Fitness in Evolutionary Art and Music: A Taxonomy and Future Prospects

Colin G. Johnson

School of Computing
University of Kent
Canterbury, Kent, CT2 7NF
England
Email: C.G.Johnson@kent.ac.uk

Abstract: This paper is concerned with the idea of *fitness* in art and music systems that are based on evolutionary computation. A taxonomy is presented of the ways in which fitness is used in such systems, with two dimensions: what the fitness function is applied to, and the basis by which the function is constructed. A large collection of papers are classified using this taxonomy. The paper then discusses a number of ideas that have not been used for fitness evaluation in evolutionary art and which might be valuable in future developments: memory, scaffolding, connotation and web search.

1 Introduction

Evolutionary algorithms—computer systems that take their inspiration from Darwinian evolution over many generations—have been an inspiration for many artistic and musical projects (Burton and Vladimirova, 1999; Miranda and Biles, 2007; Romero and Machado, 2008). One of the key aspects of evolution is that it is guided by a notion of *fitness*—i.e., some way of determining which individuals from the population survive and breed. The aim of this paper is firstly to give a broad survey and taxonomy of the ways in which artistic and musical applications of evolutionary computing have implemented this concept of fitness, and secondly to highlight some ways of measuring fitness that could guide novel directions in the artistic use of evolutionary algorithms.

McCormack (2005) has argued that there is a need for more "theory" in evolutionary art, and particularly more "art theory" rather than computational theory. Part of the aim of this paper is to contribute to that art theory, in particular by looking in Section 3 at ideas that are common in discussions of artworks as such, or of the creative process that gives rise to those works, and revisiting them in the context of evolutionary art, and in particular how these artistic ideas can be explored through fitness measures.

In particular, we are interested in whether these systems can give rise to creative outcomes. Pinning down computational creativity is difficult—a common working definition is "the study of building software that exhibits behavior that would be deemed creative in humans" (Colton et al., 2009). For the purposes of this paper we consider systems that attempt to generate some outcome that is artistically worthwhile, whether that is by evolving a work or collection of works that have some value, or by the results of the process of evolution itself being regarded as having some value.

This paper is an extended version of an earlier conference paper (Johnson, 2012c).

2 Taxonomy and Survey

This section of the paper gives an overview and taxonomy of how fitness has been used in a large number of evolutionary art and music projects.

The source material for this is a large collection of papers that used evolutionary methods to produce a creative outcome, i.e., something was generated as a result of the process or end result of the evolution, that was not merely the result of applying well-understood rules or constraints. The source material for this collection were the papers published between 2003 and 2012 in the EvoMusArt workshops, symposia and conference (Cagnoni et al., 2003; Raidl et al., 2004; Rothlauf et al., 2005, 2006; Giacobini et al., 2007, 2008, 2009; Di Chio et al., 2010, 2011; Machado et al., 2012b), which have been held every year and are the main event focused on evolutionary methods in music, art and design, together with a number of other conference papers, book chapters and journal papers on the topic. A full list of references is given in Table 1 below.

Papers that were *not* included in the collection included: review papers, papers about the theory of evolutionary art and music, papers that used evolutionary methods for the analysis of works of art or music, and papers that were about performance, rendering or re-presentation of existing works. Papers that used evolution solely as a way of implementing fixed rules for some artistic or musical system were excluded—the canonical example of these were formal rules for musical harmony. These were excluded because the aim of this study was to look into how evolution an give rise to novel, creative outcomes. However, we accept that the border between following fixed rules, and providing some kind of aesthetic measure, is vague. Papers that presented systems with a creative outcome but where no notion of fitness was involved were excluded. Most of the systems were evolutionary algorithms but some other similar algorithms, e.g. swarm algorithms, with a clear fitness function, were included.

Based on these papers we constructed a two-dimensional taxonomy of how fitness is used in this domain, consisting of a dimension called *fitness scope* and a dimension called *fitness basis*. A definition and analysis of this taxonomy, and a tally of the number of papers that fitted into each point in the taxonomy, is given in the remainder of this section.

2.1   Fitness Scope

The first part of the taxonomy is a classification of what fitness is applied to: call this the *fitness scope*. This classification contains three classes. The first, call this a *set of works*, is where each member of the population consists of a collection of individual works of art or music, and the fitness measure is applied to that collection. The second, *whole work*, is where the population consists of examples of works of art or music (or a proxy therefor, like a set of generative rules for creating such a work), and therefore fitness is applied to a single work. The final, *evolutionary process as artwork* is where the fitness evaluation is part of a process that is viewed as the work itself, and therefore the work as such is not being rated by a fitness measure.

The distribution of papers in each of these categories is given in Table 1. Note that it is possible for a paper to belong to multiple categories.

The *set of works* category is represented by just two papers. The first paper (Bird et al., 2003) is unusual, in that it is concerned with curating a collection of pre-created works, by deciding which works are to be presented together at a number of locations. The second (Machado et al., 2010) is a meta-level evolution of generative systems. Each individual in the population represents a generative system for creating artworks, and these are then judged not by creating a single example but a collection of examples, either via interaction with a human or via a proxy measure for aesthetic value such as lack of compressibility.

The most common way in which fitness scope is handled is that the individuals in the population each represent one example artwork: what we termed *whole work*. This has its origins in very early examples of the use of evolutionary processes in art, for example in Dawkins's *biomorphs* (Dawkins, 1989). Typically, the evolutionary process is carried out and then the fittest example in the final generation (or the fittest overall) is output.

Usually, the aim of such processes has been to actively search out "good" works of art. However, Reddin et al. (2009) and Dahlstedt (2007a) have argued that the role of the fitness function could be just to eliminate bad examples, leaving the "expressive power of the representation" (Reddin et al., 2009) to provide the good material (this somewhat prefigures the arguments made below about endogenous fitness functions). Some examples in this taxonomy produce tools or components that are then used by humans in the creation of more substantial works.

A particular kind of this is where the process is repeated a number of times within one piece of work, perhaps with some linkage. For example, Biles's *GenJam* jazz performance software (Biles, 2007) uses MIDI-capture to "listen" to a human performer, then evolves variants on what it has heard, alternating between human and computer phrases throughout an extended improvisation session. Members of the population are potential phrases to be played, and the fitness function—based on ratings by a human—can rate these differently as the performance progresses.

