

Technical Section

# On the development of evolutionary artificial artists

Penousal Machado<sup>a,\*</sup>, Juan Romero<sup>b</sup>, Antonino Santos<sup>b</sup>,  
Amílcar Cardoso<sup>a</sup>, Alejandro Pazos<sup>b</sup>

<sup>a</sup>CISUC, Department of Informatics Engineering, University of Coimbra, 3030 Coimbra, Portugal

<sup>b</sup>RNASA Laboratory, Faculty of Computer Science, University of Coruña, Spain

## Abstract

The creation and the evaluation of aesthetic artifacts are tasks related to design, music and art, which are highly interesting from the computational point of view. Nowadays, Artificial Intelligence systems face the challenge of performing tasks that are typically human, highly subjective, and eventually social. The present paper introduces an architecture which is capable of evaluating aesthetic characteristics of artifacts and of creating artifacts that obey certain aesthetic properties. The development methodology and motivation, as well as the results achieved by the various components of the architecture, are described. The potential contributions of this type of systems in the context of digital art are also considered.

© 2007 Elsevier Ltd. All rights reserved.

*Keywords:* Artificial art; Evolutionary computation; Artificial intelligence; Digital art

## 1. Introduction

We live in the digital art era, where machines replace brushes as tools for the expression of the artists' creativity. With the globalization of Information and Communication Technologies, most of the world has the technical resources required to create and broadcast pictures, designs, music, ideas, etc. Some of these artifacts may be considered works of art by other people or even by the society. The aim of our research is to allow autonomous computer systems to take an active part in the creation of aesthetic artifacts. One of the possible approaches to accomplish this goal is the exploration of techniques which allow the creation of art works according to user preferences, or which provide their own aesthetic evaluation, or, ideally, both.

The desire to use computational methods in order to build artistic creation systems has existed for almost 150 years, since Ada Byron dreamt of creating a computer with musical abilities [1]. Several techniques were used from that point onwards in order to study and develop artificial artistic systems (applied to music, design, visual art, poetry, etc.), among which we highlight those related to Artificial

Intelligence (AI), for example, the approaches based on mathematical models [2], knowledge [3], neural networks [4], and generative approaches [5].

One of the inherent difficulties of the design of these systems is the lack of knowledge about the process ruling human artistic creation. For this reason, the hand-coding of algorithms which model human creativity (see e.g. Refs. [6,7]) will probably be a difficult task. A possible alternative is the use of adaptive techniques that explore a solution space. In recent years, biology-inspired approaches, namely evolutionary computation techniques, have become increasingly popular in artistic and aesthetic production domains. A thorough survey on the application of biological-inspired techniques to visual art is presented by Lewis [8].

Through time, natural selection gave rise to a huge variety of species adapted to their environment. From Holland's work [9] onwards evolution has also become a source of inspiration for AI. Evolutionary Computation imitates the fundamental mechanisms of evolution: *selection*, which guarantees that the most apt individuals have greater probabilities to survive and reproduce; *reproduction*, which ensures the inheritance of the parents' genetic material, as well as variation, which allows evolution. Thus, Evolutionary Computation transforms Darwin's

\*Corresponding author. Tel.: +351 239 790 000; fax: +351 239 701 266.  
E-mail address: [machado@dei.uc.pt](mailto:machado@dei.uc.pt) (P. Machado).

ideas into algorithms, allowing the evolution of populations of solutions for specific problems [9].

Due to the subjectivity inherent to aesthetic production, and to the subsequent difficulty to create an algorithm which allows the determination of the “fitness” of an artwork, evolutionary systems related to aesthetic tasks are typically guided by the user, i.e., he/she evaluates the individuals of the current population according to his/hers preferences. The works which are better classified have greater chances to generate offspring, meaning that, gradually, the populations become closer to the user’s preferences. This technique of artifact generation has been named interactive evolutionary computation.

Following this set of ideas, Karl Sims used Genetic Programming (GP) to evolve populations of images. Each genotype is a symbolic expression that once rendered becomes an image (phenotype). The user assigns fitness to the images, thus 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.

Regardless of the techniques used, evolutionary or not, most AI systems for aesthetic creation are not autonomous. They depend on human beings, particularly with regard to the aesthetic concepts, which are, directly or indirectly, imposed by the user (e.g. by selecting images accordingly to his/her preferences) or AI researcher (e.g. by establishing a set of rules that convey his/her aesthetic beliefs or a given aesthetic theory).

This paper focuses on the development of Artificial Artists. We begin by presenting a taxonomy that establishes the main conceptual differences between different kinds of artistic oriented computational systems. Next, we point out the main characteristics that an Artificial Artist should possess and present an architecture for their development. The proposed architecture is modular, which allows the independent testing of its components. As such, we briefly describe the experimental results attained independently by each module. Finally, we present the results attained through their integration.

## 2. System classification

Based on the analysis of the state of the art of the field, we have established four main categories for the classification of artistic oriented computational systems.

