Chapter 1
Aesthetics, Artificial Intelligence, and Search-based Art

Colin G. Johnson

Abstract Why do people exhibit particular behaviour towards a class of objects called artworks? That is the topic of study in aesthetics. This paper explores how various theories of aesthetics can be interpreted in the context of artworks generated by artificial intelligence systems, in particular those that are grounded in the idea of search as a means of implementing intelligence computationally. A number of aesthetic theories are explored, including ideas of imitation, skill, form, expression, imagination, and focus. The paper concludes by highlighting a number of areas that, in light of these considerations, have been neglected by the makers of computer art systems and which provide future opportunities in this area.

1.1 Introduction

Aesthetics is the attempt to understand why and how people act in a particular set of ways to a certain class of objects in the world called artworks. Why do people exhibit particular forms of behaviour towards these objects, so called *aesthetic behaviour*? How do they form *aesthetic judgements*? Why do humans exhibit the behaviour of 'attending to objects for their own sake' (Sheppard, 1987, p72)? These questions form the core of the subject of aesthetics.

Recent decades have seen a large amount of work in the areas of computer-based art, artificial intelligence (AI) in the arts, and computational creativity, all applied to a wide variety of artistic domains. However, little attempt has been made to connect this with ideas and theories from aesthetics and the philosophy of art. The aim of this paper is to make an initial foray into this area, with a particular focus on the search-based approach to AI. The core of the paper consists of a consideration of various theories that have been proposed to explain aesthetics, and a consideration

School of Computing, University of Kent, Canterbury, UK
Institute for Advanced Studies in the Humanities, University of Edinburgh, Edinburgh, UK

of how those theories link with ideas about how creative AI systems explore search spaces.

This paper begins with an overview of the ideas of AI in artistic creative domains, focusing on those approaches, such as evolutionary algorithms, that take a search-based approach to artistic creation. Computational and AI approaches have been applied to a wide range of artistic domains, and the principles discussed in this paper are similarly designed to be across the arts, though most of the examples in the paper are drawn from the visual arts. Typically, aesthetic theories are also neutral with regard to artforms. Within the literature on aesthetics and the philosophy of art, a broad set of ideas concerned with emotion, imitation, form, expression, etc. have been explored, and then been applied to specific artforms and works.

Following this introduction, a number of topics in aesthetics are examined, and their relationship to AI systems for making art are explored. To a lesser extent the paper also explores AI systems that analyse and critique art. The aim of this exploration is to ask whether traditional aesthetic theories can be applied to such means of art-making, whether these theories need to be expanded or adapted to examine these forms of art, whether these aesthetic theories highlight gaps or hidden assumptions in our understanding of AI-based art, and whether computational and AI concepts might help us to understand human art making.

1.2 Search, AI, and Art

There is an extensive literature on the links between art and artificial intelligence. One strand is embedded in the literature on computational creativity (McCormack and d'Inverno, 2012; Veale and Cardoso, 2019; Romero and Machado, 2008), and considers both the engineering of systems that act in creative ways, or which work in a co-creative way with people, as well as the evaluation of such systems (Jordanous, 2012). Other work explores specific systems, either written by the creators of those systems (Colton, 2012; Machado and Cardoso, 2002; Draves, 2005; Cohen, 1995) or by external commentators (McCorduck, 1991).

Another strand explores theoretical and philosophical issues. These include attempts to unpack human creative processes and compare them to proposed creative AI systems (Boden, 1990), questions around attribution of creative responsibility and authorship (Broeckmann, 2019; Johnson, 2014; Mazzone and Elgammal, 2019), and works that compare the use of AI with previous technological changes in the artworld such as the introduction of photography (Hertzmann, 2018; Agüera y Arcas, 2017). One issue is that publications on AI art focus on the technical achievements and methods, and much less on developing '*art* theories of evolutionary and generative art' (McCormack, 2005). This paper begins to address that by making links between AI art and the wider literature on aesthetics.