Fitness Basis \ Fitness Scope	Set of Works	Whole Work	Process as the Work	TOT.
Aesthetic Measure	(Bird et al., 2003; Machado et al., 2010)	(Greenfield, 2004; Manaris et al., 2003; Svangård and Nordin, 2004; Collomosse and Hall, 2005; Manaris et al., 2005; Khalifa and Foster, 2006; Santarosa et al., 2006; Campolongo and Vena, 2006; Bilotta et al., 2007; Greenfield, 2008, 2009; Reddin et al., 2009; Murray Browne and Fox, 2009; den Heijer and Eiben, 2010; Li and Hu, 2010; Oh and Zajec, 2010; De Prisco et al., 2011; Romero et al., 2011; den Heijer and Eiben, 2011; Reynolds, 2011; Ekárt et al., 2011; Kaliakatsos-Papakostas et al., 2012; den Heijer and Eiben, 2012b,a; Greenfield, 2012; Bergen and Ross, 2012; Machado et al., 2012a; Li et al., 2012; Horowitz, 1994; Manzolli et al., 1999; Brown, 2004)	(Blackwell and Jefferies, 2005; Greenfield, 2005, 2006; Hochreiter, 2006; Bown and Lexer, 2006; Machado and Graca, 2008; Lopes and Urbano, 2012; Eigenfeldt and Pasquier, 2012)	41
Human Interaction	(Machado et al., 2010)	(Todd and Latham, 1992; Nemirovsky and Watson, 2003; Lutton et al., 2003; Lewis, 2004; McCormack, 2004; Draves, 2005; Basa et al., 2006; Unold and Troc, 2007; Ando et al., 2007; Hart, 2007; Oscar Sanjuán and Luque, 2007; Hiroaki Nishino and Utsumiya, 2007; Machewe and Parmee, 2007; Tsuneo Kagawa and Utsumiya, 2007; Machado and Graca, 2008; Hoover et al., 2008; Romero et al., 2009; Moniz and Jacob, 2009; Dahlstedt, 2009; Dubbin and Stanley, 2010; Li and Hu, 2010; Oh and Zajec, 2010; Vouliouri, 2011; Byrne et al., 2011; Eisenmann et al., 2011; Colton et al., 2011; Nicolau and Costelloe, 2011; Hornby and Bongard, 2012; McDermott, 2012; Li et al., 2012; Byrne et al., 2012; Kaliakatsos-Papakostas et al., 2012; Dahlstedt, 2001; Biles, 2007; Biles and Eign, 1995; Horowitz, 1994; Johanson and Poli, 1998; Thywissen, 1999; Collins, 2002; Ralley, 1995; Machado and Cardoso, 2002)	(Nemirovsky and Watson, 2003; Mandelis, 2003; Blackwell and Young, 2004; Nemirovsky and Luger-Guillaume, 2004; Yee-King, 2007; Dahlstedt and Nilsson, 2008; Jones, 2008; Machado and Graca, 2008; Dahlstedt, 2009; Eigenfeldt, 2009; Shao et al., 2010; Moroni and Manzolli, 2010; Suzuki et al., 2011)	56
Corpus		(Manaris et al., 2003; Campolongo and Vena, 2006; Somnuk Phon-Amnuaisuk and Kuan, 2007; Hervás et al., 2007; Law and Phon-Amnuaisuk, 2008; Wolkowicz et al., 2009; Phon-Amnuaisuk, 2009; Ortega de la Puente et al., 2002)		8
Seeds and Targets		(Fornari et al., 2007; Rizzuti et al., 2009)	(Gartland-Jones, 2003; Nemirovsky and Luger-Guillaume, 2004; Blackwell and Jefferies, 2005; Hochreiter, 2006; Magnus, 2006; Bown and Lexer, 2006; Tufte and Gangvik, 2008; Ariza, 2002; Horner and Goldberg, 1995)	11
Endogenous			(Dorin, 2005; Urbano, 2005, 2006; Marcelo Caetano and Zuben, 2007; Greenfield, 2006; von Mammen et al., 2008; Uozumi, 2007; McCormack and Bown, 2009; De Smedt et al., 2011; Bird et al., 2008; Jones, 2008; Eldridge and Dorin, 2009; Filipe Antunes and Fol Leymarie, 2012; Greenfield, 2010; Dorin, 2009; Eldridge et al., 2008; Wakefield and Ji, 2009; Eigenfeldt, 2009; Fornari, 2011; Eigenfeldt and Pasquier, 2011; Fornari, 2012; Waschka II, 2007)	22
Critics/Co-evol.		(Romero et al., 2003; Greenfield, 2004; Machado et al., 2004; Romero et al., 2009; Li and Hu, 2010; Law and Phon-Amnuaisuk, 2008; Colton, 2008; Zhang and Miranda, 2007; Johanson and Poli, 1998; Hornby and Bongard, 2012)		10
TOTAL	3	92	52	

Table 1 How the papers studied fit into the various definitions of fitness scope and fitness basis.

These have been categorised with the "whole work" scope, as the process is similar to a micro-scale, iterated version of what is happening in other examples of that scope.

The final type of fitness scope is where the whole process (or, occasionally, some trace of its outcome) represents the final work. We will term this *evolutionary process as work*. Usually, this is a time-based work, e.g. an evolving animation or ongoing piece of music, with just the current generation being displayed at any one point. However, some works in this class display just the end result of the process: this is illustrated, for example, in a piece by McCormack and Bown (2009), where agents deposit lines on a surface during the evolutionary process, and the final disposition of the lines generated throughout the evolutionary history forms the work.

2.2 Fitness Basis

The second part of the taxonomy considers how the fitness is evaluated: call this the *fitness basis*. This is divided into six classes: aesthetic measure (a fixed function measuring the quality of the solution), human interaction, the use of a corpus of material, the use of a guiding example, the use of some kind of seeds or targets, an endogenous or implicit fitness derived from interactions between agents, and the use of a population of critics that learn alongside the evolutionary process. The total number of papers falling into each of these categories is given in Table 1.

Again, some papers fell into more than one category. As an example, the work by Li and Hu (2010) uses notions of aesthetic measures, which are tuned by interaction with a user, which are in turn proposed as artificial critics for future works.

The remainder of this section discusses these six categories in detail.

Aesthetic Measure

The first basis for fitness is some notion of *aesthetic measure*. That is to say, there is a function which, when applied to a member of the population, measures the quality of that member as a piece of art or music. Most commonly, this function is constructed by a human system designer deciding what is to be regarded as "good" or "bad". For example, in the paper by Bilotta et al. (2007), melodies receive a score based on a sum of positive values for desirable features of a melody (e.g. commonly-used intervals) and negative values for non-desirable features. An appropriately constructed measure can be applied to multiple artworks, as in the curation example discussed earlier (Bird et al., 2003), where the measure measured the balance of works at each of a number of locations via similarity. The construction of such a measure has been discussed by Greenfield (2009), where a number of different measures were experimented with and rejected or adjusted by the system designer.

Is optimizing against a fixed aesthetic measure a creative act by the computer system? It is clearly restrictive—something that is creative, but in a way that was not anticipated by the measure chosen, will never be favoured by such a scheme. However, it is not always the case that authors of papers in these categories intended the measure to be a *universal* aesthetic measure—just something that provides enough discrimination for some creative work to be generated. For

example, the paper by Hervás et al. (2007) makes it clear that the aim of the algorithm is modest—to produce "alliterative drivel"—and we should not criticise such focused works because they fail to leap out of their designed purpose.

Similarly, Murray Browne and Fox (2009) explore a particular aesthetic pattern (expectation-violation) in the context of producing work in a particular style. The work by De Prisco et al. (2011) evolves serialist musical compositions against a set of rules for such compositions. These get close to the sort of work that we have chosen not to include here—the "uncreative" application of GAs to e.g. learning musical harmonisation in a fixed style. It seems that we can make a useful distinction between measures that are about more abstracted aesthetic characteristics (e.g. "complexity") and those that are about simple technical concepts (e.g. "avoiding parallel fifths"); but, there is a continuum and some characteristics will inevitably fall in the middle.