### 2.1. Creativity supporting tools

Systems which give support to the creativity of the user. The user is still the author of the artwork and responsible for the generation of the idea. In essence, these systems are similar to a piano or a pencil, the difference is that they are more “user friendly”, in the sense that they are designed to help the user concentrate on the creative aspects of the

task. Currently, most of the commercial artistic oriented software tends to belong to this category (e.g. Adobe’s Photoshop, Steinberg’s Cubase).

### 2.2. Computer-aided creativity

The user is no longer the only one responsible for the generation of the idea. Instead, the idea arises from an interaction between user and computer, and, as such, there is a shared authorship. A prototypical example of this type of system is the work of Sims [10].

### 2.3. Creative task analysis software

Systems that perform some sort of analysis of a creative task (e.g. musical or image analysis, classification, etc.). Although there are several works of this kind in the music field (e.g. the work of Cope on signature extraction [11]), very few results have been published in the visual domain. The work of Taylor et al. [12], where the authors use the fractal dimension estimates to classify and date Jackson Pollock’s paintings, is one few of the exceptions. Another exception is the work of Datta and co-workers [13] which uses a set of visual features and a SVM-based classifier to predict the rating of photographs, using a peer-rated online photo sharing Website (Photo.net) as data source. Some of these systems can also be seen as creativity supporting tools.

### 2.4. Artificial Artist

The computer can be seen as the creator of the art work, and it is responsible for the generation of the idea. The system should exhibit a degree of autonomy comparable to a human artist. In our opinion, systems in this category should possess the following features [14]:

#### 2.4.1. Learning

The system should not be static, it should be able to evolve through time, like human beings, who change and refine their preferences and their production in accordance with: the knowledge they acquire throughout their lives, with the surrounding cultural environment, and with its changes.

#### 2.4.2. Own aesthetics

The system must have its own aesthetic criteria, being able to evaluate its own works and those of other artists. An artificial or human artist is also, and foremost, an observer. Without the capability to perform its own aesthetic judgments, the system would be unable to appreciate the beauty of a work or to be inspired by it. Thus, it would cease to be an artist, becoming merely an image generation system.

### 2.4.3. Creativity

The capacity to produce new and innovative aesthetic artifacts is crucial for the system to be considered as a real artist. We are not interested in systems which produce minor variations of existing works, but in systems which have the potential to generate novel forms and even new aesthetic styles.

### 2.4.4. Equality

The system should not be subject to the human will; it should interact with human or artificial agents at an equal level.

### 2.4.5. Sociability

A human being does not start from scratch; he or she has access to a set of art works, to a series of experiences accumulated throughout their lives, and to the surrounding cultural environment. Similarly, an Artificial Artist should be integrated into a society and have the capacity to interact with it, thus having access to that society's aesthetic and artistic production, being able to search for inspiration in art works or events, and contributing to the society in which it is integrated.

We are not aware of any visual art system which possesses all the above-mentioned characteristics. Some works, however, do incorporate part of these properties. In what concerns learning, Baluja et al. [4] use Artificial Neural Network (ANN) in order to classify images obtained through user-guided evolution. In what concerns sociability, we may highlight several works based on co-evolutionary approaches surveyed by Greenfield [15].

Therefore, it is necessary to design an architecture which promotes the mentioned characteristics and implements the corresponding modules. Next, we describe the maturation of a system for computational artistic creation from an initial stage, when the system could be classified as computer-aided creativity application, to the later integration of aesthetic evaluation skills. An architecture called Hybrid Society is also presented. In the future, this will enable the system to expand its sociability and equality.

## 3. NEvAr: a computer-aided creativity system

Following the paradigm of interactive evolutionary computation, and inspired by Sims's work [10], we have developed a tool called NEvAr which allows the evolution of images [16] using genetic programming [17].

In the interactive mode of operation, NEvAr's functioning can be summarized as follows:

- (1) It generates a random population of images.
- (2) The user evaluates the images of the current population.
- (3) It selects a set of parents, using probabilistic roulette wheel selection (images with higher fitness have greater probabilities to be selected).




Fig. 1. A screenshot of NEvAr. The window on the left displays the current population; the one on the right displays the fittest individuals from previous populations.

- (4) It generates a new population by means of recombination and mutation of the parents' genetic material.
- (5) Back to point (2).

Like in most GP applications, the genotype of the individuals is a symbolic expression which can be represented by a tree. The trees are constructed from a lexicon of functions and terminals. The internal nodes are functions while the leaves are terminals. We use a function set composed, mainly, of simple functions such as arithmetic, trigonometric and logic operations. The terminal set is composed of the variables  $x$  and  $y$ , and of constants which can be scalar values or 3D-vectors (in order to allow the evolution of true-colour images).

We use two kinds of genetic operators: recombination and mutation. For recombination, we use the standard GP crossover operator [17], which exchanges sub-trees between individuals. Five mutation operators are used: sub-tree swap, sub-tree replacement, node insertion, node deletion and node mutation [16]. In Fig. 1, we present a screenshot of NEvAr, while in Fig. 2, we present some examples of images created through user-guided evolution and their corresponding genetic code.

