Genetic Algorithms in Visual Art and Music.

Colin G. Johnson and Juan Jesus Romero Cardalda.

Colin G. Johnson (computer scientist, composer).

Computing Laboratory.

University of Kent at Canterbury.

Canterbury, Kent, CT2 7NX, England.

Email: C.G.Johnson@ukc.ac.uk

Juan Romero (computer scientist, musician).

Information and Communications Technology Department.

University of A Coruña.

15071 A Coruña.

Spain.

Email: jj@udc.es

Abstract

This paper is an introduction to the special section of *Leonardo* on Genetic Algorithms in Visual Art and Music, which arose from a workshop at the 2000 Genetic and Evolutionary Computing Conference. This introduction gives a background review of the area, takes a look at some open questions provoked by the workshop, and summarizes the papers in the section.

1 Introduction.

Leonardo has had a long tradition of publishing work at the mutual frontier of art and science. We are pleased to present a special section of the journal which demonstrates the application of ideas from science (evolutionary biology) through technology (computing) to art (both visual art and music).

The topic of this special section is the application of genetic algorithms (GAs) and related heuristics to visual art and music. Genetic algorithms, invented by Holland in the 1970's [1], are a heuristic method which abstracts the process found in biological evolution and simulates it on the computer. However instead of using this to simulate real biology, it is instead used to solve problems in many different non-biological domains.

A typical use of genetic algorithms is in optimization, where we want to search some space for the individual which scores highest on some measure. The genetic algorithm begins by generating a random set of individuals drawn from the search space; typically individuals are represented as binary strings. There then follows a process whereby the best individuals in the population are selected, those individuals exchange some information with others (crossover), and then some small changes are made in the individuals (mutation) to produce a new population. This process is repeated until it converges or until a satisfactory solution is found (figure 1).

FIGURE 1 NEAR HERE

Procedures such as this have proven to be a powerful way to search many different kinds of search space, with many real world applications. More details can be found in [2, 3, 4]. A number of other techniques have a similar flavour, and have also been applied to artistic and musical areas. These include cellular automata [5], artificial life [6] and autonomous agents [7].

A number of researchers have investigated the use of genetic algorithms in artistic domains. At the 2000 *Genetic and Evolutionary Computation Conference* in Las Vegas [8] a workshop was held to review the state of the art in this area. This workshop consisted of a number of presented papers, general discussion of the topics presented and demonstrations of pieces of art and music created using these systems—including a live performance by Al Biles and his GenJam system. The papers in this section are extended versions of papers presented at the workshop. In this introductory article we begin by providing some general background to this area and a review of prior work on these topics. One of the aims of the workshop was to think about some of the general principles which underlie this area, and to consider if there are any general principles or major open questions in this area. Some of the papers include such topics individually, and some general questions are gathered together in the penultimate section of this article. The final section of this introduction contains a short summary of each of the papers, designed to put them into context with each other.

We would like to thank the participants and authors for their contributions to the workshop and for preparing their revised papers for this special section with utmost efficiency.

2 Review.

Many people from different approaches have dreamed of a computer with human features. This definition of a Human Computer manifests itself in various ways, depending on the conception of Man or on the feature to be highlighted. The idea of intelligence as a differentiating characteristic, which gave rise to the term Artificial Intelligence, was one of the most common designations for the old dream of creating Artificial Humanity. Other researchers have put forward creativity, learning, or adaptation capacity in general.

The present special section analyzes works which try to realize, using computers, one of the most thrilling tasks which human beings are capable of: Art. Art has a series of features which make it really interesting when trying the "human" capacity of a computational system. Art also has a number of features which make it hard to deal with using traditional computational techniques.

- The first characteristic of art is its dynamism. Artistic and aesthetic trends evolve through time from a community point of view, and they coexist at a given point in time among different societies and also within the same society.

- This description points to another interesting feature, the social character of art. Art is unconceivable without a set of interrelated individuals. If we take a look at this change inside an individual, we will see how even within a given aesthetics or concrete artistic style, there are different evaluations of a work of art according to the critic and the time.

- Art is closely linked to man's instinctive and irrational side. In many-cases, we learn by Socratic learning or environmental immersion, while learning has not been totally comprehended or formalized.

The desire to use computation for building artistic systems can be traced-back to Ada Lovelace, who dreamed of the creation of a computer with musical capabilities 150 years ago. From that moment, Artificial Artistic Systems have been studied for a long time using every kind of computational techniques, including Expert Systems, Artificial Neural Networks (ANNs), statistical and stochastic methods.

