On the Development of Critics in Evolutionary Computation Artists

J. Romero[1], P. Machado[2], A. Santos[1], A. Cardoso[3]

[1] Creative Computer Line – RNASA Lab – Fac. of Computer Science – University of Coruña, Spain. [2]Instituto Superior de Engenharia de Coimbra; Quinta da Nora, 3030 Coimbra, Portugal [3] CISUC – Centre for Informatics and Systems; University of Coimbra, Pinhal de Marrocos, 3030 Coimbra, Portugal

jj@udc.es, machado@dei.uc.pt, nino@udc.es, amilcar@dei.uc.pt

Abstract. One of the problems in the use of evolutionary computer systems in artistic tasks is the lack of artificial models of human critics. In this paper, based on the state of the art and on our previous related work, we propose a general architecture for an artificial art critic, and a strategy for the validation of this type of system. The architecture includes two modules: the *analyser*, which does a pre-processing of the artwork, extracting several measurements and characteristics; and the *evaluator*, which, based on the output of the *analyser*, classifies the artwork according to a certain criteria. The validation procedure consists of several stages, ranging from author and style discrimination to the integration of critic in a dynamic environment together with humans.

1 Introduction

The creation of art with computers is an old dream and there were several systems that create, assess or process art by using all types of computational methods. Over the last few years, many evolutionary computation (EC) systems have been dedicated to the creation of art [1]. Indeed, natural evolution has created some truly beautiful forms, and evolutionary computation techniques have proved effective in fields that require a certain degree of creativity [2].

An EC process consists of two stages, the generation of new individuals and their evaluation. Considering a multi-agent architecture, one can view the evolutionary process as an interplay between two types of agents, creators and critics. The creators are responsible for the generation of the new works; the critics evaluate the generated works, thus determining their survival probabilities. The evaluation plays a key role in EC, since it guides the search procedure. However, in fields that encompass a vast amount of subjective and cultural criteria, the development of an evaluation method presents considerable problems. As such, the main difficulty in the application of EC techniques to the field of the arts lies in the development of an appropriate fitness function.

"We are moving towards creation of machine art, or direct collaboration with machines, in which software makes aesthetic judgments." [3]. The present paper is about the development of artificial art critics (AAC), i.e. agents that evaluate artworks based on aesthetic and cultural criteria. In the real world the "value" of an artwork can only be assessed by taking into account its cultural context. One of the shortcomings of nowadays Computational Artists is precisely their isolation. In order to provide a cultural context to our agents, we propose their integration in a Hybrid Society of artists and critics, both computational and human.

Next we describe the structure of this paper. Section 2 shortly reviews the state of the art in the development of AAC's, classifying current systems into different classes, and making an analysis of the characteristics of these classes. Section 3 presents a brief description of some of our previous research efforts, directly connected with the generation and evaluation of artworks with EC techniques.

Based on the analysis presented in section 2 and on the conclusions derived from our previous work, in section 4, we propose a general architecture that tries to maximize the virtues of the different approaches and minimize their shortcomings. The proposed architecture includes two modules: the *analyser*, which does a pre-processing of the artwork, extracting several measurements and characteristics; and the *evaluator*, which, based on the output of the analyser, classifies the artwork according to a certain criteria. Thus, unlike most of the evaluation systems, the assessment does not deal with the artwork directly; instead, it is based on some of the artwork's characteristics. In our proposal, the *analyser* is static, i.e. the properties taken into account do not change over time. The *evaluator* module is adaptive, which means that the way the characteristics are evaluated, changes through time.

One of the main difficulties on the development of Computational Artist, and more specifically AAC's is their validation. We propose a multi-stage validation methodology. The first steps of this procedure, allow the objective, yet meaningful, assessment of the developed AAC's, providing a solid basis for their development. The later steps consider more dynamic criteria, and include testing the AAC's in a hybrid society of humans and artificial agents.

Finally, in section 5, we draw a series of conclusions and outline future research.