This system is described in detail in Ref. [16]. In Ref. [18] an evaluation of NEvAr as a tool for artistic creation is described.<sup>1</sup>

<sup>1</sup>Examples of images with NEvAr by university students can be found at: <http://sion.tic.udc.es/jornadas/nevar/main.php?g2itemId=800>. Further samples can be found at NEvAr's website: <http://eden.dei.uc.pt/~machado/NEvAr/>.



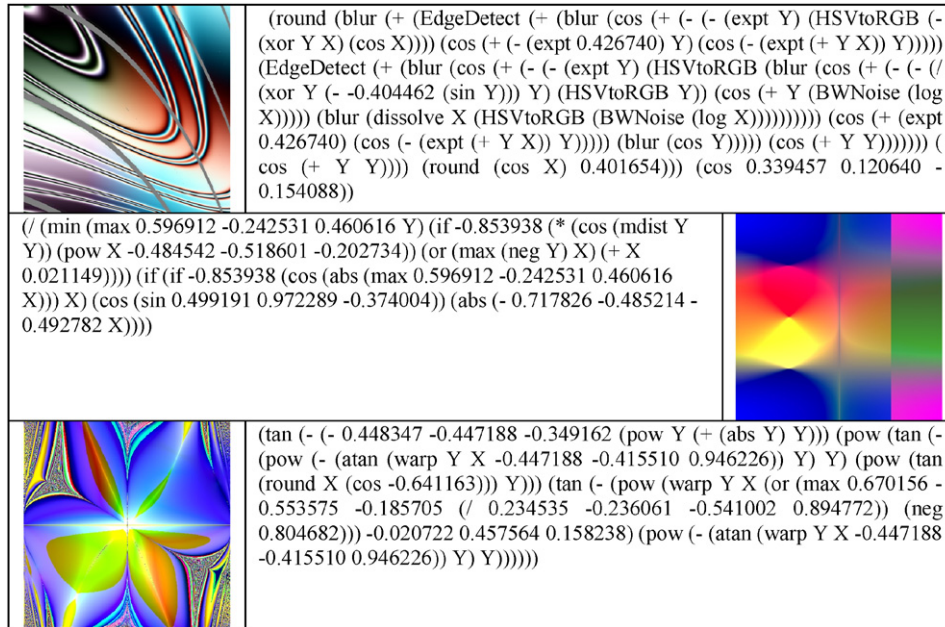


Fig. 2. Examples of images created with NEvAr by user-guided evolution and the corresponding genetic code.

In opposition to what happens with conventional art tools, using NEvAr implies that the designer/user ceases to be the only one responsible for the creation of the idea. Now, the idea is the result of an evolutionary process and of the interaction between designer and software. That is, using this system implies changes of the creative and artistic processes, since the user has a lesser degree of control. Nevertheless, the works generated comply with the user’s aesthetic preferences, which are reflected in the works created. Therefore, this is a clear example of computer-aided creativity.

#### 4. NEvAr: a first approach to autonomous computational aesthetics

According to our views, the ability to perform aesthetic judgements is a requirement for an Artificial Artist, which led us to explore and develop computational aesthetic approaches.

In a first attempt [16,19], inspired by the works of Moles [20] and Arnheim [21] and on studies of human perception that indicate a preference for simplified representations of the world as well as a tendency to perceive it in terms of regular, symmetric and constant shapes [21–24], we used complexity estimates and hand-coded formulas to evaluate aesthetics. The underlying aesthetic theory that we tried to capture with these formulas is the following: the aesthetic experience is (at least partially) related to the pleasure resulting from finding a compact percept for a visually complex stimulus.

This approach rewards images that are simultaneously visually complex and easy to perceive. For that purpose, we employ estimates for the Complexity of the Percept (CP) and for the Complexity of the Visual Stimulus (CV).

CP and CV are estimated through the division of the root mean square error (RMSE) by the compression ratio resulting, respectively, from the fractal (quadratic tree based) and JPEG encoding of the image.

Although the theory and its implementation are debatable, this approach allowed the autonomous evolution of a wide variety of images with arguable aesthetic merit (see Fig. 3). In a subsequent study [25], we explored a variation of this approach where human-made evaluations take precedence over computational ones, resulting in a semi-autonomous mode of operation which allows the user to guide evolution without requiring him/her to evaluate all the populations.

#### 5. Artificial art critics: an example of creative task analysis

All human artists resort to their ability to act as observers. They must be able to judge their works (and eventually the works of others). This ability to respond at an aesthetic level, that all humans possess, must be provided to our artificial artists. The importance of this ability and its high degree of complexity (aesthetic evaluation is subjective, social, dynamic, etc.) has led us to focus on the concept of Artificial Art Critic (AAC).<sup>2</sup>

In Ref. [26], we put forward a methodology for the development of AACs. Following this methodology, we first concentrate on systems that classify music and images according to style or author, subsequently proceeding to

<sup>2</sup>Although we use the expression “Artificial Art Critic”, we do not wish to mimic the behaviors of professional (human) art critics. What we are interested in is the modeling of the aesthetic response of a typical human observer.