There have been attempts at formalizing the act of musical creation throughout the history of music. The proposals of composition formalisms which appeared at different times in the history of western music, claimed that musical works could be created as a result of applying certain rules to some given initial material. This old idea survives in some present-day musical styles and in some musical computation projects which deal with composition [9]. These systems could be termed "algorithmic systems". On the other hand, there are systems based on the use of a simple algorithm on a numeric or symbolic series. In this case the series may have a great variety of origins and they contain non-musical material. These systems are called "mathematical or fractal systems". Generally speaking, these systems yield poor results, obtaining part of their sensations from parameters not included in the system itself.

There are other works which consider music to be a language. From this perspective, composition would consist of the creation of a message which transmits a certain content, following the rules of culture, cultural style and personal style. This conceptualization, which is more complex and complete than the previous one, allows the introduction of cultural factors (the group's) and particular ones (the individual's) in the process, adding the concept of analysis in the composition. Therefore, while the previous conceptualization was purely generative (the result springs from a set of rules which bear no relation with the environment), in this case there is an analytic component which allows the creation of some common aesthetic rules in the composing system.

However, many researches have been carried out in recent years which suggest the creation of artificial Artistic systems, from different approaches, using evolutionary techniques. Evolutionary computation is inspired by nature, taking some features of the evolution process in order to apply them to the computational field. These techniques started with Holland's work [1] in 1975. These techniques offer different solutions to a given problem, and the most highly adapted ones give rise to new generations of solutions through crossover, mutation and selection genetic operators.

Evolutionary Computation displays the following features:

- On the one hand, this technique has a high degree of adaptation capacity due to its origin; in fact few things are as adaptable as nature. This adaptation capacity is possible thanks to the use of control or feed-back structures.

- Besides, the social character is obvious, both in its parallel and emerging vision of the problem and in the importance of communication (understood as exchange) which exists in some of its techniques, such as artificial life.

- Another remarkable feature is the possibility of incorporating various forms of musical knowledge, not necessarily formal rules. Thus the abstraction of artistic information is not necessary, so information is not biased.

- They constitute dynamic models whose behavior is not totally defined by the model's creator. Art has not been understood yet. If we only had techniques limited to our understanding, art would be unapproachable right now.

Evolutionary Computational Systems have proved to be very accurate in those fields which require a certain degree of creativity [10]. Such is the case of tasks related to visual and musical art, as it will be explained in this paper.

Two roles may be distinguished in any system of artistic creation: Creator and Critic (Author and Audience). The works presented in this article have been organized according to the critic's role, while the creator's role has always been played by an evolutionary computational system. For a deep analysis of some of these implementations in the field of Music, see Burton [11] and Todd [12]; the paper [13] shows different approaches in the musical area.

This work shows a perspective of the different researches on Artificial artistic systems using evolutionary techniques. A classification is made based on the critic element of the different compositions. Four types of works are analyzed: interactive; based on examples; rule-based; and autonomous (figure 2). Finally, the integration of the various works in a common framework is proposed, where different approaches can compete and/or collaborate to create global compositions which can be adapted to different types of music, thus including the advantages of the different techniques.

FIGURE 2 NEAR HERE

2.1 Interactive Systems

The first category to be dealt with in this classification is that of Interactive Systems. In this kind of system, the critic is a human being, making an aesthetic evaluation of each piece in the system and thus conducting its evolution. The system takes these evaluations into account for the creation of the next set of compositions. The user's conducting role can be played by a single person or by a group; in the latter case, a group of people assesses the cybernetic composer's works simultaneously.

These systems, in their simplest form, pose the problem of time cost [13] (or *bottleneck* [14]) due to human participation. This problem may also tire the user, who has to evaluate a great number of musical examples. Besides, many researches think that these systems also have a high degree of subjectivity. On the other hand, the direct incorporation of the user allows to compose works with the right aesthetic conception for the individual or group with whom the system interacts.

In the musical field, there are works which make variations of a melody [14], new melodies [15, 16, 17, 18], jazz jams [19], rhythmical textures [20, 21, 22, 23] scores based on material provided by the user [24] and new sounds [25, 26].

In the visual domain, we should highlight a number of pieces of work. Firstly the work of Sims with his Creatures [27, 28], and the evolution of complex simulated structures, textures, and motions [29, 30, 31]. Secondly that of Latham, that generates live 3D forms in his "Organic Art" [32, 33, 34, 35]. There are also works related to architecture [36], 2D pictures [37, 38] and design and exploration of visual spaces [39, 40, 41, 42, 43, 44].

2.2 Systems based on examples.

The possibility of registering the user's tastes within a subsystem was suggested in some instances, with a double purpose. The first would be that of facilitating the system's learning rate, and using the present artistic works to conduct it. The second would be solving the problems of interactive systems related to slowness and specialization. This subsystem is usually integrated by an ANN trained from artistic pieces. ANNs are a type of computer system which try to simulate the behavior of the natural neuron and can learn from samples. These pieces are examples of some artistic style or author, or else they stem from some interactive system.