2 State of Art

We can distinguish two roles in any system of artistic creation: the creator (or author) and the critic (or audience). This section briefly analyses systems in which the generation role is played by an evolutionary algorithm, focusing on the classification of the approaches used to evaluate the artworks. Due to lack of space, we cannot include a wide list of references. A more elaborate survey can be found in [1]. In the field of music generation, several works analyse this aspect, such as: Burton [4], Romero [5], and Todd [6].

Taking into account the great number of evolutionary systems devoted to artistic tasks, it is convenient to classify them according to the type of approach used to implement the AAC. Following this criterion we obtain four main classes: interactive; example based; rule based; and EC based.

In the first type of system – the interactive type – the role of the critic is played by a human being, who evaluates the pieces generated by the system, and thus guides the evolutionary process. The critic's role can be played by a single person or by a group;

in the latter case the generated works are evaluated simultaneously by several people, and the system uses the average value. Some examples of works in this category are [7-15], in the musical domain, and [16-25], in the visual domain.

The main drawback of this type of system is the time involved in the user's evaluation of the artworks, which severely limits the number of generations [17, 26, 27]. This drawback is particularly severe in the musical domain. Additionally, the lack of consistency of the human evaluation also raises difficulties.

These negative consequences of human evaluation have led researches to develop systems that try to learn the user preferences from a set of examples. The most common approach relies on the use of Artificial Neural Networks (ANN's). Typically, the ANN's are trained using, either, a set of pieces of a particular style and/or author, or, a set resulting from an interactive evolutionary system. Some examples of evolutionary systems that resort to example based evaluation are [28, 29, 30], in the musical domain, and [31], in the visual domain.

This type of systems appears to have great potential, however, the results achieved are often disappointing, particularly in the visual arts domain. The main difficulties in the development of these systems are: the creation of a representative training set; and the huge amount of information present in the training artwork.

In rule-based systems, the AAC is formed by a set of rules that conduct the system. The set of rules tries to express some knowledge of the concerned domain, and typically, results from stylistic requirements of a particular type of artwork. In the musical domain, one can find several examples of this type of system, some examples being [27, 32, 33]. We weren't able to find any system using this type of approach in the field of visual arts, which can't be considered surprising, when we take into account that musical theory is way more developed than visual art theory.

The main problem of this type of approach is their lack of generality. These systems are based around a particular vision of an artistic style or theory. As such, their adaptation to other styles is difficult, if not impossible. Their main advantage lies in the possibility of using a set of formalizations, structures, metrics, and knowledge, which makes the analysis of the pieces of art easier.

A further possibility is to use an evolutionary system as AAC. A typical approach would be to co-evolve two populations: one of creators and one of AAC's. In this type of system the fitness of a given individual (creator or critic) is determined by the interplay of between agents, giving rise to new and isolated aesthetics. Some examples of this type of approach are: [34] that composes music; [35, 36] that uses artificial life techniques to create musical themes and sounds; and [37] that resorts to co-evolution to evolve images.

3 Background Work

This section shortly describes some of our previous work, which is the basis for our ongoing research. This description comprises examples of some of the classes of systems introduced in the previous section. We will focus on the analysis of the critics used.

3.1 Music Generation and Hybrid Society

The Tribe project started in 1998 [12, 13] and aims at building artificial models of human music composers. The underlying idea of this project is to follow the evolution of human music through time, i.e., start by focusing on the most primitive and simple forms of music, and gradually build more complex models.

Tribe is a typical interactive evolutionary system that composes this very primeval music, namely purely percussive. Each *tribe* consist on one percussive pattern. The user assigns the fitness of each tribe, in an interactive way.

The experience with the Tribe project revealed the need to facilitate the design of social artificial beings, able to develop and assess new creative products. Following this goal, a model of an egalitarian society populated by humans and artificial beings was proposed, giving rise to the Hybrid Society Project[1]. Since the performance of the agents acting in Hybrid Society, is evaluated by a dynamic society of both artificial and human agents, this paradigm allows an adequate validation of a social system.