Fig. 3. Examples of images evolved autonomously in evolutionary run.

systems able to perform aesthetic evaluation tasks, which entail subjective criteria.

Given the high subjectivity of the application fields we deal with, the AACs must possess adaptive skills. The architecture used integrates two modules: (i) feature extractor and (ii) adaptive classifier system. The metrics used in the visual feature extractor can be classified into three groups.

#### 5.1. Based on image compression

Similar to the ones used on previous work [16,19], these metrics estimate image complexity by considering the compression rate and error associated with the fractal and JPEG compression of the images.

#### 5.2. Based on Zipf's law

Inspired by previous work in the musical domain [27], these measurements consider the slope of the trend line of the Zipf distribution [28] of the pixels intensities and the linear correlation with the trend line.

#### 5.3. Based on the fractal dimension

Inspired by the work of Taylor et al. [12], the fractal dimension of the image and of the image's edges was calculated through the box counting method.

The selection of these metrics was inspired by several studies which associate aesthetics with complexity [21], Zipf's law [27], and fractal dimension [12]. This set of metrics is recursively applied to the entire image and to five image partitions (the quadrants and an overlapping central region). The extraction of metrics from the partitions allows their subsequent comparison, which can arguably provide insight into relevant aesthetic characteristics such as balance, symmetry and rhythm. A more thorough description of the underlying rationale for the features and of the feature extractor can be found in [29].

Once feature extraction is performed, the role of the adaptive classifier system is to perform a classification according to a given criterion (e.g. author, style, success

achieved, interest for a given user, etc.) based on the collected measurements.

In our case, the classifier is implemented by means of a feed-forward multilayer ANN trained with the standard back propagation algorithm. The ANN receives as input 198 measurements resulting from the feature extraction process.

In order to validate the AAC in tasks allowing an objective analysis of the results, its capacity to identify the author of a set of paintings was analysed. A corpus containing a total of 802 paintings created by six artists was used. In the set of experiments performed a set composed of 70% or 85% of the images was used for training, while the remaining images were used for testing purposes. The success rates ranged from 90.9% to 96.7% [30] in the test set (success rates in the training set were always of 100%). These results demonstrate that, in the considered experimental settings

- the set of considered features is sufficient to discriminate between stylistically different images, and
- the ANN is able to generalize from a relatively small set of examples.

Prompted by this success, and following the idea of using tests where performance can be objectively measured, we used a psychological test, entitled Design Judgement Test (DJT) [31], which may be applied to individuals or groups, allowing the calculation of certain aptitude components for the evaluation of aesthetic elements. With this aim, the test tries to measure the degree of recognition and reaction to aesthetic principles. The test includes 90 pages. Each comprises two or three similar designs. One of the designs follows general aesthetic guidelines, while the others violate aesthetic principles.

Due to the small number of training patterns, we decided to use a cross-evaluation technique. The 90 pages were grouped randomly into 20 sets (each containing 4–5 pages). Using a 20-fold leave-one-out cross validation approach, we independently trained 20 ANNs (for each ANN, 19 sets were used for training and the remaining one for testing purposes) attaining an average success rate of 74.49% in the test set [32].

A feature extractor for the music domain was also developed and thoroughly tested. This component was developed by researchers at the College of Charleston [27,33–35]. This system uses music metrics based on Zipf's law [28] and fractal dimension. It has been applied to the automatic classification of music pieces based on author identification [30]. It has also been applied to aesthetics-based classification involving emotional responses and data collected from humans [34,35]. Success rates are higher than 85% across all experiments.

## 6. Artificial artist—the first steps

Motivated by the results obtained with standalone AACs, we started a new research stage characterized by the development of systems where AACs are used in conjunction with the evolutionary creator described in the third section. This approach constitutes a step towards the development of an Artificial Artist.

The idea is to create an arms-race between an AAC and NEvAr. The role of the AAC is to discriminate between “paintings” and images created by NEvAr. The role of NEvAr is to find images that are classified by the AAC as paintings.

In an initial stage, the AAC is trained so as to differentiate between images generated randomly by NEvAr and paintings of well-known artists. In a subsequent stage, NEvAr is used to create populations of images that are evaluated by the AAC, which supplies the fitness values, guiding evolution. Images that are classified as paintings by the AAC attain higher fitness scores. Once the evolutionary run is over, the created images are added to the training set, the AAC is re-trained, and the process is repeated.

Succinctly, the approach used can be described as follows:

- (1) A training set composed of works by famous authors and images generated at random by the NEvAr is constructed.
- (2) The network is trained to distinguish between NEvAr-generated images and human works of art.
- (3) Once the network has been trained, it is used to guide the evolutionary process. Fitness is proportional to the degree of activation of the neuron, indicating that the image is a human work of art, and inversely proportional to the degree of activation of the neuron, indicating that the image is produced by NEvAr. That is, the goal is to develop images classified by the AAC as human works of art.
- (4) After a pre-established number of generations, the evolutionary algorithm is interrupted. The images produced during the process are added to the training set, being classified as NEvAr-produced images.
- (5) The process is repeated from step 2.