In the musical world, we may find examples in the rhythmic domain [45, 46], with music inspired by Charlie Parker's songs [47, 48], in Jazz music [49] and in four-part Western harmony compositions [50]. In the visual domain there are works in 2D pictures [51, 52].

2.3 Rule-based Systems

In rule-based systems, the critic is built from a set of rules which conduct the system. This set of rules is built by the system's author from his/her musical knowledge or from artistic studies.

In the musical domain, most of the works are related to harmonization [53, 54, 55, 56]. There are also works which make jazz solos [53, 56], a work which makes minimalist music [57] and another one with two critics, a set of rules and the user [58].

2.4 Autonomous Systems

The radical change in the separation between system and user occurs in those systems which have their own autonomous aesthetics. In this case, artistic works evolve following their own path, which may have nothing to do with human aesthetics. They are usually regarded as models of social evolution.

In the domain of music there are works where a group of elements work as evaluators and others work as composition creators, while both evolve simultaneously [58, 59]; and works which compose melodies [60] and make sounds [61, 62] with artificial life techniques.

2.5 Conclusions

The state of the Art shown in this paper reflects the thriving moment that this research field is going through. There is a promising diversity, quantity and quality of works. One of the problems in this field is the high degree of dispersion of these works, given that there are few conferences on this specific area. This makes the spread of field-work difficult. But this situation is beginning to change, thanks to conferences such as the "Musical Creativity" Symposium, which was part of AISB'99, and the workshop on Genetic Algorithms in Visual Art and Music (GAVAM), which was part of GECCO'2000. Such events will also enable a closer collaboration among researchers.

3 Outstanding questions.

One of the aims of the workshop on which this special issue is based was not just to allow people to talk about and hear about developments in this exciting area, but also to try to identify important questions which span this whole area of research. As indicated by the review above, much work has been done in this field, and we feel that there is a sufficient body of work from which to draw together some common experiences and consolidate the major questions in this area. In this section we discuss some of these questions.

One difficulty which is found in the musical applications of GAs which isn't found in visual applications is that it takes a certain amount of time to listen to the various members of the population. This makes assigning a score to members of the population somewhat time-consuming. Trying to resolve this problem is an important question for the application of these techniques. This is less of a problem for visual applications, as the user can examine more than one picture at once, and can compare two images directly by looking from one to the other. Nonetheless comparing and rating visual items over a period of time can be somewhat tiring.

3.1 Fitness bottlenecks.

One approach which shows promise is to wrap the scoring system up in the context of an activity. This idea is demonstrated in a number of the papers in this section. A musical example is Biles's *GenJam* system, where the human plays a passage, then the computer plays (whilst being rating by the human performer or the audience), then the human plays again, and so on. This involves the listener(s) in an experience which is much closer to ordinary music-making than a process where the listener has to rate each member of a population in a separate process. A piece of work which similarly embeds the evolution of a visual objects in a context is described in the paper of Rowland and Biocca, which places their evolving objects in the context of a virtual sculpture park. A future direction for such models would be to explore fitness creation in an indirect way, for example a sculpture park which created new sculptures based on those that the user has paid most visual attention to in a virtual environment.

Another approach which may have promise, which is not explored in any of the papers here, would be to preprocess the population so that the user would get an idea of the scope of what can be generated by the algorithm in advance. It may be possible to do this using a non-interactive genetic algorithm or other search technique. In this case we would construct the fitness automatically by giving a high score to those elements of the population which are most different from the previous population members. Clearly defining "difference" would be the biggest challenge here. The aim of this would be to drive the population around the search space, and reconstructing this tour would enable the user to see a sampling of the different types of object that the search space could produce.

There are other potential approaches. One could involve using some kind of machine learning mechanism to monitor the user choices and attempt to "second guess" the user. This could use a learning method, such as a neural network—however one study [49] has shown that it is difficult to train a neural network to learn some features of music, due perhaps to the complexity of the interplay between individual musical gestures and the different meanings of those gestures in different contexts. Alternatively this could involve some form of analysis of the chosen objects and attempt to draw out features common to many of the chosen objects.

Another alternative would involve starting with objects of low complexity, evolving a low complexity object, then evolving more complex features later; this would avoid the problem of evolving detailed features which only work in context, only to have those detailed features vanish when we change the large-scale structure of the object. Reducing the load (or apparent load) on the user, whilst not reducing the quality of the exploration, is a major area of exploration in this area on which research has only just begun to touch.