Moreover, this paradigm provides a natural learning approach, intermediary between learning by discovering and learning by reinforcement, as it mimics the dynamics of a human society.

Although this approach presents some logistic difficulties since it needs the participation of several humans. However, it is more flexible than standard interactive evolutionary systems, since it is designed to allow the change of the human participants through time. In standard interactive evolutionary systems, the interaction of one human, or a close group, it is required during the all duration of the experiment. Moreover, Hybrid Society takes into account the “opinion” of each of the humans (and artificial systems) in a group and not just an average opinion.

Hybrid Society allows the incorporation of computer systems that manage several artificial agents, which can have different “genetic code” and distinct behaviour at the same time. In this sense, each computer system can be seen as a species with different individuals. This capacity provides the necessary conditions to allow the use of hybrid society in conjunction with evolutionary computation systems.

In the first experiments this paradigm was applied in the musical domain [38]. Two artificial species were considered: one of artificial creators, based on the previously mentioned Tribe project; and one of AAC called “Oreja” (Spanish word for ear). Additionally, the Hybrid society also included human critics and creators.

The AAC resorted to a set of ANN’s, which were trained using evolutionary computation techniques. The weights of the connections between the nodes, and the architecture of the net were included in the genetic code, and, could thus evolve through time. Each rhythmic pattern is codified as an array of 160 binary elements, 10 percussive instruments in 16 slots of time. The short dimension of the rhythmic patterns allows us to use them, directly, as input for the ANN’s. In the future we intend to apply Hybrid Society to the evolution of more complex structures, the size of these structures will, eventually, make it unfeasible to use them directly as input. As stated before, our idea is to extract relevant characteristics, properties and

[1] More information about the Hybrid Society Project can be found in http://www.hybridsociety.net

measures, of the artworks and use them as input for the ANN's. This idea will be discussed in more detail in section 4.

3.2 Visual art

NEvAr (Neuro Evolutionary Art) is a research project that aims at building an artificial artist in the field of visual arts. In the simplest mode of operation NEvAr is an evolutionary art tool, it allows the evolution of populations of images, which are evaluated by a user that guides evolution.

NEvAr was inspired in the works of Dawkins [39] and Sims [21] and shares several similarities with the latter. It resorts to Genetic Programming [40], and as such the genotypes are trees constructed from a lexicon of functions and terminals. The function set is composed mainly of simple functions such as arithmetic, trigonometric and logic operations. The terminal set is composed of a set of variables x and y and random constants. The phenotype (image) is generated by evaluating the genotype for each (x,y) pair belonging to the image. Thus, the images generated by NEvAr can be seen as graphical portrayals of mathematical expressions.

As usual, the genetic operations (recombination and mutation) are performed at the genotype level. In order to produce colour images, NEvAr resorts to a special kind of terminal that returns a different value depending on the colour channel – Red, Green or Blue – that is being processed.

In [17] we made an assessment of NEvAr as a tool. According to our analysis the artworks evolved with NEvAr reflect the aesthetic and artistic principles of the user[2]. This analysis also revealed the importance of the individual's database. In NEvAr the user has the possibility to store highly fit individuals in a database. Later these individuals can be injected in an ongoing experiment, or used to create a non-random initial population. By resorting to this database, one can significantly decrease the amount of time necessary to create "good" images. Since the recombination possibilities are virtually infinite, the generated images will still be new, and, in most cases, innovative. However, they will share several characteristics with the selected database individuals, which can be considered an inspiring set [17].

With time, the database size has increased drastically, which has led to the development of automatic seeding procedures. The basic idea was to select a set of database images that resemble one supplied by the user, and then make this set as part of the initial population. To achieve this goal we needed to develop a way to compare images. After testing several measures of distance among images, we came to the conclusion that a direct comparison was not appropriate [17]. Instead, we extracted some characteristics of the images that were deemed important (in this case several complexity estimates), and then performed the comparison based on these characteristics [17]. Although limited, this approach gave promising results, especially if we take into account that we used a very small number of characteristics.