The competition promoted makes both systems (co) evolve from iteration to iteration. The AAC is forced to

learn new ways to distinguish between human and NEvAr works. NEvAr is forced to find new images that the AAC classifies as created by humans. While the iterative refinement of the training set promotes “stylistic” change, the fixed set of external images is meant to ensure that the evolved imagery relates to human aesthetics.

We conducted a series of experiments using this approach. In this paper, we present a short overview of the experimental results attained in the evolution of grayscale images.

The feature extractor used collects 108 metrics, which are fed to a feed-forward ANN with 108 neurons in the input layer, 12 neurons in the hidden layer, and two output neurons. The activation value of the first output neuron determines the degree of belonging to the class of paintings. Conversely, the activation value of the second neuron determines the degree of belonging to the class of NEvAr generated images.

The initial set comprised 1000 paintings and 1000 NEvAr images (randomly generated). We randomly selected 70% of the images of the initial set for training, 10% for testing, and 20% for validation purposes. The ANN was trained using the standard backpropagation algorithm. Training was halted when the RMSE in the test set was lower than 0.05 or when the number of training cycles reached 500.

To measure the recognition rates of the ANN we consider a winner-takes-all strategy, i.e. the output neuron with the highest activation value determines if the image was classified as a painting or as a NEvAr one. The experimental results show that the ANN was able to discriminate between the set of paintings and NEvAr generated images, attaining recognition rates of 100% in training and above 99% in the test and validation sets, which indicates that it is able to generalize properly.

Each evolutionary run comprises 100 generations, with a population size of 200. We used a fixed random seed, which implies that the initial population is the same for all iterations. The function and terminal sets used are equal to those traditionally employed when using NEvAr in interactive mode.

As previously described, the ANN is re-trained in each iteration. In spite of the fact that training took substantially longer as the number of iterations increased (due to the higher number of training patterns, and also to the increased difficulty of the task), the ANNs were always able to attain recognition rates higher than 98% in test and validation.

In what concerns the evolutionary component, NEvAr was always able to find images that the ANN used to assign fitness classifies as paintings. That is, although in the long run one of the systems—evolutionary creator or AAC—will eventually fail to cope, that stage has not been reached so far.

In Fig. 4, we present prototypical images of the different iterations of one of the first experiments carried out using this approach. Fig. 5 shows one of the images from first