3.2 Agents.

Another area with much of potential for future research has been identified in Biles's paper. His system, unlike most systems of this type, evolves an *agent* which goes on to create pieces of music, rather than working directly in the artistic medium itself. This idea of evolving agents to carry out the task is also explored in the paper by Santos et al. This allows us to use such systems in different kinds of contexts, e.g. we can train such an agent in private, then use it in a live public performance, as is done with GenJam. We can also imagine a new kind of collaborative musical activity where instead of improvising directly with other musicians, we create agents which represent us in a group music-making activity. This has similarities to computer games such as *Creatures*, where the game works by the player creating autonomous agents and releasing them into an environment, rather than sitting directly at the controls of the agent.

3.3 Comparisons with traditional applications of GAs.

Another interesting question is to what extent these applications of genetic algorithms are similar to traditional genetic algorithms, what the differences are, and how we might design genetic algorithms for these application areas. This remains a largely open field of enquiry. Firstly we can consider how the exploratory nature of many of these applications makes them different from the traditional optimization algorithm? Does this require more diversity in the population? Do we need more mutation so that we keep being presented with new areas to explore, or new operators entirely? Do we need to encode members of the population differently for exploratory-type applications than we do for optimization-type applications?

Secondly we can consider the extent to which the generation of fitness by having a "human in the loop" makes a difference. This makes a number of changes to the algorithm, for example we are no longer guaranteed that when an individual appears twice in a population, that individual gets the same fitness. Indeed there is evidence [63] that the fitness declines as users become bored with the repeated individual. Questions about how fitness gets allocated, how users allocate fitness relative to other individuals and to absolutes, and how we can influence user behaviour so that they make best use of the algorithms, are interesting questions.

3.4 Coherence of populations.

Another question which ties in with theory is the extent to which we can take individuals and drag and drop them between different populations. In Unemi's paper there is the possibility of taking one picture and dropping it into another population. This is used in his application to add diversity to a population which has undergone premature convergence. There is the potential to use this for more subtle effects, e.g. putting certain images aside, and then introducing them into another population to steer that population towards images which are more like the introduced images. This question is discussed further in Johnson's paper.

3.5 Combining sound and visuals.

Another issue is brought up by Unemi's use of both sound and visuals in his software. At present his software only evolves the visuals, the sound being pre-composed. However this opens up the possibility of evolving both sound and music by some kind of joint process. One possibility would be to use evolutionary algorithms for sound and vision, using separate populations but with a common fitness rating given to sound and music played together. A more interesting possibility is to create both from some common source. One way in which this could be done would be to generate the images and sound from the same set of parameters, but using different algorithms. At the crudest level this is little more than lights flashing on and off in beat with the music, but it has potential for a more sophisticated system to be built on it. Perhaps more interesting would be to generate the image and the sound not just using the same parameters, but using the same *algorithm*. This raises the question of whether our aesthetic feelings about a piece of algorithmic art are generated by the algorithm or by the way that algorithm is realized. Are we likely to see some kind of “harmony” between two realizations of the same algorithms through different media? Or is it entirely down to the way that the algorithm is realized? Or a mixture of the two? This revisits questions which were raised a while ago (see e.g. [64]), when fractal pictures were first generated, and musicians attempted to create musical “analogies” of these images by taking the underlying algorithms. The paper by Soddu also discusses such issues. These new ways of exploring these spaces have the potential to provide new tools for the investigation of such phenomena.

3.6 Analysis of music and art.

So far we have discussed the use of GAs in the *creation* of music and visual art. However papers such as that of Federman provide systems which could provide the foundation not just for creation, but also for *analysis*. Federman describes a system which is able to anticipate the next note in a melody. This kind of system could be used in an analytical way, to extract common patterns from different kinds of music, and classify these patterns. There is scope for extracting data from this sort of classification about what the differences are between different kinds of music, and creating a data-mining-driven style of musical analysis. There are also potential applications in the analysis of style, e.g. in the attribution of authorship, or in studying which musicians have influenced others.

4 The papers.

The papers in this special section are drawn from the various talks and demonstrations given during the workshop. The papers give a good balance of current thinking in this area, balancing music and visual art, and balancing descriptions of systems with more speculative sections on the subject as a whole.

The first paper is by Al Biles, who describes recent progress on his ongoing *GenJam* project [14, 19, 49, 65], which aims to create a computer system for jazz improvisation. In contrast to much work on computer music, which attempts to re-examine the nature of music in the light of new technology, Biles is concerned with producing a system which is capable of creating a "straight-up" musical style, and as such we can compare this directly with human performances. In this paper Biles gives a brief outline of how the system works, and then goes on to suggest a number of dimensions along which we can classify this kind of system.

