Fitness in Evolutionary Art and Music: What Has Been Used and What Could be Used?

Colin G. Johnson

University of Kent, Canterbury, Kent
C.G.Johnson@kent.ac.uk,
http://www.cs.kent.ac.uk/people/staff/cgj/

Abstract. This paper considers the notion of *fitness* in evolutionary art and music. 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. Papers from a large collection are classified using this taxonomy. The paper then discusses a number of ideas that have not be used for fitness evaluation in evolutionary art and which might be valuable in future developments: memory, scaffolding, connotation and web search.

1 Introduction

A large number of artistic and musical projects have used evolutionary algorithms as a way of generating their source material. One important component of evolutionary algorithms is a way of evaluating the fitness of individuals in the population. The aim of this paper is twofold. Firstly, we carry out a substantial survey, based on the nine previous EvoMusArt proceedings, of the different ways in which fitness has been used in creative music and art applications of evolutionary algorithms. This informs the creation of a taxonomy of how fitness has been used, and we discuss this taxonomy. Secondly, we point out gaps in this taxonomy—aspects of creativity that are not captured by these applications of fitness—and discuss whether these other aspects could be brought into a fitness-based framework.

McCormack [30] 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.

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" [8]. As such, we will not devote any space to many other interesting areas of evolutionary computation in art and music.

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 study were the papers published between 2003 and 2011 in the EvoMusArt workshops and symposia [6, 35, 39, 40, 18–20, 14, 15], which have been held every year and are the main event focused on evolutionary methods in music and art. The papers were reviewed, and those that used evolutionary methods to produce a creative outcome—i.e. something generated as a result of the process or end result of the evolution, that is not the result just of applying well-understood rules or constraints—were selected. Therefore, we exclude papers that were e.g. 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, papers that used evolutionary algorithms within a framework of fixed rules (e.g. musical harmony rules) and papers that were about performance, rendering or re-presentation of existing works. We also excluded papers that presented systems where no notion of fitness was involved. Most of the systems were evolutionary algorithms but some other 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. Due to space limitations a full listing of the papers and how they fit into the taxonomy is not possible, but these details can be obtained by contacting the author.

2.1 Fitness Scope

The first part of the taxonomy is a classification of what fitness is applied to: we will call this the *fitness scope*. This is one of three classes. The first, which we will call a *set of works* is where each member of the population consists of a collection of individual artworks, and 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 which is viewed in some fashion 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, so the total does not necessarily add up to the total number of papers.

The *whole set of works* category is represented by just two papers. The first paper [3] 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 [26] is a meta-level evolution of generative systems. Each

Scope	Set of Works	Whole Work	Process as the Work
Number of papers	2	64	23

Table 1. The scope of fitness application. This is the count of the total number of papers that used each of the three scope types.

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* [12]. Typically, the evolutionary process is carried out and then the fittest example in the final generation is output.

Usually, the aim of such processes has been to actively search out "good" works of art. However, Reddin et al. [36] and Dahlstedt [10] have argued that the role of the fitness function could be just to eliminate bad examples, leaving the "expressive power of the representation" [36] 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 future version of the taxonomy might separate this out as a distinctive category.

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 artwork*. 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 genre display just the end result of the process: this is illustrated, for example, in a piece by McCormack and Brown [31], agents deposit lines on a surface during the evolutionary process, and the final disposition of the lines generated throughout evolution forms the work.

2.2 Fitness Basis

The second part of the taxonomy considers how the fitness is evaluated: call this *fitness basis*. This is divided into five classes: aesthetic measure (a fixed function measuring the quality of the solution), human interaction, the use of a corpus of material or guiding example, 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 2. Again, some papers fell into more than one category. The remainder of this section discusses these five categories in detail.

Basis	Aesthetic Measure	Human Interaction	Corpus or Example	Endogenous	Critics
Number of papers	29	27	20	18	7

Table 2. The basis of fitness application. This is the count of the total number of papers that used each of the basis types.

Aesthetic measure. The first basis for fitness is some notion of *aesthetic measure*. That is to say, some function is applied to the member of the population that measures the quality of that population 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. [2], melodies receive a score based on a sum of positive values for desirable features of a melody (e.g. well-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 [3], 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 [21], 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. [22] 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.

However, one aspect of creative development that such 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 a human. 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 & Watson [33] and Dahlstedt & Nilsson [11]) 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.

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. [1], 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.