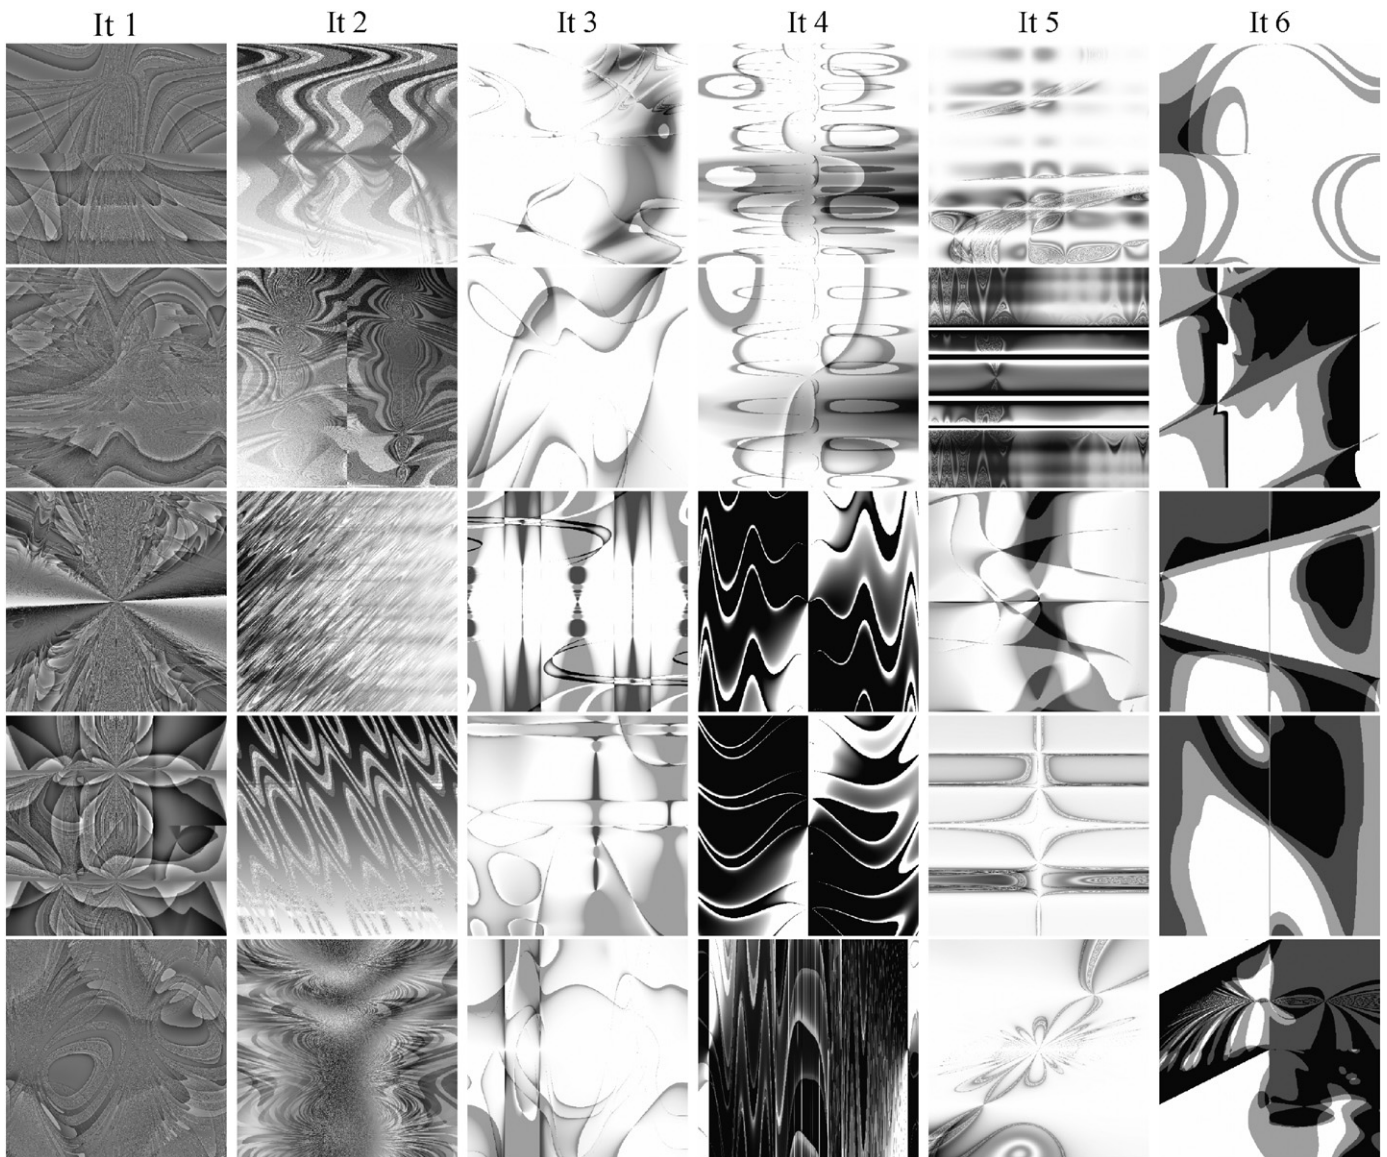


Fig. 4. Examples of images from the first six iterations of a particular experiment.

iteration. A thorough description of this iterative approach can be found in Ref. [29], where we present results attained in the evolution of true-colour images.

This iterative approach has also been applied in the music domain. In Ref. [34], various experiments are described involving a genetic-programming system called NEvMuse. NEvMuse autonomously “composed” novel variations of J.S. Bach’s Invention #13 in A minor (BWV 784). These variations have been judged as aesthetically pleasing by human listeners. A comparison of aesthetic judgments from an artificial music critic with emotional responses from 23 human subjects has revealed significant correlations.

In the current stage of development, our system possesses three of the five characteristics mentioned in the second section, namely: (i) learning abilities: since it is a system that evolves through time, refining its preferences and creations autonomously; (ii) own aesthetics: since it

establishes some aesthetic criteria which allows the evaluation of its own works, as well as those of others; (iii) creativity: the dynamic approach ensures that the system explores new, and in some cases novel, paths.

In the next section, we describe the framework in which these components will be integrated so as to comply with the remaining characteristics (sociability and equality).

## 7. Hybrid society

The final goal of our research is to build artificial artists which have an autonomous functioning within a social environment, interacting both with artificial and with human agents. For that purpose, we have designed a framework called Hybrid Society (HS) [36]. In this environment, artificial agents have access to the same tools as their human collaborators. By using them, one has access to the collective “cultural memory” of the society,




Fig. 5. Example of an image created in the first iteration.

being able to search for inspiration for its own artistic production. This search takes on board several aspects, including the machine's aesthetic preferences, the current social trends, the degree of innovation of a particular work, etc. Thus, we ensure equality and sociability in HS.

HS was designed specifically for social domains and is, therefore, based on a social conception. According to this view, only those products which are found to be interesting in a given cultural surrounding are valued. A cultural surrounding may be defined as a set of participants with a high degree of cultural affinity.