The next paper, by Francine Federman, presents an elegant application of learning classifier systems (LCSs) to developing a system which is able to anticipate the next note in a melody. In particular it studies the effect of different representations on the performance of the system. Searching for representations for musical objects which work well with GAs and similar systems is an important general question in the application of these systems to these domains, particularly as it is being realized that GAs cannot work well without a proper consideration of representational issues (as discussed in e.g. [66]). This suggests a whole world of new ways in which technology could be used in the analysis of musical style, as well as providing a solid foundation for programs such as improvisation systems.

The next three papers demonstrate how ideas from GAs and Alife can be used in computer music making at all levels from the generation of individual sounds to the structuring of whole compositions. The paper by Johnson applies GAs to the lowest level of sound production, that of synthesizing individual sounds to use in electronic music. Learning to program synthesizers is a complex task, and the system described in this paper describes the production of a program which uses GAs to provide a new way to explore the gamut of timbres that can be produced by a complex synthesis system.

The next paper, by Eduardo Reck Miranda, describes two systems, both based around cellular automata. One of these is also a system for controlling the timbre of individual sounds, and creating interesting new timbres. The second works at the level of notes and chords, being a computerized composition system.

The *Vox Populi* system described by Artemis Moroni *et al* is another system which uses GAs as the basis for the generation of whole compositions.

This works on a larger scale still, the population consisting of many sound events within a space which can be explored with the mouse, which in turn gives feedback which drives the evolution of the population. This kind of ongoing, hidden fitness feedback provides a way to avoid the bottleneck of having to alternate between listening to and rating individuals. This paper also discusses a number of theoretical issues.

The next paper by Rowland and Biocca is the first of the papers to tackle visual art. In this paper they describe how to add context to a system by embedding the evolution in a virtual environment, the "Genetic Sculpture Park".

The systems by Miranda and Moroni *et al* both present radically new models for automated composition. The system of Santos *et al* is an agent-based system which draws on a traditional idea—that of a population of "musicians" and a population of critics—in a new way.

The paper by Celestino Soddu considers further the idea discussed above about the extent to which a code representing an artistic idea can be realized in different ways. His "Generative Art" system creates pieces of art which are different every time, but which nonetheless represent different realizations of the same core idea by the artist. His paper both discusses these issues in general, and with reference to his programs which generate striking visual images.

The final paper is by Tatsuo Unemi, and demonstrates another system for the generation of images driven by GAs. It is fitting that this is the last paper in the section, as one application of his system, as was demonstrated in the workshop, is to allow a "video jockey" to create live works of video art alongside a traditional DJ mixing records. This combination of music and visual art through technology is an area fecund with future promise.

We hope that you enjoy the papers in this section.

References

[1] John H. Holland. *Adaptation in Natural and Artificial Systems.* MIT Press, 1975. Second edition 1992.

[2] Melanie Mitchell. *An Introduction to Genetic Algorithms.* Series in Complex Adaptive Systems. Bradford Books/MIT Press, 1996.

[3] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning.* Addison-Wesley, 1989.

[4] Zbigniew Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs.* Springer, 1996. Third edition.

[5] Stephen Wolfram. *Cellular Automata and Complexity.* Addison Wesley, 1994.

[6] Christopher G. Langton, editor. *Artificial Life: An Overview.* MIT Press, 1997.

[7] Michael N. Huhns and Munimdar P. Singh, editors. *Readings in Agents*. Morgan Kaufmann, 1998.

[8] Darrell Whitley, David Goldberg, Erick Cantú-Paz, Lee Spector, Ian Parmee, and Hans-Georg Beyer, editors. *Proceedings of the Genetic and Evolutionary Computation Conference*. Morgan Kaufmann, 2000.

[9] E. Inverson and R. Hartley. Metabolizing music. In *Proceedings of the 1990 International Computer Music Conference*, pages 298–301, San Fransisco, 1987.

[10] P.J. Bentley. Is evolution creative? In Patrizio et al. [67].

[11] Anthony R. Burton and Tanya Vladimirova. Generation of musical sequences with genetic techniques. *Computer Music Journal*, 23(4):59–73, Winter 1999.

[12] P.M. Todd. Evolving musical diversity. In Patrizio et al. [67], pages 40–48.

[13] G. Papadoulus and G. Wiggins. AI methods for algorithmic composition: A survey, a critical view and future prospects. In Patrizio et al. [68], pages 110–117.

[14] John A. Biles. GenJam: A genetic algorithm for generating jazz solos. In *Proceedings of the 1994 International Computer Music Conference*, 1994.

[15] D. Ralley. Genetic algorithms as a tool for melodic development. In *Proceedings of the 1995 International Computer Music Conference* [69], pages 501–502.

[16] B.L. Jacob. Composing with genetic algorithms. In *Proceedings of the 1995 International Computer Music Conference* [69], pages 452–455.

[17] B.L. Jacob. Algorithmic composition as a model of creativity. *Organized Sound*, 1(3):157–165, 1996.