An interesting paper by Machado et al. (2012a) uses visual similarity to faces as its basis for fitness. This can be seen as a simple aesthetic choice—the aim of the system is to evolve things that look face-like as that is an aesthetically desirable outcome. However, the actual results of the process are much more interesting, as many of the images evolved by the algorithm look nothing like faces. Instead, the system evolved images that exploit pathologies in the face-recognition algorithm used to make fitness decisions. The resultant images have a conceptual depth that is much greater than the mere evolution of obviously face-like objects. This emphasises the point made by Johnson (2012b) that, when computational creativity becomes sufficiently rich, art that has something to say about its own embeddedness in a technological domain might become interesting (rather than just being the "technical fetish" explored by McCormack (2005)).

One aspect of creative development that aesthetic measures fail to address is the change in what is regarded as creative over time. In light of the idea that computational creativity is behaviour that would be regarded as creative if achieved by humans, it would be interesting to apply these measures to existing human-produced artworks, or more interesting still to a sequence of artworks that are recognised in art history as working from the beginnings to the pinnacle of some artistic style or technique.

Another critique of such approaches (though, this could be seen as a criticism of fitness applied to creativity as a whole) is that there is an assumption that an aesthetically valuable outcome can be achieved by passing through a succession of aesthetically less valuable outcomes. This will be discussed further below.

Some papers use the concept of an aesthetic measure, but it is one that is *learned* from a corpus of examples; this is discussed in the discussion of corpus-based methods below.

Human Interaction

This basis is where the fitness of individuals in the population is guided by one or more people. Many of these are of the generational type, where the population at each generation is displayed to the user of the system, and the user then scores or ranks the population members. Clearly, there is scope for this "fitness function" to be adaptive over time; contrasted to a fixed fitness function, the user can change their mind over the course of the run about what aspects of the outcomes they

prefer. Furthermore, there is the capability (indeed, the expectation) of rewarding surprises generated by the system; the user can pick out something that they would not have expected in advance of the run and reward it.

The usual aim of this is to find a final outcome that will then be displayed or performed as the artwork. Another approach (as illustrated by Nemirovsky and Watson (2003), Dahlstedt and Nilsson (2008) and Eigenfeldt (2009)) is to use this during improvised performance, where the user is exploring the space using interactive evolution and making use of promising solutions as they are created. This has a flavour of the *evolutionary process as work* scope above, but also the *whole work* scope: not everything generated is used, as the improviser listens to material being generated, but doesn't need to make everything available to the audience, nor do they need to use it in the order generated (this, and other ideas of "temporary storage", have been discussed by Dahlstedt (2009)).

An alternative to conscious human selection is to use a human-in-the-loop but use some feature of that person's reaction as the fitness guide. For example, eye-tracking or a physiological measurement could be used. This is illustrated by the work of Basa et al. (2006), who use measurements of EEG signals as the basis for a fitness function which distinguishes "positive" emotional states, and then uses this to evolve art that provokes these states.

Other work, such as that by Biles and Eign (1995), takes input from a population of users rather than a single rater, e.g. taking input from an audience either in a single rating action or else over the course of an exhibition or performance.

Human interaction can also consist of a system responding to human sounds or gestures, rather than the direct provision of fitness scores or rankings to the system by the human. For example, the JaVOX/AURAL system by Moroni and Manzolli (2010) contains robots whose movements guide the generation of sound, and the movement of those robots is guided by the movement of people within their environment. This relates to ideas of Lewis (2000), who envisages a computer music improvisation system which interacts with humans by being an equal partner with human players and receiving information from those players by building and understanding of what they play, rather than getting direct positive/negative feedback. However, in the context of GA systems, this does then provoke the question of how this gets converted into fitness; this can end up just being naive, such as measuring fitness by similarity to what how the human is playing, whereas we would like to encourage a more sophisticated form of interaction.

Another way in which human interaction has been used is in building up a model of the user, which then exercises fitness choices, rather than (or as well as) those choices acting directly on the population. A good example of this is the paper by Hornby and Bongard (2012), where user preferences are used to generate a model of that specific user by training an ensemble of neural networks—similar ideas have been used by Johanson and Poli (1998). The aim of this is to allow the user model to make some decisions in place of the user once it has developed sufficient confidence in its predictions, and therefore allow valuable user effort to be focused on more subtle discriminations.

One problem with human-interactive fitness evaluations is that they can be tedious, with the human interactor having to sort away a large number of examples

that are very far from being acceptable. One approach to address this has been proposed by Dahlstedt (2007b) (and prefigured by Horowitz (1994)), which is to use a set of basic guidelines for rejecting population members, and only present to the human evaluator those members that get through this initial filter. This can be seen as a kind of crude aesthetic measure; perhaps more accurately, an *anti*-aesthetic measure that filters out those individuals that have very low merit or relevance.

Corpus

Some fitness functions are guided by an existing example or corpus of material. There are a number of different ways in which this basic idea has been used.

One way in which this can be used is in providing a set of exemplars for the evolutionary process to use as "good" examples. Typically, a fitness function is *learned* from this corpus, which is then used to guide an evolutionary process, with no direct reference to the corpus whilst the evolutionary process is running. There is no reason why this abstraction should be so—a system based on instance-based learning (Daelemans and van den Bosch, 2005) might well be plausible—but, most systems to date have worked with this layer of abstraction, perhaps to provide greater generalisation.

For example, Phon-Amnuaisuk et al. (2007) use a self-organising map to learn features from a corpus of music, and then learn new music which exhibits similar features. Sometimes, more structure is put into the design of the fitness function by the system designer. For example, in the work of Manaris et al. (2003), the fitness function is in the form of a Zipf-Mandlebrot law, the parameters of which are tuned so that the function would give a high value when applied to a specific corpus of musical material. This latter approach combines the corpus-based approach with the aesthetic measure approach discussed above; the system designer is asserting that the functional form of the fitness function is a meaningful measure, and then the parameter details are learned from the corpus.

It could be argued that this approach is rather limited. Whilst it is fine for learning to imitate previous examples, it doesn't provide any guidance as to how to go beyond the kinds of things illustrated in the corpus. Perhaps a very sophisticated learning system would be able to generalise in a very high-level way about features of these good examples and apply them to a radically different example—for example, it would be interesting to see if such learning could be applied to assessing fitness of a different artform. There are similarities here with the arguments of Evans (1987) who discusses the idea of whether two generative art algorithms based on the same underlying algorithm would provide the same aesthetic impact; however, he is concerned with the whole generative process, whereas we are focused on fitness.

Seeds and Targets

A contrasting use of the use of example sounds or images is to use examples as some kind of "seeds" or "targets" to guide the evolution in some form. Often, this is combined with the Evolutionary Process as Artwork scope, where the evolutionary progress towards the target that is viewed as the work. For example, in the sculptural piece *Trans<–>Former #13* (Tufte and Gangvik, 2008), an

articulated sculpture is given a number of target configurations, and applies a genetic algorithm to learn how to move between these configurations. In the *MusicBlox* system (Gartland-Jones, 2003), a target piece of music is provided, and the evolutionary algorithm learns a path from a random starting population towards this target. However, in this example, the aim is "to explore the space, and find musically interesting solutions" (Gartland-Jones, 2003)—and samples from this exploration are then chosen by a human user for incorporation into the final work. In both of these examples, the targets are *not* provided as exemplars, just as a way of provoking the system into doing something that might be of interest.