The concept of scenario is central to HS. A scenario defines the rules of the game and of communication among agents. The agents (artificial or human) can play the roles of creator or critic. Creators send products to the scenario, while critics judge products, communicating these judgements through bets on products. The “energy exchange” and “affinity” mechanisms become crucial for establishing social relations [37]. The energy of a creator depends on the bets placed on its products. The energy of a critic depends on the success of the works of art on which it bets (by placing a bet the critic ensures a percentage of future bets). Each agent and artifact have a spatial position. The spatial representation ensures that critics are more likely to receive products which are spatially close, and that products, critics and creators are displaced according to the evaluations performed. All together, this allows the emergence of sub-societies that share common interests and preferences.

We are currently developing an installation which fulfils these requirements, allowing the creation of a hybrid society. The Golem project stands out as a first approach.

Financed by Microsoft,<sup>3</sup> this project aims at fostering the creation and evaluation of web services in an environment that combines learning, competition and leisure. Although all agents are human, the mechanisms employed in Golem are similar to those of HS, and as such this project shows the feasibility of using them to create an “artistic community” of software developers.

## 8. Conclusions

Given the heterogeneity of current systems and applications related to artistic creation, we provide a classification that allows a global analysis of their characteristics. As basis for the research, and bearing in mind the characteristics which should be incorporated to an Artificial Artist, we developed a system composed by two modules. The first module is an evolutionary engine which has been thoroughly tested in several modes of operation—interactive, semi-autonomous and autonomous [16].

The second module is an adaptive classifier. The experimental results attained in its independent testing show that the proposed set of features is able to capture relevant stylistic properties, and that the ANNs are able to use these features to discriminate successfully between authors and styles. Additionally, the results attained by the adaptive classifier in the DJT show that it is also able to capture some principles of aesthetic order.

The integration of the evolutionary engine and the adaptive classifier allowed us to build an autonomous system which partially fulfils the requirements for an Artificial Artist. The experimental results attained indicate that the system presented has learning abilities, is being able to build its own aesthetic criteria (based on a set of positive examples), and is being able to develop its aesthetic criteria and “artistic” production through time.

To better fulfil the requirements of an Artificial Artist, we propose the integration of the system proposed in a hybrid egalitarian society.

It is relevant to highlight the contributions which can be made by the systems and subsystems discussed in this paper in the context of Digital Art. Computer-aided creativity systems can be used to democratize creation, allowing the exploration of the creative skills of people who might lack artistic knowledge.

Automatic classification systems, such as the ones presented, can be used for the classification of multimedia contents in organizations and search tools. If these systems have the ability to perform aesthetic judgements (e.g. by modelling user preferences) they may also be used for personal searches of musical or visual contents, or even as part of an online sales platform.

Finally, frameworks such as Hybrid Society, even when used only by human participants, may constitute a tool for the dissemination of creations, and for building sub-societies of users who share affinity relations.

<sup>3</sup><http://www.golemproject.com>.



## Acknowledgements

The authors would like to acknowledge Bill Manaris for the valuable ideas and suggestions provided in the course of our collaborative work. In particular, the idea of using Zipf and fractal dimension metrics in the visual domain, and the help provided in their implementation, were particularly valuable. Santiago Gonzalez integrated the AAC and NEvAr, and helped on the collection of data and on the setting of the experimental environment. Antonio Seoane provided the scripts used in the analysis of the experimental results. Patrick Roos wrote the code upon which the calculation of the FD is based. Dwight Krehbiel, Dallas Vaughn, Luca Pellicoro and Marisa Santos made relevant contributions to previous projects upon which the present work is based, and to related ones. This work was partially supported by the research project XUGA-PGIDIT04TIC105012PR.

