357	7.057	10.697	16.189
5	5.965	21.663	6.504	10.650	16.365
6	4.694	13.395	4.988	9.976	20.486
7	5.744	19.373	6.795	10.503	16.981
8	12.125	91.074	25.331	16.948	8.349
9	2.399	3.413	2.200	9.239	100.000
10	4.593	12.839	5.883	9.989	20.359
11	5.113	14.434	6.244	10.129	19.170
12	8.736	42.895	13.636	13.765	11.673
13	7.978	45.669	13.523	14.062	11.506
14	11.518	71.164	21.835	100.000	9.239
15	6.891	34.791	10.861	12.181	13.032

The initialisation of a new population is based on the similarity to an image chosen by the user. The seeding procedure uses the similarity values as fitness, and roulette wheel selection for choosing which images will become part of the initial population. We do not allow the repetition of images.

Although the experimental results are still preliminary, the seeding procedure based on compression quality seems to produce good results, indicating that the comparison of images based on their characteristics, namely on their complexity, is appropriate. This, of course, suggests that taking into consideration other types of features of the images (e.g. edges, colour, outline, etc.) may also prove useful.

It is also important to notice that the compression methods described can be applied to any image. Therefore, we can compare the database images with ones that where not generated with NEvAr. We still haven't explored this possibility, nevertheless we believe it can produce interesting results.

6 Conclusions and Further Work

From the artistic point of view, we consider NEvAr to be a tool with great potential. Through the use of NEvAr, the artist is no longer responsible for the generation of the idea, which results from an evolutionary process and from the interaction of artist and tool. Thus, the use of NEvAr implies a change to the artistic and creative process. It is important to notice that, in spite of these changes, the artworks obey the aesthetic and artistic principles of the artist. The use of NEvAr implies an abdication of control; however, this lack of control isn't necessarily negative. The artist can express her/himself through the use of the tool and review her/himself in the works created.

One of the erroneous conceptions about evolutionary art tools is that the generation capabilities of a system are deeply connected with the used primitives. Our experience with NEvAr shows that this is wrong. What is necessary is a set of "basic" primitives that can be combined in a powerful way.

The inclusion of a long term memory mechanism is extremely important, since it allows the reuse of previously generated ideas. It is also the basis for the inclusion of Case Based Reasoning mechanisms in NEvAr. Preliminary experiments indicate that the inclusion of these mechanisms can create interesting results and further extend the capabilities of our system.

Interactive evolution proved to be a very powerful technique. This can be explained by the fact that the user can use other criteria besides fitness to evaluate the individuals, and thus guide the evolutionary algorithm more efficiently. We are currently studying ways to automate fitness assignment. Our initial idea was to train a neural network and use it to automate this task. We currently feel that full automation is not attainable on short term. Our current idea is to use neural networks (as well as other techniques) as a filter that eliminates individuals that are clearly undesirable.

Acknowledgements

This work was partially funded by the Portuguese Ministry of Science and Technology, under Program PRAXIS XXI.

References

P. J. Angeline. Evolving Fractal Movies. *Genetic Programming: Proceedings of the First Annual Conference*, MIT Press, Stanford University, CA, USA, pp. 503-511, 1996.

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

S. Baluja, D. Pomerlau, and J. Todd. Towards Automated Artificial Evolution for Computer-Generated Images. *Connection Science*, 6, 325-354, 1994.

P. Bentley, (Ed.) *Evolutionary Design by Computers*, Morgan Kaufman, 1999.

J. Biles A genetic algorithm for generating jazz solos. *International Computer Music Conference*, Aarhus, Denmark, 1994.

R. Dawkins, R. *The Blind Watchmaker*, W.W. Norton & Company, Inc., New York, 1987.

E. De Bono. *El pensamiento lateral. Manual de la creatividad*. Barcelona, Spain: Paidós (in Spanish), 1986.

J. Dewey. *How we think*. Boston: D. C. Heath, 1919.

D. Goldberg. The design of innovation: lessons from genetic algorithms, lessons for the real world, *IlliGAL Report Nº 98004, Department of General Engineering*, University of Illinois at Urbana-Champain, 1998.

J. Graf, and W. Banzhaf. Interactive Evolution for Simulated Natural Evolution. *Artificial Evolution*, Vol. 1063, Springer Verlag, pp. 259-272, 1996.

J. P. Guilford. *The Nature of Human Intelligence*. McGraw-Hill, New York, 1967.

P. Hodgson. Understanding Computing, Cognition, and Creativity. *University of the West of England*, 1990.

P. Hodgson. *Gene-e Evolutionary Dance Music Program*, 1996.

P. Hodgson. Understanding Computing, Cognition, and Creativity. *AISB Symposium on Musical Creativity*, Edinburgh, UK, 1999.

A. Horner, D. Goldberg. Genetic algorithms and computer-assisted composition. Genetic Algorithms and Their Applications: Proceedings of the *Fourth International Conference on Genetic Algorithms*, pp. 427-441, 1991.

D. Horowitz. Generating Rythms with Genetic Algorithms. *International Computer Music Conference*, Aarhus, Denmark, 1994.

B. L. Jacob. Composing with Genetic Algorithms. *International Computer Music Conference*, 1995.

W. Latham, and S. Todd. *Evolutionary Art and Computers*, Academic Press, Winchester, UK, 1992.

R. A. McIntyre. Bach in a Box: The Evolution of Four-Part Barroque Harmony Using Genetic Algorithms. In *IEEE Conference on Evolutionary Computation*, 1994.

G. Papadopoulos, and G. Wiggins. AI Methods for algorithmic composition: A Survey, a Critical View and Future Prospects. *AISB Symposium on Musical Creativity*, Edinburgh, UK, 1999.

D. Ralley. Genetic algorithm as a tool for melodic development. *International Computer Music Conference*, 1995.

S. Rooke. The Evolutionary Art of Steven Rooke, *http://www.azstarnet.com/~srooke/*, 1996.

K. Sims. Artificial Evolution for Computer Graphics. *ACM Computer Graphics*, 25, 319--328, 1991.

L. Spector, and A. Alpern. Criticism, culture, and the automatic generation of artworks. *Proceedings of Twelfth National Conference on Artificial Intelligence*, AAAI Press/MIT Press, Seattle, Washington, USA, pp. 3--8, 1994.

L. Spector, and A. Alpern. Induction and Recapitulation of Deep Musical Structure. *International Joint Conference on Artificial Intelligence*, IJCAI'95 Workshop on Music and AI, Montreal, Canada, 1995.

J. Ventrella. Animated Artificial Life. *Virtual Worlds - Synthetic Universes, Digital Life, and Complexity*, Heudin, J. C. (ed.) Perseus Books, pp. 67-94, 1999.

G. Wallas. *The art of thought*. Nova Iorque: Harcourt Brace, 1926.

G. Wiggins et al. Evolutionary Methods for Musical Composition. *International Journal of Computing Anticipatory Systems*, 1999.