[18] Jeffery B. Putnam. Genetic programming of music. Report, New Mexico Institute of Mining and Technology, 1994.

[19] John A. Biles. GenJam Populi: Training an IGA via audience-mediated performance. In *Proccedings of the 1995 International Computer Music Conference*, 1995.

[20] Damon Horowitz. Generating rhythms with genetic algorithms. In *Proceedings of the 1994 International Computer Music Conference*, 1994.

[21] A. Pazos, A. Santos, J. Dorado, and J.J. Romero. Genetic music compositor. In P. Angeline, Z. Michalewicz, M. Schoenhauer, X. Yao, and A. Zalzala, editors, *Proceedings of the 1999 Congress on Evolutionary Computation*, pages 885–890. IEEE Press, 1999. Volume 2.

[22] A. Pazos, A. Santos, J. Dorado, and J.J. Romero. 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, editors, *Proceedings of the Genetic and Evolutionary Computation Conference*, page 1794. Morgan Kaufmann, 1999. Volume 2.

[23] A. Pazos and J.J. Romero. Musical adaptive systems. In *Genetic and Evolutionary Computation Conference 1999, Research Student Workshop*, pages 343–344, 1999.

[24] J.H. Moore. Gamusic: Genetic algorithm to evolve musical melodies. Windows 3.x Software available on the WWW at `http://www.cs.cmu.edu/afs/cs/project/ai-repository/ai/areas/genetic/ga/%systems/gamusic`, 1994.

[25] Colin G. Johnson. Exploring the sound-space of synthesis algorithms using interactive genetic algorithms. In Geraint A. Wiggins, editor, *Proceedings of the AISB Workshop on Artificial Intelligence and Musical Creativity, Edinburgh*, 1999.

[26] Colin G. Johnson. Exploring sound-space with interactive genetic algorithms. (this volume).

[27] Karl Sims. Evolving virtual creatures. In *Computer Graphics SIGGRAPH'94 Proceedings*, pages 15–22, 1994.

[28] Karl Sims. Evolving 3D morphology and behavior by competition. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 28–39. MIT Press, 1994.

[29] Karl Sims. Artifical evolution for computer graphics. *Computer Graphics*, 25(4):319–328, 1991.

[30] Karl Sims. Interactive evolution of dynamical systems. In F.J. Varela and P. Bourgine, editors, *Towards a Practice of Autonomous Systems : Proceedings of the First European Conference on Artificial Life*, pages 171–178. MIT Press, 1992.

[31] Karl Sims. Interactive evolution of equations for procedural models. *The Visual Computer*, 9(9):466–476, 1993.

[32] M. Haggerty. Evolution by esthetics, an interview with W. Latham and S. Todd. *IEEE Computer Graphics*, 11(2):5–9, March 1991.

[33] S.J.P. Todd and W. Latham. Mutator, a subjective human interface for evolution of computer sculptures. IBM United Kingdom Scientific Centre Report 248, 1991.

[34] Stephen Todd and William Latham. *Evolutionary Art and Computers*. Springer, 1992.

[35] R. Barter. William Latham. *Arts Magazine*, 66(1):92, 1992.

[36] C. Soddu. Recognizability of the idea: the evolutionary process of Argenia. In Patrizio et al. [67], pages 18–27.

[37] Richard Dawkins. The evolution of evolvability. In Christopher G. Langton, editor, *Artificial Life*, pages 201–220. Addison-Wesley, 1989.

[38] Richard Dawkins. *The Blind Watchmaker*. Penguin, 1990.

[39] Peter J. Bentley, editor. *Evolutionary design by computers*. Academic Press, 1999.

[40] D. I. Perrett, K. J. Lee, I. Penton-Voak, D. Rowland, S. Yoshiwaka, D. M. Burt, S. P. Henzi, D. L. Castles, and S. Akamatsu. Effects of sexual dimorphism on facial attractiveness. *Nature*, 394:884–887, 1998.

[41] M.A. Rosenman. The generation of form using evolutionary approach. In D. Dasgupta and A. Michalewicz, editors, *Evolutionary algorithms in Engineering Applications*. Springer, 1997.

[42] Duncan A. Rowland. *Computer Graphics Control over Human Face and Head Appearance, Genetic Optimisation of Perceptual Characteristics*. PhD thesis, University of St. Andrews, 1998. Available at `http://www.mindlab.org/documents/Thesis_Rowland/title.htm`.

[43] P.J.B. Hancock and C.D. Frowd. Evolutionary generation of faces. In Patrizio et al. [67], pages 93–99.

[44] Craig Caldwell and Victor S. Johnston. Tracking a criminal suspect through "face-space" with a genetic algorithm. In Belew and Booker [70].