## References

- [1] Moore DL. *Ada: countess of lovelace*. John Murray; 1977.
- [2] Papadopoulos G, Wiggins G. AI methods for algorithmic composition: a survey, a critical view and future prospects. In: Wiggins G, editor. *AISB symposium on musical creativity*, Edinburgh, UK, 1999.
- [3] Harold C. The further exploits of aaron, painter. *SEHR, Constructions of the Mind* 1995;4(2).
- [4] Baluja S, Pomerleau D, Todd J. Towards automated artificial evolution for computer-generated images. *Connection Science* 1994; 6(2):325–54.
- [5] Stiny G, Gips J. Shape grammars and the generative specification of paintings and sculpture. In: Freiman CV, editor, *Information processing*, vol. 71. Amsterdam: North Holland Publishing Co; 1971. p. 1460–5.
- [6] Margaret AB. *The creative mind: myths and mechanisms*. New York: Basic Books; 1990.
- [7] Csikszentmihalyi M. Implications of a systems perspective for the study of creativity. In: *Handbook of creativity*. Press Syndicate—Cambridge University, 1999, p. 313–35.
- [8] Lewis M. *Evolutionary visual art and design*. In: *The art of artificial evolution*. Springer, 2007.
- [9] Holland JH. *Adaptation in natural and artificial systems*. Ann Arbor, MI: University of Michigan Press; 1975.
- [10] Sims K. Artificial evolution for computer graphics. *ACM Computer Graphics* 1991;25:319–28.
- [11] Cope D. The composer's underscoring environment: Cue. *Computer Music Journal* 1997;21(3):20–37.
- [12] Taylor RP, Micolich AP, Jonas D. Fractal analysis of Pollock's drip paintings. *Nature* 1999;399:422.
- [13] Li J, Datta R, Joshi D, Wang JZ. Studying aesthetics in photographic images using a computational approach. *Lecture Notes in Computer Science* 2006;3953:288–301.
- [14] Castro L, Romero J, Machado P, Pazos A. *Mquinas estticas*. In: *Proceedings of ARTTECH 2006*. Pontevedra: Terceira conferencia Internacional de Arte Digital y Electrónica; 2006.
- [15] Greenfield G. Co-evolutionary methods in evolutionary art. In: *The art of artificial evolution*. Springer, 2007.
- [16] Machado P, Cardoso A. All the truth about NEvAr. *Applied Intelligence, Special Issue on Creative Systems* 2002;16(2):101–19.
- [17] Koza JR. *Genetic Programming: on the programming of computers by natural selection*. Cambridge, MA: MIT Press; 1992.
- [18] Machado P, Cardoso A. NEvAr—the assessment of an evolutionary art tool. In: Wiggins G, editor. *Proceedings of AISB'00 symposium on creative and cultural aspects and applications of AI and cognitive science*, Birmingham, UK, 2000.
- [19] Machado P, Cardoso A. Computing aesthetics. In: Oliveira F, editor. *XIVth Brazilian Symposium on Artificial Intelligence SBIA'98*. LNAI Series. Porto Alegre, Brazil: Springer; 1998. p. 219–29.
- [20] Moles A. *Thorie de l'information et perception esthetique*. Denoel, 1958.
- [21] Arnheim R. *Towards a psychology of art/entropy and art—an essay on disorder and order*. The Regents of the University of California; 1966.
- [22] Field DJ, Hayes A, Hess RF. The roles of polarity and symmetry in the perceptual grouping of contour fragments. *Spatial Vision* 2000; 13(1):51–66.
- [23] Tyler CW, editor. *Human symmetry perception and its computational analysis*. Lawrence Erlbaum Associates, 2002.
- [24] Wertheimer M. Laws of organization in perceptual forms. In: Ellis WD, editor. *A source book of Gestalt psychology*. New York: Harcourt Brace; 1939. p. 71–88.
- [25] Machado P, Romero J, Cardoso A, Santos A. Partially interactive evolutionary artists. *New generation computing—special issue on interactive evolutionary computation* 2005;23(42):143–55.
- [26] Romero J, Machado P, Santos A, Cardoso A. On the development of critics in evolutionary computation artists. In: *Applications of evolutionary computing*, LNCS. Essex, UK: Springer; 2003.
- [27] Manaris B, Purewal T, McCormick C. Progress towards recognizing and classifying beautiful music with computers—midi-encoded music and the zipfmandelbrot law. In: *Proceedings of the IEEE southeastcon 2002 conference*, Columbia, 2002.
- [28] Zipf GK. *Human behaviour and the Principle of Least Effort: an introduction to human ecology*. Addison-Wesley, 1949.
- [29] Machado P, Romero J, Manaris B. Experiments in computational aesthetics. In: *The art of artificial evolution*, Springer, 2007.
- [30] Machado P, Romero J, Santos A, Cardoso A, Manaris B. Adaptive critics for evolutionary artists. In: *Applications of evolutionary computing*, vol. 3005 of LNCS. Coimbra, Portugal: Springer; 2004. p. 435–44.
- [31] Graves M. *Design Judgement Test*. New York: The Psychological Corporation; 1948.
- [32] Machado P. *Inteligencia Artificial e Arte*. Ph.D. thesis, University of Coimbra, Coimbra, Portugal (in Portuguese), 2007.
- [33] Manaris B, Vaughan D, Wagner C, Romero J, Davis RB. Evolutionary music and the Zipf-Mandelbrot Law: progress towards developing fitness functions for pleasant music. *EvoMU-SART2003—1st European Workshop on Evolutionary Music and Art*, Essex, UK, *Lecture Notes in Computer Science, Applications of Evolutionary Computing*, LNCS 2611, 2003, Springer, p. 522–34.
- [34] Manaris B, Roos P, Machado P, Krehbiel D, Pellicoro L, Romero J. A corpus-based hybrid approach to music analysis and composition. In: *Proceedings of 22nd conference on Artificial Intelligence (AAAI-07)*, Vancouver, 2007.
- [35] Manaris B, Romero J, Machado P, Krehbiel D, Hirzel T, Pharr W, et al. Zipf's law, music classification and aesthetics. *Computer Music Journal* 2005;29(1):55–69.
- [36] Pazos A, Santos A, Arcay B, Dorado J, Romero J, Rodriguez J. An application framework for building evolutionary computer systems in music. *Leonardo* 2003;36(1):61–4.
- [37] Romero J. *Metodologia Evolutiva para la construccin de modelos cognitivos complejos*. Exploracin de la creatividad artificial en composicin musical. Ph.D. thesis, University of Corunha, Corunha, Spain (in Spanish), 2002.