Endogenous

One class of evolutionary art systems are those where the fitness measure is implicit or endogenous. That is, an evolutionary process is set up where the fitness of an individual depends upon its context in some artificial environment. This is invariably, in the papers studied here, found in conjunction with the presentation of the whole process as the artwork (see Table 1): what is presented to the audience is an unfolding of the evolutionary process.

A good example of this is given in the paper of Bird et al. (2008), where a drawing robot is provided with a fitness function that "correlated the changes in state of their line detector and pen position" (Bird et al., 2008). That is, the fitness function was not taking an "outside look" at the drawing (being) produced and assigning a score to it; instead, it is rewarding local behaviour. There are many similar projects, including many examples of "a-life art" where some simulation of (an abstraction of) life processes is visualised or used to generate sound (see e.g. the ecosystem models of Bown and McCormack (2010)).

Clearly, this contrasts with the three previous methods in that the aim of the evolutionary process is *not* to guide the creativity of the system. The evolution is not evolving towards "better works of art" or "more creative outcomes". Instead, the creativity of the process is in the human decision of which interactions will make an interesting work—but, unlike most artistic decisions, the results of this are emergent and hard to predict.

Another way of seeing this is that this is an implied *aesthetic measure*—that is, the designer of the system has decided that the kind of complexity found in nature is aesthetically valuable or otherwise engaging, and rather than measuring the closeness of evolved systems to this complexity, has built this in by design. As such, this kind of work provides one of the few examples where evolutionary art references anything in the "outside world"—this kind of art often makes us think about life processes in the world, to ask questions like "what is special about biological life?".

It might be interesting to explore the idea of a meta-level system that used an evolutionary system with explicit creative aims to evolve the structure and parameters of a second level of evolutionary systems that are themselves considered as the works themselves.

One of the barest examples of this is *GenDash* by Waschka II (2007). This consists of a crossover-only genetic algorithm with no fitness measure, the chance of any two individuals being crossed over being equal. Individuals consist of short musical fragments, and the pieces created are the unfolding of this evolutionary

process with each fragment being played as it is created. The author points out that this means that the initial population becomes much more important than in a traditional genetic algorithm; indeed, the way in which the algorithm produces significantly different pieces is by the hand-curation of the starting population rather than by creating a fitness function.

It is interesting that many of the examples of works where the evolutionary process was displayed as the art itself were of this type—there is no *a priori* reason why a goal-directed evolutionary process could not be displayed in this way, apart from the desire to have an ongoing process to display rather than something that reaches a stopping point and then needs to be restarted. *Trans<->Former #13* (Tufte and Gangvik, 2008) discussed above is perhaps the closest to this; within the musical domain, the target-directed evolution of electroacoustic music by Magnus (2006) also.

Critics and Co-evolution

A small number of papers used two activities which developed together during the evolutionary process. The first activity was the generation of the artwork itself, whilst the second set provided some notion of *criticism* or *theory-formation*.

The aim of these papers is to emulate some of the process of aesthetic development that occurs during the development of an artform, whereby artworks are created and critics make comments and evaluations on these works, which then provoke artists to make different work, which *changes the perception of the critics* so that they make different evaluations. For example, in the work by Romero et al. (2003) critics are proposed that build through a number of stages—starting from author and style recognition, building through aesthetic measure induction from a corpus, to developing a society of critics—perhaps including both human and computer critics (Romero et al., 2009)—who are able to evaluate artworks being created via another system. However, ideas for how the later stages of this process would be implemented are not given in much detail. A later paper (Machado et al., 2004) gives some implemented examples of the earlier stages—but, these are still essentially classification algorithms.

An attempt to develop a meta-level automated analyst of automatically generated images has been carried out by Colton (2008). This is the closest attempt that has been made to the kind of critical reasoning discussed earlier. The system begins by generating a number of random images (within a particular high-level style). Features of these images are then given to a theory-formation engine, which attempts to abstract theories that explain links between these images. These theories are then trialled as fitness functions for the generation of new images. The aim is that the reasoning system will pick out some coherent theory in the production of the images. Perhaps this reflects how some students of art work—they begin by attempting a standard task, and realise that certain things are common to their works, and decide to refine these further.

This is an interesting piece of work, but we can criticise it in a couple of ways. Firstly, the theories developed are not guided by any aesthetic *appreciation* of the works created—the engine is just extracting a compact *description* of the works (though is *human* appreciation of this any more sophisticated?). Secondly, it could be argued that this is just a rather abstracted form of aesthetic measure—the

author is arguing that *ability to form a coherent theory* is why we should choose one set of works over another. This is not unreasonable—but, we should acknowledge that we have made this choice explicitly. Nonetheless, there may always need to be a point at which we need to "cash out" and apply some kind of (suitably abstract) aesthetic measure if we are ever to make a fitness evaluation.

3   Future Possibilities

This final section of the paper turns its attention to what *hasn't* been used as the basis for fitness in evolutionary art and music. More specifically, what features of artworks, or of the process of creating art, have not been considered, or very little considered, in the body of work on evolutionary art? In the remainder of this section we consider a number of topics—memory, scaffolding, connotation and web search—that could provide a meaningful basis for fitness computation in evolutionary art.

3.1   Memory and Scaffolding

Evolutionary algorithms, including those used for the creation of evolutionary art, are typically memoryless. Each member of the population is evaluated relative to a fixed fitness function or by comparison with other current population members. The members of previous generations have no influence on the current population. This contrasts with the process that would be carried out by a human artist or designer, who might well use their memory of the entire process-so-far to influence decisions in the current stage of the artistic/design process. For example, a component that does not fit with the current prototypes might be revisited at a later stage when an opportunity to fit it in with a later version of the overall design occurs.

A related aspect is the need for every aspect of the work to be a part of some highly-rated member of the current population in order to survive into the next generation. Many human creative processes begin with a process of exploration or rumination (Webb Young, 1943), where the designer or artist explores a large number of concepts and their connections before bringing these together to form a design or artwork. This idea of *bringing together* is different from standard notions of crossover in that the individual components do not have to be a part of an already coherent individual in order to be crossed in.

When we are creating something—this could be something that we know how to build, or this could be a creative process designed to build something new—we often use objects that are not part of the final design, but which somehow support the process of getting to that design. This is a large category of things, including everything from physical supports to mental concepts.

We could envision a population structure consisting of an interconnected set of different kinds of elements: potential artworks, components thereof, and *scaffolding* components that might be considered important (e.g. intangible components such as moods or emotions, which have meaning by being connoted by concrete components in the work). Parts of this structure would gain or lose fitness according both to measures such as those discussed above (i.e. an aesthetic

measure or interaction with a human), but would also gain or lose fitness by being associated with other components in the network. Eventually, unfit components would be removed, and fit components duplicated, mutated and these mutants connected into the structure.