Corpus or Example. 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 [9] 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. [34] 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. [29], 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 de-

signer is asserting that the functional form of the fitness function is a meaningful measure, and then the details are set by the corpus.

It could be argued that this approach is rather limited. Whilst 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 [16] who discusses the idea of whether two generative art algorithms based on the same underlying algorithm could provide the same aesthetic impact; however, he is concerned with the whole generative process, whereas we are focused on fitness.

A second approach to using examples is to use an example 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* [41], an articulated sculpture is given a number of target configurations, and applies a genetic algorithm to learn how to get into these configurations. In the *MusicBlox* system [17], 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" [17]—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 typically found in conjunction with the presentation of the whole process as the artwork: 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. [4], where a drawing robot is provided with a fitness function that "correlated the changes in state of their line detector and pen position" [4]. 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 [5]).

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.

It is interesting that most 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* [41] discussed above is perhaps the closest to this; within the musical domain, the target-directed evolution of electroacoustic music by Magnus [28] is similar.

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. [38] 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 [37]—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 [27] gives some implemented examples of the earlier stages—but, these are essentially classification algorithms.

An attempt to develop a meta-level automated analyst of automatically generated images has been carried out by Colton [7]. 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, an 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 [42], 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 [32, 24], and the idea of a fitness measure based on what a population member *connotes* as well as what it *denotes* has been explored briefly elsewhere [23]. 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 [13], as a steady climb up a smooth gradient, by contrast with the saltational leap up of 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 [30] 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. [25]. 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

1. Basa, T., Go, C., Yoo, K.S., Lee, W.H.: Using physiological signals to evolve art. In: Rothlauf et al. [40], pp. 633–641
2. Bilotta, E., Pantano, P., Cupellini, E., Rizzuti, C.: Evolutionary methods for melodic sequences generation from non-linear dynamic systems. In: Giacobini et al. [18], pp. 585–592
3. Bird, J., Faith, J., Webster, A.: *Tabula Rasa*: A case study in evolutionary curation. In: Cagnoni et al. [6], pp. 981–995
4. Bird, J., Husbands, P., Perris, M., Bigge, B., Brown, P.: Implicit fitness functions for evolving a drawing robot. In: Giacobini et al. [19], pp. 473–478
5. Bown, O., McCormack, J.: Taming nature: tapping the creative potential of ecosystem models in the arts. Digital Creativity 21(4), 215–231 (2010)
6. Cagnoni, S., et al. (eds.): Applications of Evolutionary Computing, Lecture Notes in Computer Science, vol. 2611. Springer Berlin / Heidelberg (2003)
7. Colton, S.: Automatic invention of fitness functions with application to scene generation. In: Giacobini et al. [19], pp. 381–391
8. Colton, S., López de Mántaras, R., Stock, O.: Computational creativity: Coming of age. AI Magazine 30(3), 11–14 (2009)
9. Daelemans, W., van den Bosch, A.: Memory-Based Language Processing. Cambridge University Press (2005)
10. Dahlstedt, P.: Autonomous evolution of complete piano pieces and performances. In: Workshop on Music and Artificial Life (2007)
11. Dahlstedt, P., Nilsson, P.: Free flight in parameter space: A dynamic mapping strategy for expressive free impro. In: Giacobini et al. [19], pp. 479–484
12. Dawkins, R.: The Selfish Gene. Oxford University Press (1989), second edition. Original edition 1976
13. Dawkins, R.: Climbing Mount Improbable. Penguin (1997)
14. Di Chio, C., et al. (eds.): Applications of Evolutionary Computation, Lecture Notes in Computer Science, vol. 6025. Springer Berlin / Heidelberg (2010)
15. Di Chio, C., et al. (eds.): Applications of Evolutionary Computation, Lecture Notes in Computer Science, vol. 6625. Springer Berlin / Heidelberg (2011)
16. Evans, B.: Integration of music and graphics through algorithmic congruence. In: Proceedings of the 1987 International Computer Music Conference. pp. 17–24 (1987)
17. Gartland-Jones, A.: Musicblox: A real-time algorithmic composition system incorporating a distributed interactive genetic algorithm. In: Cagnoni et al. [6], pp. 145–155
18. Giacobini, M., et al. (eds.): Applications of Evolutionary Computing. Lecture Notes in Computer Science 4448, Springer Berlin / Heidelberg (2007)
19. Giacobini, M., et al. (eds.): Applications of Evolutionary Computing, Lecture Notes in Computer Science, vol. 4974. Springer Berlin / Heidelberg (2008)
20. Giacobini, M., et al. (eds.): Applications of Evolutionary Computing, Lecture Notes in Computer Science, vol. 5484. Springer Berlin / Heidelberg (2009)
21. Greenfield, G.: Evolved ricochet compositions. In: Giacobini et al. [20], pp. 518–527
22. Hervàs, R., Robinson, J., Gervàs, P.: Evolutionary assistance in alliteration and allelic drivel. In: Giacobini et al. [18], pp. 537–546
23. Johnson, C.G.: Search and notions of creativity. In: Veale, T., Pease, A., Wiggins, G. (eds.) Proceedings of the IJCAI-2005 Workshop on Computational Creativity (2005)