As stated in the beginning of this section the ultimate goal of this project is to create a computational artist. This means that NEvAr should be able to work

---

[2] For examples of artworks generated with NEvAr can be found at: http://www.dei.uc.pt/~machado.

autonomously, and hence, assign fitness to the individuals based on aesthetic criteria. In other words, NEvAr must be able to act, at least to some extent, as its own critic.

The approach used in NEvAr to evaluate images, is based on our personal views about aesthetics, and relies on the idea that the aesthetic value of an image is connected with the sensorial and intellectual pleasure resulting from its perception. Moreover, this pleasure is deeply related with the perceived complexity of the sensorial stimulus, and with the complexity of the percept (the representation of what is perceived).

Based on this notion, and using estimates of image and processing complexity, we developed an evaluation procedure that, basically, attributes high fitness values to images that are, simultaneously, visually complex and easy to process. A full description of this procedure, along with some experimental results can be found in [41], for a discussion of the importance of complexity in art see, e.g., [42] [43].

The evaluation procedure only takes into account the lightness information of the images, discarding the hue and saturation information. Therefore, in this mode of execution, we are limited to greyscale images.

An analysis of the role of colour and the way colour is assigned, particularly in abstract art, leads to the conclusion that artists (certainly not all, but at least a significant proportion) usually work with a limited colour palette, and that the spatial relation between colours usually follows a set of rules. This is consistent with the view that each artist constructs its own artistic language, which complies with an implicit grammar.

The idea of creating a program to give colour to the greyscale images created by NEvAr emerged naturally. Unfortunately the development of such a program poses several problems. We are therefore developing a system that learns to colour images from a set of training ones. This approach has, potentially, several advantages over a built-in colouring procedure, namely: we do not need to code by hand a set of colouring rules; the results of the system are less predictable; and we can use paintings made by well-known artists as training set, thus learning to colour images according to their style.

In [44] we present our current approach in which, we employ Genetic Programming to evolve computer programs that mimic the colourings of the training instances. As before, the experimental results indicate that making a direct comparison between the colourings yields, poor and uninteresting results. By taking into consideration some of their underlying characteristics, we were able to improve these results significantly, obtaining interesting colouring programs [44].

4 Model proposal

In this section we propose a model for the development of an AAC. The design of this model was based on a set of characteristics that we consider desirable:

- Generality – We want a model that allows the development of AAC's for different domains; the domain specific tasks should be carried out by specialized modules, allowing an "easy" adaptation of the AAC to new domains.

- Independence – We are interested in AAC's that are able to perform autonomously in an egalitarian society of humans and artificial beings. Therefore, the agents should be able to perceive the artworks in a standard and uniform representation (binary files as bitmaps or midi) and evaluate them. In other words, they should be able to "see" (or ear, touch, etc., depending on the type of piece) the artworks and form a judgement base on what they "see". An AAC can build its own internal representation of an artwork, but it cannot access the original artwork representation (assuming that this representation exists).
- Adaptability – The AAC's should evolve and adapt with time. This however is not a strict requirement, a static AAC can also be successful; we are primarily interested in this type of AAC, because they mimic better the behaviour of human critics.
- Sociability – Ideally, the AAC's should be able to adjust their behaviour according to the demands of the society in which they are integrated.

4.1 Architecture

Figure 1 presents a rough outline of the proposed architecture. The AAC is composed by two main modules: an *analyser* and an *evaluator.*

The *analyser* is static and receives as input a direct representation of the artwork (e.g., bitmap file, midi sequence, etc.) producing some sort of analysis. The *evaluator* is a purely adaptive system, based, for instance, on ANN's or Evolutionary Computation techniques, and forming a judgement based on the analysis created by the first module.