This concept of giving a fitness value to components by virtue of their role in a network rather than by their value in isolation has been explored by Berthold and colleagues (Nagel et al., 2011; Kötter and Berthold, 2011), and the idea of a fitness measure based on what a population member *connotes* as well as what it *denotes* has been explored elsewhere (Johnson, 2012a). This kind of exploration incorporates many of the valuable aspects of how human designers and artists carry out the creation process; however, this is not to say that the aim of an evolutionary creative exploration is to replicate point-for-point the human-based design/creation process.

This idea of scaffolding represents a break with the tradition in evolution of gradual development through putative solutions. Evolution has been presented, e.g. by Dawkins (1997), as a steady climb up a smooth gradient, by contrast with the saltational leap up the cliffside of creationism. This is a good analogy for understanding biological evolution, where every component required at every generation must serve some useful function to the current population members. However, an artificial evolutionary process is not bound by this constraint: instead, it can store components away for later access, give fitness valuations to things that could never be realised physically, and generally incorporate items that are designed to scaffold the evolutionary search.

3.2   Engagement with the Outside World

Most art takes place within a community—both an artistic community, and in the context of a broader culture. Attempts have been made within evolutionary art to engage with a world outside the immediate evolving population, but typically this has been a second artificial world running within the same overall computer system. Examples of this are the critics-based systems discussed earlier.

An important aspect of many works of art is their engagement with the outside world—artworks frequently comment on the world, either in a very direct way, but also via indirection, connotation and allusion. One way in which evolutionary artworks could engage more strongly with a human audience would be to include this kind of reference and allusion.

One of the arguments about evolutionary art made by McCormack (2005) is that such systems should be able to "produce art recognized by humans for its *artistic* contribution (as opposed to any purely technical fetish or fascination)". Whilst it is not a pre-requisite for such production that it must engage with the world outside the computer system, this would seem to be one way of working towards such an aim.

Something along these lines—albeit not within an evolutionary framework— has been proposed by Krzeczkowska et al. (2010). They refer to the above principle as the idea that "good art makes you think"—that is, makes you think about the work in the "context of the world around". This system extends Colton's *Painting Fool* system (`http://www.thepaintingfool.com/`) by creating collage works based on two kinds of web searches. The first extracts text from current news

articles; the second extracts images based on the keywords in these articles. These images are then assembled into visually-coherent collages using the heuristics in *Painting Fool.*

The idea of using web searches as a source of artistic material is the most obvious way in which computer art could engage directly with the real world. We can imagine a system where the fitness function is in some way based on a web search. A component within a piece of art will usually have a number of *connotations*—that is, things that are brought to the mind (consciously or subconsciously) of the viewer when they see it. One principle that we could adopt is that a piece of work that has a number of items that share connotations has a coherence that we might regard as being artistically valuable—the piece is *about something*, but without this "something" being represented directly. This could be represented via a fitness function that used web search to gain a list of such connotations, and assigned a high fitness to items (or sets of items) where there was a large amount of sharing of connotations.

One particular aspect of web search that has only become available in recent years is the ability to search using *images* rather than words. Therefore, such a search could begin from images, rather than evolution always working on words that are then translated into images.

One difficulty for computer art that attempts to engage with the outside world is judging when something has enough *significance* to make it worth making art about. Krzeczkowska et al. resolve this by beginning from current news stories—but, this is only one example. Another approach might be to find topics that are densely connected from different areas of the web: nexuses where many ideas meet and about which much can be said (but, this is an aesthetic judgement in its own right!). Another issue is how to judge the significance of links—there is a danger that we just end up making "bad puns", collections of things with a shallow connection. These are challenging areas of work.

References

Ando, D., Dahlsted, P., Nordahl, M. G., and Iba, H. (2007) 'Interactive GP with tree representation of classical music pieces', in Giacobini et al. (2007), pages 577–584.

Ariza, C. (2002) 'Prokaryotic groove: Rhythmic cycles as real-value encoded genetic algorithms', in *Proceedings of the 2002 International Computer Music Conference*, pages 561–567. International Computer Music Association.

Basa, T., Go, C., Yoo, K.-S., and Lee, W.-H. (2006) 'Using physiological signals to evolve art', in Rothlauf et al. (2006), pages 633–641.

Bergen, S. and Ross, B. J. (2012) 'Aesthetic 3D model evolution', in Machado et al. (2012b), pages 11–22.

Biles, J. and Eign, W. (1995) 'Genjam *Populi*: Training an IGA *via* audience-mediated performance', in *Proceedings of the 1995 International Computer Music Conference*. International Computer Music Association.

Biles, J. A. (2007) 'Improvizing with genetic algorithms: GenJam', in Miranda and Biles (2007), pages 137–169.

Bilotta, E., Pantano, P., Cupellini, E., and Rizzuti, C. (2007) 'Evolutionary methods for melodic sequences generation from non-linear dynamic systems', in Giacobini et al. (2007), pages 585–592.

Bird, J., Faith, J., and Webster, A. (2003) '*Tabula Rasa*: A case study in evolutionary curation', in Cagnoni et al. (2003), pages 981–995.

Bird, J., Husbands, P., Perris, M., Bigge, B., and Brown, P. (2008) 'Implicit fitness functions for evolving a drawing robot', in Giacobini et al. (2008), pages 473–478.

Blackwell, T. and Jefferies, J. (2005) 'Swarm tech-tiles', in Rothlauf et al. (2005), pages 468–477.

Blackwell, T. and Young, M. (2004) 'Swarm granulator', in Raidl et al. (2004), pages 399–408.

Bown, O. and Lexer, S. (2006) 'Continuous-time recurrent neural networks for generative and interactive musical performance', in Rothlauf et al. (2006), pages 652–663.

Bown, O. and McCormack, J. (2010) 'Taming nature: tapping the creative potential of ecosystem models in the arts', *Digital Creativity*, Vol. 21 No. 4, pp. 215–231.

Brown, A. R. (2004) 'An aesthetic comparison of rule-based and genetic algorithms for generating melodies', *Organised Sound*, Vol. 2 No. 4, pp. 191–197.

Burton, A. R. and Vladimirova, T. (1999) 'Generation of musical sequences with genetic techniques', *Computer Music Journal*, Vol. 23 No. 4, pp. 59–73.

Byrne, J., Fenton, M., Hemberg, E., McDermott, J., O'Neill, M., Shotton, E., and Nally, C. (2011) 'Combining structural analysis and multi-objective criteria for evolutionary architectural design', in Di Chio et al. (2011), pages 204–213.

Byrne, J., Hemberg, E., Brabazon, A., and O'Neill, M. (2012) 'A local search interface for interactive evolutionary architectural design', in Machado et al. (2012b), pages 23–34.

Cagnoni, S. et al., editors (2003) *Applications of Evolutionary Computing,*, volume 2611 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg.

Campolongo, G. and Vena, S. (2006) 'Science of networks and music: A new approach on musical analysis and creation', in Rothlauf et al. (2006), pages 642–651.

Collins, N. (2002) 'Experiments with a new customisable interactive evolution framework', *Organised Sound*, Vol. 7 pp. 267–273.

Collomosse, J. and Hall, P. (2005) 'Genetic paint: A search for salient paintings', in Rothlauf et al. (2005), pages 437–447.