[45] A.R. Burton and T. Vladimirova. Genetic algorithm utilising neural network fitness evaluation for musical composition. In G.D. Smith, N.C.

Steele, and R.F. Albrecht, editors, *International Conference on Artificial Neural Networks and Genetic Algorithms*, pages 220–224. Springer, 1997.

[46] A.R. Burton. *A Hybrid Neuro-Genetic Pattern Evolution System Applied to Musical Composition.* PhD thesis, University of Surrey, 1998.

[47] L. Spector and A. Alpern. Criticism, culture, and the automatic generation of artworks. In *Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 3–8, 1994.

[48] L. Spector and A. Alpern. Induction an recapitulation of deep musical structure. In *IJCAI-95 Workshop on Artificial Intelligence and Music*, pages 41–48, 1995.

[49] John A. Biles, Peter G. Anderson, and Laura W. Loggi. 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), March 26–28, Reading, UK*, pages B39–B44. ICSC Academic Press, 1996.

[50] P.M. Gibson and J.A. Byrne. Neurogen, musical composition using genetic algorithms and cooperating neural networks. In *Second International Conference on Artificial Neural Networks*, pages 309–313, 1991.

[51] S. Baluja, D. Pomerleau, and T. Jochem. Towards automated artificial evolution for computer-generated images. *Connection Science*, 6(2):325–354, 1994.

[52] S. Baluja, D. Pomerleau, and J. Todd. Towards automated artificial evolution for computer-generated images. In Griffith and Todd [71].

[53] G.A. Wiggins, G. Papadopoulos, S. Phon-Amnuaisuk, and A. Tuson. Evolutionary methods for musical composition. *International Journal of Computing Anticipatory Systems*, 1999.

[54] S. Phon-Amnuaisuk and G.A. Wiggins. The four-part harmonisation problem: A comparison between genetic algorithms and a rule-based system. In Patrizio et al. [68], pages 28–34.

[55] R.A. McIntyre. 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*, pages 852–857. IEEE Press, 1994.

[56] G. Papadopoulos and G.A. Wiggins. A genetic algorithm for the generation of jazz melodies. In *SteP'98*, 1998.

[57] Andrew Horner and David E. Goldberg. Genetic algorithms and computer-aided music composition. In Belew and Booker [70].

[58] K. Thywissen. Genotator: An environment for investigating the application of genetic algorithms in computer assisted composition. In *Proceedings of the 1996 International Computer Music Conference*, 1996.

[59] G.M. Werner and P.M. Todd. Too many love songs: Sexual selection and the evolution of communication. In *Proceedings of the Fourth European Conference on Artificial Life*, pages 434–443. MIT Press, 1997.

[60] P.M. Todd and G.M. Werner. Frankensteinian methods for evolutionary music composition. In Griffith and Todd [71].

[61] K. McAlpine, E. Miranda, and S. Hoggar. Making music with algorithms: A case-study system. *Computer Music Journal*, 23(2):19–30, 1999.

[62] Eduardo Reck Miranda. Granular synthesis of sounds by means of a cellular automaton. *Leonardo*, 28(4):297–300, 1995.

[63] Tony D. May. Music and computers: The design and implementation of a musical genetic algorithm. Master's thesis, University of Kent, 2000.

[64] Brian Evans. Integration of music and graphics through algorithmic congruence. In *Proceedings of the 1987 International Computer Music Conference*, pages 17–24, 1987.

[65] John A. Biles. Interactive GenJam: Integrating real-time performance with a genetic algorithm. In *Proceedings of the 1998 International Computer Music Conference*, 1998.

[66] Andrew L. Tuson. *No Optimisation Without Representation*. PhD thesis, University of Edinburgh, 1999.

[67] Andrew Patrizio, Geraint A. Wiggins, and Helen Pain, editors. *Proceedings of the AISB'99 Symposium on Creative Evolutionary Systems*. Society for Artifical Intelligence and the Simulation of Behaviour, 1999.

[68] Andrew Patrizio, Geraint A. Wiggins, and Helen Pain, editors. *Proceedings of the AISB'99 Symposium on Musical Creativity*. Society for Artifical Intelligence and the Simulation of Behaviour, 1999.

[69] *Proceedings of the 1995 International Computer Music Conference*, 1995.

[70] Richard K. Belew and Lashon B. Booker, editors. *Proceedings of the Fourth International Conference on Genetic Algorithms*. Morgan Kaufmann, 1991.

[71] N. Griffith and P.M. Todd, editors. *Musical Networks: Parallel Distributed Pereption and Performance*. MIT Press, 1998.