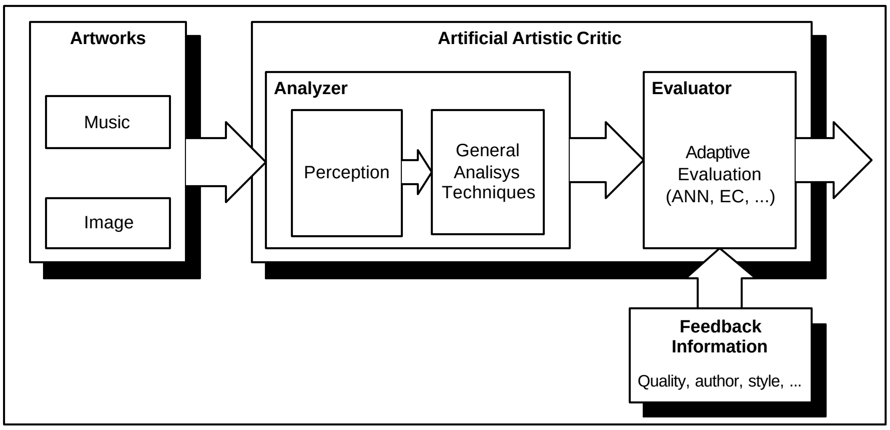

Figure 1. Outline of the proposed model.

The analysis comprises two stages, called *perception* and *general analysis*. The first step is domain specific, while the second is mostly domain independent. The *perception* stage builds a percept of the artwork, i.e. some sort of internal representation, outputting the percept and information about the perception task (e.g. the complexity of the task). The *general analysis* stage uses generic analysis techniques to extract relevant information from its input.

In some cases this division between stages is more conceptual than real. The main idea is that, the perception acquires information about domain specific parameters, which are then analysed. We do not impose any kind of constraints to the type of internal representation, nor to the range of techniques used on the *general analyser* module. Therefore, the range of techniques that can be used includes statistical, rule based, algorithmic, symbolic, or sub symbolic techniques, etc.

The second module, *evaluator*, takes as input the characterization done by the previous analyser, and outputs an assessment of the artwork. In order to allow adaptation it receives feedback information, which reflects the quality of its appraisal. The role of this information will become clearer in the next sub-section dedicated to the validation methodology.

The proposed architecture allows a certain degree of independence between the search for new relevant features, analysis, and evaluation. Moreover, the adaptive *evaluator* module can give information about the features that are relevant in the assessment of an artwork.

4.2 Validation

The validation of an AAC involves great difficulties, due mainly to the subjective nature of the task. In this section we propose a multi level validation methodology. In each level the AAC is presented with a different task. The idea is to begin with tasks in which the response of the AAC can be evaluated easily, and then move to task that involve a higher subjectivity and dynamics. In each level it is possible to compare distinct methods and approaches in order to obtain better AAC's, without losing the adaptation capacity.

Currently, we consider four levels of validation:

1. Author Identification – In this task, the AAC is presented with several artworks of different authors. Its task is to determine the author of the each piece. The *evaluator* module is trained using the feedback information, which provides the correct answer. Afterwards the system can be tested with a different set of artworks. The main goal of this validation step is to assess the quality of the analyser module, since a failure would imply that the extracted features are not sufficient to discriminate between authors.
2. Style identification – This level is very similar to the previous one, basically the AAC should be able to recognize the style of a given artwork. The training and testing is done in the same manner, the only difference being that the feedback indicates the style instead of the author. Depending on the domain and authors/styles considered, this level can be more or less difficult than level 1. The main advantage of these levels is the objectivity of the tasks being performed by the AAC.
3. Static Evaluation – The task of the AAC is to assign an aesthetic value to a series of artworks, which were previously evaluated by humans. The main difficulty in performing this test is the construction of a database of consistently

evaluated artworks. One option is to use some sort of evolutionary art tool to generate the artworks. Alternatively, we can resort to pieces of classical art. However, we need works of varying quality (which includes "bad" pieces). Depending on the set of positive and negatives examples used, the task of the ACC can be either difficult or easy.