Colton, S. (2008) 'Automatic invention of fitness functions with application to scene generation', in Giacobini et al. (2008), pages 381–391.

Colton, S., Cook, M., and Raad, A. (2011) 'Ludic considerations of tablet-based evo-art', in Di Chio et al. (2011), pages 223–233.

Colton, S., López de Mántaras, R., and Stock, O. (2009) 'Computational creativity: Coming of age', *AI Magazine*, Vol. 30 No. 3, pp. 11–14.

Daelemans, W. and van den Bosch, A. (2005) *Memory-Based Language Processing,* Cambridge University Press.

Dahlstedt, P. (2001) 'A MutaSynth in parameter space: Interactive composition through evolution', *Organised Sound*, Vol. 6 No. 2, pp. 121–124.

Dahlstedt, P. (2007a) 'Autonomous evolution of complete piano pieces and performances', in *Workshop on Music and Artificial Life.*

Dahlstedt, P. (2007b) 'Evolution in creative sound design', in Miranda and Biles (2007), pages 79–99.

Dahlstedt, P. (2009) 'On the role of temporary storage in interactive evolution', in Giacobini et al. (2009), pages 478–487.

Dahlstedt, P. and Nilsson, P. (2008) 'Free flight in parameter space: A dynamic mapping strategy for expressive free impro', in Giacobini et al. (2008), pages 479–484.

Dawkins, R. (1989) *The Selfish Gene,* Oxford University Press (second edition, original edition 1976).

Dawkins, R. (1997) *Climbing Mount Improbable,* Penguin.

De Prisco, R., Zaccagnino, G., and Zaccagnino, R. (2011) 'A genetic algorithm for dodecaphonic compositions', in Di Chio et al. (2011), pages 244–253.

De Smedt, T., Lechat, L., and Daelemans, W. (2011) 'Generative art inspired by nature, using NodeBox', in Di Chio et al. (2011), pages 264–272.

den Heijer, E. and Eiben, A. E. (2010) 'Comparing aesthetic measures for evolutionary art', in Di Chio et al. (2010), pages 311–320.

den Heijer, E. and Eiben, A. E. (2011) 'Evolving art using multiple aesthetic measures', in Di Chio et al. (2011), pages 234–243.

den Heijer, E. and Eiben, A. E. (2012a) 'Evolving pop art using scalable vector graphics', in Machado et al. (2012b), pages 48–59.

den Heijer, E. and Eiben, A. E. (2012b) 'Maintaining population diversity in evolutionary art', in Machado et al. (2012b), pages 60–71.

Di Chio, C. et al., editors (2010) *Applications of Evolutionary Computation,,* volume 6025 of *Lecture Notes in Computer Science.* Springer Berlin / Heidelberg.

Di Chio, C. et al., editors (2011) *Applications of Evolutionary Computation,*, volume 6625 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg.

Dorin, A. (2005) 'Artificial life, death and epidemics in evolutionary, generative electronic art', in Rothlauf et al. (2005), pages 448–457.

Dorin, A. (2009) 'Habitat: Engineering in a simulated audible ecosystem', in Giacobini et al. (2009), pages 488–497.

Draves, S. (2005) 'The electric sheep screen-saver: A case study in aesthetic evolution', in Rothlauf et al. (2005), pages 458–467.

Dubbin, G. A. and Stanley, K. O. (2010) 'Learning to dance through interactive evolution', in Di Chio et al. (2010), pages 331–340.

Eigenfeldt, A. (2009) 'The evolution of evolutionary software: Intelligent rhythm generation in kinetic engine', in Giacobini et al. (2009), pages 498–507.

Eigenfeldt, A. and Pasquier, P. (2011) 'A sonic eco-system of self-organising musical agents', in Di Chio et al. (2011), pages 283–292.

Eigenfeldt, A. and Pasquier, P. (2012) 'Populations of populations: Composing with multiple evolutionary algorithms', in Machado et al. (2012b), pages 72–83.

Eisenmann, J., Schroeder, B., Lewis, M., and Parent, R. (2011) 'Creating choreography with interactive evolutionary algorithms', in Di Chio et al. (2011), pages 293–302.

Ekárt, A., Sharma, D., and Chalakov, S. (2011) 'Modelling human preference in evolutionary art', in Di Chio et al. (2011), pages 303–312.

Eldridge, A. and Dorin, A. (2009) 'Filterscape: Energy recycling in a creative ecosystem', in Giacobini et al. (2009), pages 508–517.

Eldridge, A., Dorin, A., and McCormack, J. (2008) 'Manipulating artificial ecosystems', in Giacobini et al. (2008), pages 392–401.

Evans, B. (1987) 'Integration of music and graphics through algorithmic congruence', in *Proceedings of the 1987 International Computer Music Conference*, pages 17–24.

Filipe Antunes, R. and Fol Leymarie, F. (2012) 'Generative choreography: Animating in real-time dancing avatars', in Machado et al. (2012b), pages 1–10.

Fornari, J. (2011) 'Path of patches: Implementing an evolutionary soundscape art installation', in Di Chio et al. (2011), pages 323–333.

Fornari, J. (2012) 'A computational environment for the evolutionary sound synthesis of birdsongs', in Machado et al. (2012b), pages 96–107.

Fornari, J., Maia Jr., A., and Manzolli, J. (2007) 'Creating soundscapes using evolutionary spatial control', in Giacobini et al. (2007), pages 517–526.

Gartland-Jones, A. (2003) 'MusicBlox: A real-time algorithmic composition system incorporating a distributed interactive genetic algorithm', in Cagnoni et al. (2003), pages 145–155.

Giacobini, M. et al., editors (2007) *Applications of Evolutionary Computing,*, Lecture Notes in Computer Science 4448. Springer Berlin / Heidelberg.

Giacobini, M. et al., editors (2008) *Applications of Evolutionary Computing,*, volume 4974 of *Lecture Notes in Computer Science.* Springer Berlin / Heidelberg.

Giacobini, M. et al., editors (2009) *Applications of Evolutionary Computing,*, volume 5484 of *Lecture Notes in Computer Science.* Springer Berlin / Heidelberg.

Greenfield, G. (2004) 'Tilings of sequences of co-evolved images', in Raidl et al. (2004), pages 427–436.

Greenfield, G. (2005) 'Evolutionary methods for ant colony paintings', in Rothlauf et al. (2005), pages 478–487.

Greenfield, G. (2006) 'Robot paintings evolved using simulated robots', in Rothlauf et al. (2006), pages 611–621.

Greenfield, G. (2008) 'Evolved diffusion limited aggregation compositions', in Giacobini et al. (2008), pages 402–411.

Greenfield, G. (2009) 'Evolved ricochet compositions', in Giacobini et al. (2009), pages 518–527.

Greenfield, G. (2010) 'Generative art and evolutionary refinement', in Di Chio et al. (2010), pages 291–300.

Greenfield, G. (2012) 'A platform for evolving controllers for simulated drawing robots', in Machado et al. (2012b), pages 108–116.

Hart, D. A. (2007) 'Toward greater artistic control for interactive evolution of images and animation', in Giacobini et al. (2007), pages 527–536.

Hervás, R., Robinson, J., and Gervás, P. (2007) 'Evolutionary assistance in alliteration and allelic drivel', in Giacobini et al. (2007), pages 537–546.