24. Kötter, T., Berthold, M.R.: (missing) concept discovery in heterogeneous information networks. In: Proceedings of the Second International Conference on Computational Creativity. pp. 135–140 (2011), http://www.inf.uni-konstanz.de/bioml2/publications/Papers2011/KoBe11.pdf
25. Krzeczkowska, A., El-Hage, J., Colton, S., Clark, S.: Automated collage generation—with intent. In: Ventura, D., et al. (eds.) Proceedings of the International Conference on Computational Creativity. pp. 36–40 (2010), available at http://creative-systems.dei.uc.pt/icccx
26. Machado, P., Nunes, H., Romero, J.: Graph-based evolution of visual languages. In: Di Chio et al. [14], pp. 271–280
27. Machado, P., Romero, J., Santos, M., Cardoso, A., Manaris, B.: Adaptive critics for evolutionary artists. In: Raidl et al. [35], pp. 437–446
28. Magnus, C.: Evolutionary musique concrète. In: Rothlauf et al. [40], pp. 688–695
29. Manaris, B., Vaughan, D., Wagner, C., Romero, J., Davis, R.: Evolutionary music and the Zipf-Mandelbrot law: Developing fitness functions for pleasant music. In: Cagnoni et al. [6], pp. 65–72
30. McCormack, J.: Open problems in evolutionary music and art. In: Rothlauf et al. [39], pp. 428–436
31. McCormack, J., Bown, O.: Life's what you make: Niche construction and evolutionary art. In: Giacobini et al. [20], pp. 528–537
32. Nagel, U., Thiel, K., Kötter, T., Piatek, D., Berthold, M.R.: Bisociative discovery of interesting relations between domains. In: Advances in Intelligent Data Analysis X: 10th International Symposium. pp. 306–317. Lecture Notes in Computer Science Volume 7014 (2011), http://www.inf.uni-konstanz.de/bioml2/publications/Papers2011/NTKP+11.pdf
33. Nemirovsky, P., Watson, R.: Genetic improvisation model a framework for real-time performance environments. In: Cagnoni et al. [6], pp. 547–558
34. Phon-Amnuaisuk, S., Law, E., Kuan, H.: Evolving music generation with som-fitness genetic programming. In: Giacobini et al. [18], pp. 557–566
35. Raidl, G., et al. (eds.): Applications of Evolutionary Computing, Lecture Notes in Computer Science, vol. 3005. Springer Berlin / Heidelberg (2004)
36. Reddin, J., McDermott, J., ONeill, M.: Elevated pitch: Automated grammatical evolution of short compositions. In: Giacobini et al. [20], pp. 579–584
37. Romero, J., Machado, P., Santos, A.: On the socialization of evolutionary art. In: Giacobini et al. [20], pp. 557–566
38. Romero, J., Machado, P., Santos, A., Cardoso, A.: On the development of critics in evolutionary computation artists. In: Cagnoni et al. [6], pp. 559–569
39. Rothlauf, F., et al. (eds.): Applications of Evolutionary Computing, Lecture Notes in Computer Science, vol. 3449. Springer Berlin / Heidelberg (2005)
40. Rothlauf, F., et al. (eds.): Applications of Evolutionary Computing, Lecture Notes in Computer Science, vol. 3907. Springer Berlin / Heidelberg (2006)
41. Tufte, G., Gangvik, E.: Trans<–>former #13: Exploration and adaptation of evolution expressed in a dynamic sculpture. In: Giacobini et al. [19], pp. 509–514
42. Webb Young, J.: A Technique for Producing Ideas. McGraw-Hill (2003 (original edition 1943))