4. Dynamic Evaluation – In this fourth level the AAC is part of a society of agents, which can be artificial or human (see the description of the Hybrid Society in section 3). The success of the AAC depends on the appraisal of his judgements by the other members of the society. This type of test introduces a new, social and dynamic dimension to the validation, since the value of an artwork varies over time and depends on the agents that compose the society. Although challenging, we think that this type of validation is the most natural one, since it tries to mimic the conditions in which human critics and authors perform. Furthermore, it is the only test that takes into consideration that the value of an artwork depends on its cultural context, and that consequently, the critic must be sensible to this context and adapt to its cultural surroundings.

The validation methodology presented here tries to find a compromise between automated and human-like validation. We are fully aware of the difficulty of the proposed tasks. However, it is our belief that AAC's, which are capable to overcome only some of the levels, can still be interesting and useful.

In the first levels (1-3) of validation it is possible to assess the performance of the *analyser* and *evaluator* module independently, since the output of the *analysis* module (in conjunction with the feedback information), can be seen as a training instance to the evaluator. In the latest level, this is no longer possible since the feedback information does not reflect directly the quality of the artworks, but only an appraisal of the AAC actions by the society, and it changes dynamically in time. So, the system must work in real time.

5 Conclusions and further work

This paper has analysed the current state of the art in the development of art critics, focusing on the ones employed in evolutionary art systems. Based on this analysis, and on the experience acquired in the development of previous systems, we have proposed a generic model for the development of artificial art critics. In order to allow an easy adaptation to different domains, the proposed architecture separates generic from domain specific components. Furthermore, it also establishes a boundary between static and adaptive modules.

The validation of an artificial art critic is a complex task. We proposed a multi level validation methodology that is aimed at simplifying and automating this task. Since the proposed architecture is not domain dependent, we would also like to test the possibility of mapping aesthetical principles between different domains, e.g. visual art and music.

The research in the area of artificial art critics and artists is still on an embryonic stage. It is our hope that the ideas presented here may help its development, by supplying a common framework that allows the comparison of different techniques, and facilitates the collaboration between researchers.

Acknowledgements

This work was partially supported by grants from Xunta de Galicia, project "Hybrid Society" from University of A Coruña, and also by the Portuguese Ministry of Education, program PRODEP III, Action 5.3.