Hiroaki Nishino, Takuya Sueyoshi, T. K. and Utsumiya, K. (2007) 'An interactive graphics rendering optimizer based on immune algorithm', in Giacobini et al. (2007), pages 459–469.

Hochreiter, R. (2006) 'Audible convergence for optimal base melody extension with statistical genre-specific interval distance evaluation', in Rothlauf et al. (2006), pages 712–716.

Hoover, A. K., Rosario, M. P., and Stanley, K. O. (2008) 'Scaffolding for interactively evolving novel drum tracks for existing songs', in Giacobini et al. (2008), pages 412–422.

Hornby, G. S. and Bongard, J. (2012) 'Learning comparative user models for accelerating human-computer collaborative search', in Machado et al. (2012b), pages 117–128.

Horner, A. and Goldberg, D. E. (1995) 'Genetic algorithms and computer-asssisted music composition', in Belew, R. and Booker, L., editors, *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 1–8. Morgan Kauffman.

Horowitz, D. (1994) 'Generating rhythms with genetic algorithms', in *AAAI-94 Proceedings*, page 1459.

Johanson, B. and Poli, R. (1998) 'GP-music: An interactive genetic programming system for music generation with automated fitness raters', in *Proceedings of the 3rd International Conference on Genetic Programming*, pages 181–186. MIT Press.

Johnson, C. G. (2012a) 'Connotation in computational creativity', *Cognitive Computation*, Vol. 4 No. 3, pp. 280–291.

Johnson, C. G. (2012b) 'The creative computer as romantic hero? computational creativity systems and creative personæ', in Maher, M. L. et al., editors, *Proceedings of the Third International Conference on Computational Creativity*, pages 57–61.

Johnson, C. G. (2012c) 'Fitness in evolutionary art and music: What has been used and what could be used?', in Machado et al. (2012b), pages 129–140.

Jones, D. (2008) 'Atomswarm: A framework for swarm improvisation', in Giacobini et al. (2008), pages 423–432.

Kaliakatsos-Papakostas, M. A., Epitropakis, M. G., Floros, A., and Vrahatis, M. N. (2012) 'Interactive evolution of 8-bit melodies with genetic programming towards finding aesthetic measures for sound', in Machado et al. (2012b).

Khalifa, Y. and Foster, R. (2006) 'A two-stage autonomous evolutionary music composer', in Rothlauf et al. (2006), pages 717–721.

Kötter, T. and Berthold, M. R. (2011) '(Missing) concept discovery in heterogeneous information networks', in *Proceedings of the Second International Conference on Computational Creativity*, pages 135–140.

Krzeczkowska, A., El-Hage, J., Colton, S., and Clark, S. (2010) 'Automated collage generation—with intent', in Ventura, D. et al., editors, *Proceedings of the International Conference on Computational Creativity*, pages 36–40.

Law, E. H. H. and Phon-Amnuaisuk, S. (2008) 'Towards music fitness evaluation with the hierarchical SOM', in Giacobini et al. (2008), pages 443–452.

Lewis, G. (2000) 'Too many notes: Computers, complexity and culture in *Voyager*', *Leonardo Music Journal*, Vol. 10 pp. 33–39.

Lewis, M. (2004) 'Aesthetic video filter evolution in an interactive real-time framework', in Raidl et al. (2004), pages 409–418.

Li, Y., Hu, C., Chen, M., and Hu, J. (2012) 'Investigating aesthetic features to model human preference in evolutionary art', in Machado et al. (2012b), pages 153–164.

Li, Y. and Hu, C.-J. (2010) 'Aesthetic learning in an interactive evolutionary art system', in Di Chio et al. (2010), pages 301–310.

Lopes, P. and Urbano, P. (2012) 'The traveling percussionist', in Machado et al. (2012b), pages 165–175.

Lutton, E., Cayla, E., and Chapuis, J. (2003) 'ArtiE-Fract: The artist's viewpoint', in Cagnoni et al. (2003), pages 125–130.

Machado, P. and Cardoso, A. (2002) 'All the truth about NEvAr', *Applied Intelligence*, Vol. 16 No. 2, pp. 101–118.

Machado, P., Correia, J., and Romero, J. (2012a) 'Expression-based evolution of faces', in Machado et al. (2012b), pages 187–198.

Machado, P. and Graca, F. (2008) 'Evolutionary pointillist modules: Evolving assemblages of 3D objects', in Giacobini et al. (2008), pages 453–462.

Machado, P., Nunes, H., and Romero, J. (2010) 'Graph-based evolution of visual languages', in Di Chio et al. (2010), pages 271–280.

Machado, P., Romero, J., and Carballal, A., editors (2012b) *Evolutionary and Biologically Inspired Music, Sound, Art and Design,*, volume 7247 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg.

Machado, P., Romero, J., Santos, M., Cardoso, A., and Manaris, B. (2004) 'Adaptive critics for evolutionary artists', in Raidl et al. (2004), pages 437–446.

Machewe, A. T. and Parmee, I. C. (2007) 'Towards an interactive, generative design system: Integrating a 'build and evolve' approach with machine learning for complex freeform design', in Giacobini et al. (2007), pages 449–458.

Magnus, C. (2006) 'Evolutionary musique concrète', in Rothlauf et al. (2006), pages 688–695.

Manaris, B., Machado, P., McCauley, C., Romero, J., and Krehbiel, D. (2005) 'Developing fitness functions for pleasant music: Zipf's law and interactive evolution systems', in Rothlauf et al. (2005), pages 498–507.

Manaris, B., Vaughan, D., Wagner, C., Romero, J., and Davis, R. (2003) 'Evolutionary music and the Zipf-Mandelbrot law: Developing fitness functions for pleasant music', in Cagnoni et al. (2003), pages 65–72.

Mandelis, J. (2003) 'Genophone: Evolving sounds and integral performance parameter mappings', in Cagnoni et al. (2003), pages 253–270.

Manzolli, J., Moroni, A., von Zuben, F., and Gudwin, R. (1999) 'An evolutionary approach to algorithmic composition', *Organised Sound*, Vol. 4 No. 2, pp. 121–126.

Marcelo Caetano, J. M. and Zuben, F. V. (2007) 'Self-organizing bio-inspired sound transformation', in Giacobini et al. (2007), pages 477–487.

McCormack, J. (2004) 'Aesthetic evolution of L-systems revisited', in Raidl et al. (2004), pages 477–488.

McCormack, J. (2005) 'Open problems in evolutionary music and art', in Rothlauf et al. (2005), pages 428–436.

McCormack, J. and Bown, O. (2009) 'Life's what you make: Niche construction and evolutionary art', in Giacobini et al. (2009), pages 528–537.

McDermott, J. (2012) 'Graph grammars as a representation for interactive evolutionary 3D design', in Machado et al. (2012b), pages 199–210.

Miranda, E. R. and Biles, J. A., editors (2007) *Evolutionary Computer Music,*. Springer.

Moniz, R. D. and Jacob, C. (2009) '*Fractal Evolver*: Interactive evolutionary design of fractals with grid computing', in Giacobini et al. (2009), pages 442–451.