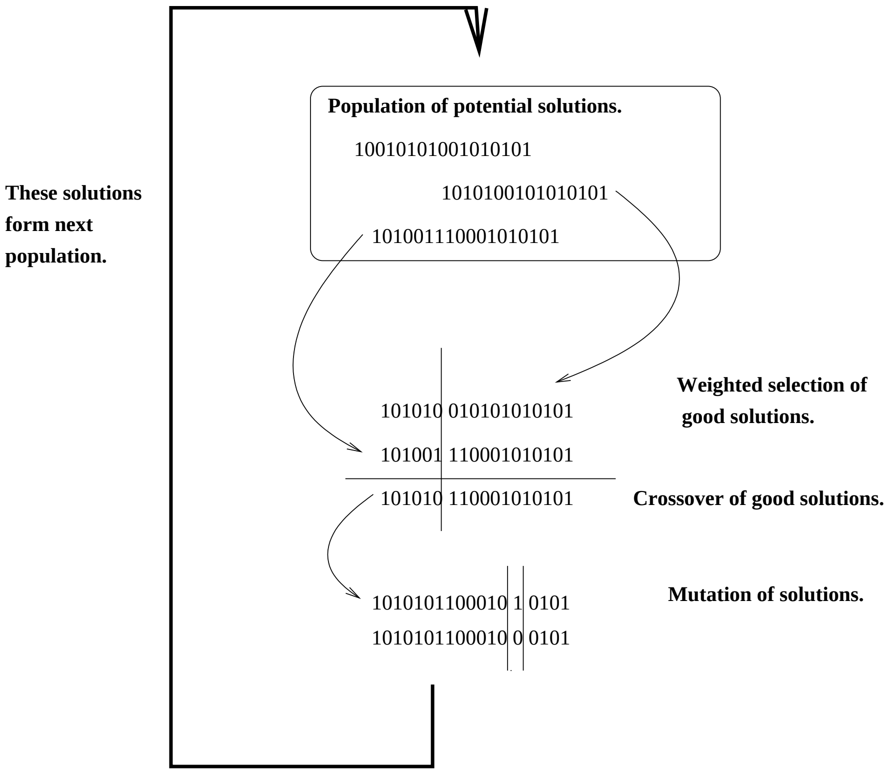
Population of potential solutions.
10010101001010101
1010100101010101
101001110001010101
These solutions
form next
population.
101010 010101010101
101001 110001010101
101010 110001010101
Weighted selection of
good solutions.
Crossover of good solutions.
1010101100010 1 0101
1010101100010 0 0101
Mutation of solutions.

Figure 1

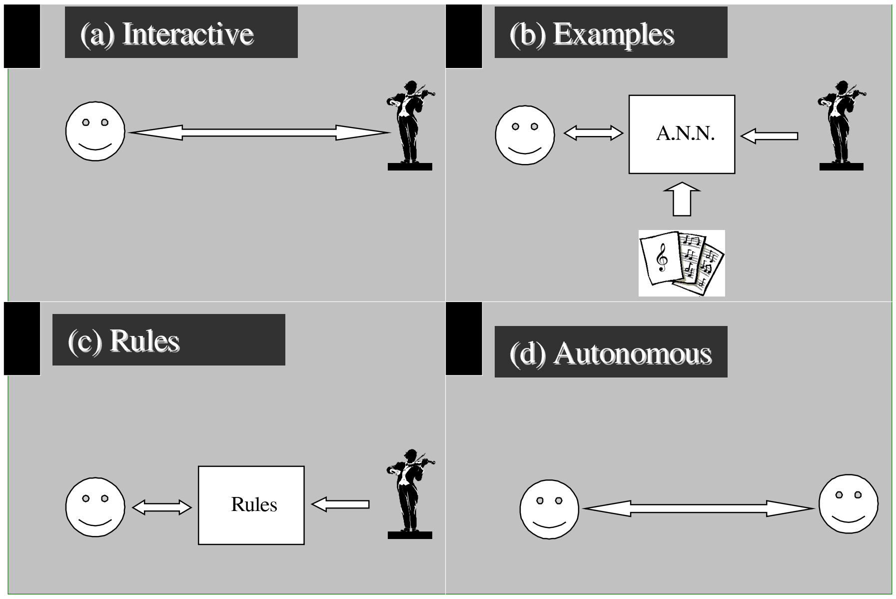

figure 2

Figure captions.

Figure 1. Genetic algorithms—a summary.

Figure 2. Different ways of interacting with the evolutionary composer oriented to musical works. (a) The user acts as critic of the system's compositions. (b) The user introduces a series of examples in order to train the ANN, which will work as a critic of the evolutionary system's compositions. (c) The user defines a set of rules used to evaluate the system's compositions. (d) In this case, composer and critic are part of the system,and they evolve simultaneously. (e) Finally, the use of different paradigms in a single system.