Bibliography

1. Johnson, C., Romero, J.: "Genetic Algorithms in Visual Art and Music". In Leonardo. MIT Press. Cambridge MA Vol 35, (2). pp. 175-184. 2002
2. Bentley, P. J.: "Is Evolution Creative?". In P. J. Bentley and D. Corne (Eds.), Proceedings of the AISB'99 Symposium on Creative Evolutionary Systems (CES), pp. 28-34. Published by The Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB). Edinburgh.1999.
3. Karnow, C. E.A.: "In collaboration with machines". Leonardo, (30):4 pp 248. 1997
4. Burton, A. R., Vladimirova, T.: "Generation of musical sequences with genetic techniques". Computer Music Journal, 23(4):59--73, Winter 1999.
5. Romero, J., Santos, A., Dorado, J., Arcay, B., Rodriguez, J.: "Evolutionary Computation System for Musical Composition". In Mathematics and Computers in Modern Science, pp. 97--102. World Scientific and Engineering Society Press. 2000
6. Todd, P.M.: "Evolving musical diversity". In Andrew Patrizio, Geraint A. Wiggins, and Helen Pain, Eds. Proceedings of the AISB'99 Symposium on Creative Evolutionary Systems. Society for Artificial Intelligence and the Simulation of Behaviour, pp. 40-48, 1999.
7. Biles, J. A.: "GenJam in Perspective: A Tentative Taxonomy for Genetic Algorithm Music and Art Systems". In Wu, A.S. (Ed.), Workshop GAVAM 2000 Genetic and Evolutionary Computation Conference, pp. 133-136. Morgan Kaufmann. San Francisco CA. 2000.
8. Horowitz, D.: "Generating rhythms with genetic algorithms". In Proceedings of the 1994 International Computer Music Conference, 1994.
9. Jacob, B.L.: "Composing with genetic algorithms". In Proceedings of the 1995 International Computer Music Conference, pages 452—455, 1995.
10. Jacob, B.L.: "Algorithmic composition as a model of creativity". Organized Sound, 1(3):157--165, 1996.
11. Moroni, A., Zuben, Von F., Manzoli, J.: "Arbitration: Human machine interaction in musical domains". In Leonardo. MIT Press. Cambridge MA, Vol. 35, (2). pp. 185-189. 2002.
12. Pazos, A., Santos, A., Dorado, J. and Romero, J.J.: "Genetic music compositor". In P. Angeline, Z. Michalewicz, M. Schoenhauer, X. Yao, and A. Zalzala, Eds, Proceedings of the 1999 Congress on Evolutionary Computation, pp. 885--890, Vol. 2. IEEE Press. 1999.
13. Pazos, A., Santos, A., Dorado, J. and Romero, J.J.: "Adaptive aspects of rhythmic composition: Genetic music". In W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela, and R. E. Smith, Eds., Proceedings of the Genetic and Evolutionary Computation Conference, page 1794, Vol. 2. Morgan Kaufmann, 1999.
14. Putnam , J. B.: "Genetic programming of music". Technical Report, New Mexico Institute of Mining and Technology, 1994.

15. Ralley, D.: "Genetic algorithms as a tool for melodic development". In Proceedings of the 1995 International Computer Music Conference, pp. 501--502. 1995.
16. Haggerty, M.: "Evolution by aesthetics, an interview with W. Latham and S. Todd". .IEEE Computer Graphics, 11(2):5--9, March 1991.
17. Machado, P., Cardoso, A.: "NEvAr - The Assessment of an Evolutionary Art Tool". In: Wiggins, G. (Ed.). Proceedings of the AISB'00 Symposium on Creative & Cultural Aspects and Applications of AI & Cognitive Science, Birmingham, UK, 2000.
18. Rosenman, M.A.: "The generation of form using evolutionary approach". In D. Dasgupta and A. Michalewicz, Eds., Evolutionary algorithms in Engineering Applications. Springer, 1997.
19. Rowland, D. A.: "Computer Graphics Control over Human Face and Head Appearance", Genetic Optimisation of Perceptual Characteristics. PhD thesis, University of St. Andrews, 1998.
20. Rowland, D. A, Biocca, F.: "Evolutionary Cooperative Design Methodology: The Genetic Sculpture Park". In Leonardo. MIT Press. Cambridge MA Vol 35, (2). Pp. 193-196. 2002.
21. Sims, K.: "Artificial evolution for computer graphics".Computer Graphics, 25(4):319--328, 1991.
22. Sims, K.: "Interactive evolution of dynamical systems". In F.J. Varela and P. Bourgine, Eds., Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life, pages 171--178. Mit Press, 1992.
23. Soddu, C.: "Recognizability of the idea: the evolutionary process of Argenia". In A. Patrizio, G. A. Wiggins, and H. Pain, Eds., Proceedings of the AISB'99 Symposium on Creative Evolutionary Systems. Society for Artifical Intelligence and the Simulation of Behaviour, pages 18--27. 1999.
24. Todd, S. and Latham, W.: "Mutator, a subjective human interface for evolution of computer sculptures". IBM United Kingdom Scientific Centre Report 248, 1991.
25. Unemi, T.: "SBART 2.4 An IEC Tool for Creating Two-Dimensional Images, Movies and Collages". In Leonardo. MIT Press. Cambridge MA Vol 35, (2). pp. 189-192. 2002
26. Biles, J. A.: "GenJam: A genetic algorithm for generating jazz solos", In Proceedings of the 1994 International Computer Music, 1994.
27. Papadopoulos, G. and Wiggins, G.A.: "A Genetic Algorithm for the Generation of Jazz Melodies", In SteP'98. Proceedings of STeP'98: 8th Finnish Conference on Artificial Intelligence. September 7-9, 1998, Jyväskylä, Finland. 1998.
28. Biles, J. A., Anderson, P. G. and Loggi, L. W., "Neural Network Fitness Functions for an IGA", in Proceedings of the International ICSC Symposium on Intelligent Industrial Automation (ISA'96) and Soft Computing (SOCO'96), 26--28 March 1996 (Reading, U.K.: ICSC Academic Press, 1996) pp. B39--B44. 1996.
29. Burton, A. R. and Vladimirova , T., "Genetic Algorithms Utilising Neural Network Fitness Evaluation for Musical Composition", in G.D. Smith, N.C. Steele and R.F. Albrecht Eds., International Conference on Artificial Neural Networks and Genetic Algorithms Norwich, UK. pp. 220--224, Springer, 1997.
30. Spector, L. and Alpern, A.: "Criticism, Culture and the Automatic Generation of Artworks", In Proceedings Twelfth National Conference on Artificial Intelligence (AAAI-94) August 1-4, pp. 3—8. AAAI Press. 1994.
31. Baluja, S., Pomerleau, D. and Todd, J.: "Towards Automated Artificial Evolution for Computer-Generated Images", In "Connection Science 6, No. 2, pp. 325--354. 1994.
32. Wiggins, G.A., Papadopoulos, G., Phon-Amnuaisuk, S. and Tuson, A.: "Evolutionary Methods for Musical Composition", In International Journal of Computing Anticipatory Systems, 1(1), 1999.
33. McIntyre, R.A.: "Bach in a Box: The Evolution of Four-Part Baroque Harmony Using Genetic Algorithm", In Proceedings of the 1994 IEEE International Conference on Evolutionary Computation, pp. 852--857, IEEE Press, 1994.