Moroni, A. and Manzolli, J. (2010) 'From evolutionary composition to robotic sonification', in Di Chio et al. (2010), pages 401–410.

Murray Browne, T. and Fox, C. (2009) 'Global expectation-violation as fitness function in evolutionary composition', in Giacobini et al. (2009), pages 538–546.

Nagel, U., Thiel, K., Kötter, T., Piatek, D., and Berthold, M. R. (2011) 'Bisociative discovery of interesting relations between domains', in *Advances in Intelligent Data Analysis X: 10th International Symposium*, Lecture Notes in Computer Science Volume 7014, pages 306–317.

Nemirovsky, P. and Luger-Guillaume, R. (2004) 'Improvisational media space: Architecture and strategies for evolution', in Raidl et al. (2004), pages 457–466.

Nemirovsky, P. and Watson, R. (2003) 'Genetic improvisation model: A framework for real-time performance environments', in Cagnoni et al. (2003), pages 547–558.

Nicolau, M. and Costelloe, D. (2011) 'Using grammatical evolution to parameterise interactive 3D image generation', in Di Chio et al. (2011), pages 374–383.

Oh, J. C. and Zajec, E. (2010) 'Evolving artistic styles through visual dialogues', in Di Chio et al. (2010), pages 261–270.

Ortega de la Puente, A., Sánchez Alfonso, R., and Alfonseca Moreno, M. (2002) 'Automatic composition of music by means of grammatical evolution', in *Proceedings of the 2002 Conference on APL*, pages 148–155. ACM Press.

Oscar Sanjuán, Gloria García, Y. S. and Luque, C. (2007) 'Human mosaic creation through agents and interactive genetic algorithms applied to videogames movements', in Giacobini et al. (2007), pages 470–476.

Phon-Amnuaisuk, S. (2009) 'Composing using heterogeneous cellular automata', in Giacobini et al. (2009), pages 547–556.

24　　*C.G. Johnson*

Phon-Amnuaisuk, S., Law, E., and Kuan, H. (2007) 'Evolving music generation with som-fitness genetic programming', in Giacobini et al. (2007), pages 557–566.

Raidl, G. et al., editors (2004) *Applications of Evolutionary Computing,*, volume 3005 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg.

Ralley, D. (1995) 'Genetic algorithms as a tool for melodic development', in *Proceedings of the 1995 International Computer Music Conference*. International Computer Music Association.

Reddin, J., McDermott, J., and O'Neill, M. (2009) 'Elevated pitch: Automated grammatical evolution of short compositions', in Giacobini et al. (2009), pages 579–584.

Reynolds, C. (2011) 'Evolving textures from high level descriptions: Gray with an accent color', in Di Chio et al. (2011), pages 384–393.

Rizzuti, C., Bilotta, E., and Pantano, P. (2009) 'A GA-based control strategy to create music with a chaotic system', in Giacobini et al. (2009), pages 585–590.

Romero, J. and Machado, P., editors (2008) *The Art of Artificial Evolution,*. Springer.

Romero, J., Machado, P., Carballal, A., and Osorio, O. (2011) 'Aesthetic classification and sorting based on image compression', in Di Chio et al. (2011), page RMCO11.

Romero, J., Machado, P., and Santos, A. (2009) 'On the socialization of evolutionary art', in Giacobini et al. (2009), pages 557–566.

Romero, J., Machado, P., Santos, A., and Cardoso, A. (2003) 'On the development of critics in evolutionary computation artists', in Cagnoni et al. (2003), pages 559–569.

Rothlauf, F. et al., editors (2005) *Applications of Evolutionary Computing,*, volume 3449 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg.

Rothlauf, F. et al., editors (2006) *Applications of Evolutionary Computing,*, volume 3907 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg.

Santarosa, R., Moroni, A., and Manzolli, J. (2006) 'Layered genetical algorithms evolving into musical accompaniment generation', in Rothlauf et al. (2006), pages 722–726.

Shao, J., McDermott, J., O'Neill, M., and Brabazon, A. (2010) 'Jive: A generative, interactive, virtual, evolutionary music system', in Di Chio et al. (2010), pages 341–350.

Somnuk Phon-Amnuaisuk, E. H. H. L. and Kuan, H. C. (2007) 'Evolving music generation with SOM-fitness genetic programming', in Giacobini et al. (2007), pages 557–566.

Suzuki, R., Yamaguchi, S., Cody, M. L., Taylor, C. E., and Arita, T. (2011) 'iSoundScape: Adaptive walk on a fitness soundscape', in Di Chio et al. (2011), pages 404–413.

Svangård, N. and Nordin, P. (2004) 'Automated aesthetic selection of evolutionary art by distance based classification of genomes and phenomes using the universal similarity metric', in Raidl et al. (2004), pages 447–456.

Thywissen, K. (1999) 'GeNotator: an environment for exploring the application of evolutionary techniques in computer assisted composition', *Organised Sound*, Vol. 4 No. 2, pp. 127–133.

Todd, S. and Latham, W. (1992) *Evolutionary Art and Computers,* Academic Press, London.

Tsuneo Kagawa, Yukihide Tamotsu, H. N. and Utsumiya, K. (2007) 'Interactive texture design using IEC framework', in Giacobini et al. (2007), pages 439–448.

Tufte, G. and Gangvik, E. (2008) 'Trans<−>former #13: Exploration and adaptation of evolution expressed in a dynamic sculpture', in Giacobini et al. (2008), pages 509–514.

Unold, O. and Troc, M. (2007) 'Curve, draft, and style: Three steps to the image', in Giacobini et al. (2007), pages 601–608.

Uozumi, Y. (2007) 'GISMO2: An application for agent-based composition', in Giacobini et al. (2007), pages 609–616.

Urbano, P. (2005) 'Playing in the pheromone playground: Experiences in swarm painting', in Rothlauf et al. (2005), pages 527–532.

Urbano, P. (2006) 'Consensual paintings', in Rothlauf et al. (2006), pages 622–632.

von Mammen, S., Wong, J., and Jacob, C. (2008) 'Virtual constructive swarm compositions and inspirations', in Giacobini et al. (2008), pages 491–496.

Vouliouri, E. (2011) 'Merging aesthetics with functionality: An interactive genetic algorithm based on the principle of weighted mutation', in Di Chio et al. (2011), pages 424–433.

Wakefield, G. and Ji, H. H. (2009) 'Artificial nature: Immersive world making', in Giacobini et al. (2009), pages 597–602.

Waschka II, R. (2007) 'Composing with genetic algorithms: *GenDash*', in Miranda and Biles (2007), pages 117–136.

Webb Young, J. (2003 (original edition 1943)) *A Technique for Producing Ideas,* McGraw-Hill.

Wolkowicz, J., Heywood, M., and Keselj, V. (2009) 'Evolving indirectly represented melodies with corpus-based fitness evaluation', in Giacobini et al. (2009), pages 603–608.

Yee-King, M. J. (2007) 'An automated music improviser using a genetic algorithm driven synthesis engine', in Giacobini et al. (2007), pages 567–576.

Zhang, Q. and Miranda, E. R. (2007)   'Experiments in generatic musical performance with a genetic algorithm', in Miranda and Biles (2007), pages 100–116.