A key position in much artificial intelligence research and practice is that intelligent action and thought can be achieved by a search process operating in an appropriate search space. For example, much machine learning can be seen as a search in a space of hypotheses for a hypothesis that makes successful predictions based on a set of training data (Mitchell, 1997). As another example, reinforcement learning (Sutton and Barto, 2018) reframes the problem of learning of act intelligently as the search in a space of possible action policies for a policy that links appropriate actions to each situation.

A similar approach can be taken to creative, artistic domains. The key idea is that intelligent creative behaviour can be carried out by some kind of search process. Music creation can be seen as a search is a space of possible configurations of notes over a period of time. Literature is can be reframed as the task of deciding which point the space of all permutations of words up to a certain length. In practice, such search domains are too large to be searched directly, and so some additional structures are included which give the search some additional domain knowledge. Choosing this search space therefore becomes one of the main design activities for the human creator of an AI art system.

Boden (1990) has discussed these ideas at length, considering different ways in which AI search algorithms relate to the *conceptual spaces* in which creative AI systems work. In particular, she distinguishes between *exploratory* creativity, in which the conceptual space is fixed, and *transformational* creativity, in which the creative agent expands or transforms the space being explored. The latter are associated with the transformative moments in the history of an artform where assumptions are dropped, new media explored, or an expanded conceptual vocabulary becomes available.

An example of this from art history is the transition from perspective painting, which is already a vast space of possibilities, into into cubist painting, which allows a whole new space to be explored, where multiple points of perspective are combined into a single work. As discussed by Wiggins (2006a,b), these conceptual spaces might not be identical to traditional AI search spaces; they are more abstract spaces of ideas rather than enumerations of possible works. Also, it is important to get the level of granularity right when talking about such spaces; little is to be gained from considering literature as the exploration of all possible combinations of words up to a certain length; there is a whole ontology of concepts that help to subdivide this search space.

One difficulty with a traditional AI search approach to creative domains is that search algorithms typically assume that a particular point in the search space will be assigned the same objective function value regardless of when it is evaluated. For creative domains, it is not obvious that this assumption holds. Danto (1981) has argued that most statements about aesthetic judgement also involve the context of the artwork being examined; both the personal knowledge of the person exhibiting the aesthetic behaviour, but more importantly a broader ‘artworld’ which influences individual’s aesthetic judgements. A work can validly receive a particular judgement at one point in the history of art, but a different one at a different time. This points towards interesting areas of creative AI in a broader simulated world containing both AI artists and AI critics (Machado et al., 2004; Romero et al., 2003).

1.2.1 What Drives Search in Creative Domains?

In order to apply a search-based AI technique, it is necessary to have some means of driving the search, which is typically a way of assigning a value or ranking to a specific point in the space. Such methods are known variously as objective functions, fitness functions, error measures, or loss functions. Going beyond the notion of assigning a simple score or rank, recent research has introduced the more general notions of *search drivers* (Krawiec et al., 2016), where the assessment of a point in the search space provides a richer judgement than a simple score or ranking.

The vast majority of applications tackled using search-based AI make use of an objective, clearly-defined function to act as this measure. For subjective application areas in the arts, a wider variety of fitness drivers has been used. In two earlier papers (Johnson, 2012b, 2016), I outlined a taxonomy of fitness measures used in search-based (primarily evolutionary) art systems, based on a survey of many papers describing such systems. A major dimension of that taxonomy was the kind of fitness used—the *fitness basis*. This dimension consists of six classes:

AESTHETIC MEASURE. These are where the search is driven by a formula that gives a score or ranking to an object that reflects some aspect of its aesthetics. Such measures range from very generic measures that can be applied across different artforms, such as ways of measuring symmetry or balance, to measures that capture some aspect of a particular artform or style. Such measures can be critiqued from the standpoint of creativity, because as soon as an idea of the aesthetically valuable is fixed, this provides a limitation on artistic creativity. Furthermore, once an aesthetic measure is fixed, it provides a challenge to future artists to find ways of making art that either are not measured well by, or are incapable of being measured by, those aesthetic criteria. Nonetheless, such measures retain value, in particular for the creation of novel works in a known style.