34. Todd, P.M. and Werner, G.M.: "Frankensteinian Methods for Evolutionary Music Composition", In N. Griffith and P.M. Todd, Eds., Musical Networks: Parallel Distributed Perception and Performance. Cambridge, MA: MIT Press, 1998.
35. McAlpine, K., Miranda, E. and Hoggar, S.: "Making Music with Algorithms: A Case-Study System", Computer Music Journal, 23, No. 2, 19--30. 1999.
36. Miranda, E.: "Granular Synthesis of Sounds by Means of a Cellular Automaton", Leonardo, No. 4, 297--300. 1995.
37. Greenfield, G. R.: "On the Co-Evolution of Evolving Expressions, International Journal of Computational Intelligence and Applications", Vol. 2, No. 1, pp. 17-31. 2002
38. Romero, J: "Metodología Evolutiva para la construcción de modelos cognitivos complejos. Exploración de la 'creatividad artificial' en composición musical" (In Spanish). Ph.D. Thesis. University of A Coruña. 2002.
39. Richard Dawkins: "The Blind Watchmaker". Penguin, 1990.
40. Koza, J.: "Genetic Programming: On the Programming of Computers by Means of Natural Selection". MIT Press. 1992.
41. Machado, P. and Cardoso, A.: "All the truth about NEvAr". Applied Intelligence, Special issue on Creative Systems, 16(2):101–119. 2002
42. Arnheim, R.: "Entropy and Art", University of California Press, 1971.
43. Arnheim, R.: "Art and Visual Perception: A Psychology of the Creative Eye", University of California Press, 1954.
44. Machado, P., Dias, A., and Cardoso, A.: "Learning to Colour Greyscale Images". The Interdisciplinary Journal of Artificial Intelligence and the Simulation of Behaviour, 1(2), 2002.