HUMAN INTERACTION. In these systems, the search is driven by a human collaborator or audience member examining candidate points in the space and assigning them some kind of score or rank, which then drives the search.

CORPUS. This is where the fitness is guided by a collection of material, the corpus. Most typically, this will be in the form of an 'inspiring set' (Ritchie, 2007), where the search is guided by similarity to that set, or similarity to patterns abstracted from that set. Alternatively, that corpus might be used as a collection of material for the search to plunder and reassemble.

SEEDS AND TARGETS. In these systems, the movement through the search space is driven by one or more externally provided examples. However, in contrast to the 'inspiring set', approach, these are used either as something from the search to work away from or towards. This includes systems that are given some initial state (a 'seed') and then the system elaborates on this or moves away from it; alternatively, a 'target' where the system starts from an arbitrary starting state and moves towards that state. Typically, it is not the final state that is the artistic focus of interest: it is the trajectory of the system as it moves through the search space that forms the artwork.

ENDOGENOUS. These systems use the idea of a fitness measure or driver, but that measure does not attempt to assign aesthetic value to an object. Instead, the measure drives some kind of evolutionary or learning process, and it is the result of that process that is observed as the object of aesthetic attention.

CRITICS AND CO-EVOLUTION. The creation of artworks doesn't happen in isolation. Audiences and critics respond to works, and artists respond to this network of feedback. Broader social concerns both inform, and are sometimes influenced by, such work. This category of systems consists of those systems that simulate or make use of such ideas of art-making being contextualised in a wider society. This class includes attempts to create agents that observe and critique work, the feedback cycles that are therefore generated, and systems that have some notion of co-evolution between different populations of agents.

Potentially, all of these kinds of search drivers can be applied at different levels. Most typically, they are applied to individual candidate artworks in a search space. However, they can also be applied to components of these artworks, and in (Johnson, 2016) it was noted that a particularly large number of the systems using ENDOGENOUS fitness measures applied fitness to components of a work. Indeed, they have occasionally been applied to whole collections of works, for example in Bird et al. (2003), where an evolutionary algorithm is applied to curating an exhibit, and each point in the search space represents a structured collection of works.

These AI agents are concerned with the creation of artworks. There is also the, contrastingly, a category of AI agents that are designed that exhibit aesthetic behaviour, appreciate and critique artworks, and come to (perhaps simulated) aesthetic judgements. This is one role that agents can play as part of a CRITICS AND CO-EVOLUTION fitness basis.

1.3 Aesthetic Theories

Plato's dialogues in the *Greater Hippias* (Plato, 1926, 287b–292e) represent the earliest attempt to explore the idea that there is some general aesthetic theory—in this case, a theory of the 'beautiful'—abstracted from specifics. Having been asked by Socrates to explain 'what the beautiful is', Hippias responds with various examples: 'a beautiful maiden is beautiful', 'beautiful mares are bred in our country', and so on for pots, lyres, etc. However, for each specific example given, Socrates responds with a request for a generality, to unpack the idea that 'there is something by reason of which these things would be beautiful'.

Aesthetic theories attempt to unpack human aesthetic behaviour and aesthetic judgement, and to analyse the kinds of things that are the objects of such attention. One problem with such theories is that any attempt to pin down what makes something a worthy object of aesthetic attention eventually results in the space of objects thus defined being gradually exhausted; as a result, creative individuals seek means of aesthetic expression that either contradict, or cannot be readily analysed, by current aesthetic theories.

As a result, the broad history of aesthetic theory is one of increasing abstraction. Initially—as pointed out by Plato—aesthetic value was something ascribed to individual objects, each in their own way. Gradually, domain specific theories of aesthetic value emerged. For example, in music, theories of harmony were developed which attempted to identify features of harmony that were supposed to be pleasing to the ear. Over time, these rules were broken and expanded. In the last couple of centuries, aesthetic theories have become more abstracted still, focusing on broad features of artworks across different domains. Furthermore, this has been influenced by changes in technology—ideas such as skilful imitation and representation have been downplayed as parts of aesthetic judgement as mechanical reproduction reduces the skill level required for such representation. Changes in social mores also influence the focus of aesthetic theories: the attempt to distinguish between morally improving art, and pseudo-art which appeals to baser instincts (Tolstoy, 1897; Kant, 1790) seems less relevant as society places less emphasis on art as a driver for these improving aspects.

This has produced a long history of theories that attempt to explain the aesthetic effect of artwork and to define what art is and why it is a distinctive part of human experience. Many of these are focused on the relationship between art-maker, artwork, and audience. The introduction of AI into this space adds a layer of complexity, particularly if the AI system is considered to be an additional component, whether as a sophisticated tool used by the art-maker or as a separate intelligent agent in that set of interactions. These theories have included:

- Ideas about imitation of the natural world and about whether items in the natural world can be considered works of art
- Ideas about the necessity of skill or expertise in the production of the work
- The idea that expression of emotion from the creator to the audience is a key role of art
- The idea that the form or structure of the work is the key aspect that makes soomething a work of art.
- Ideas to do with the intention of the art-maker and the attitude of the audience towards the work.
- Ideas that explain art in the wider context of a social system and where the aesthetic effect is (partially) contingent on the particular social context in which the work is made or exhibited.

Some of the these authors have asserted that their particular criterion is the sole defining criterion for assessing the aesthetic impact of a work. Others, such as Carroll (1999) have argued that art is a family resemblance concept (Wittgenstein, 1953), where there are a number of defining characteristics, and examples of the concept typically exhibit most of those characteristics, but there is not a single one that can be seen as defining by itself. The scope of what are considered artworks shifts over the centuries and in different cultural contexts, which provides challenges to a theory that is dependent on a single core idea; these family resemblance concepts are more robust to these changes.

One such attempt to explain art as a 'cluster' concept is that of Dutton (2009, 2013), who has attempted to outline a set of 'universal features' for art. He is clear that these are not 'criterial for the presence of art', but that they represent a set of practices which, when brought together, broadly characterise much artistic practice. That is, not that every work will have all of these, nor will any of them be found in all art, but that most art will feature most of these characteristics. In particular, he argues that they characterise the uncontroversial core of art, and that too much attention has been paid to difficult edge-cases and examples, in particular those where the point of the work is deliberately to problematise the concept of art. The existence of such extreme cases shouldn't detract from a core of ideas about how the majority of art produces an aesthetic effect. These characteristics are (Dutton, 2009):

Direct Pleasure	The artwork is valued directly in itself as a source of pleasure or aesthetic engagement.
Skill and Virtuosity	The production of the art requires some specialised skills.
Style	Artworks are produced within broad styles, which change and emerge with time.
Novelty and Creativity	An aspect of aesthetic appreciation is appreciation of the novelty of the work.
Criticism	There is a discourse of critique around the world of making and appreciating art.
Representation	Artworks represent aspects of the world (we could, though Dutton does not explicitly do this, extend this to the idea of representing mental and emotional characteristics of the world).
Special Focus	Art is appreciated in special places and times and is the focal point of the audience's attention.
Expressive Individuality	Art is expressive of the individual's personality.
Emotional Saturation	Emotional effect on the audience is a major part of the aesthetic effect of art.
Intellectual Challenge	Aesthetic appreciation involves the exercise of human intellectual capacity.
Art Traditions and Institutions	The production and appreciation of art is embedded in a social network of institutions and traditional practices.
Imaginative Experience	Both the production and appreciation of art involve the exercise of the imagination.

It is interesting to compare these with the various attempts that have been made in the computational creativity literature (starting from (Boden, 1990)) to define the actions typical of a creative agent. The starting point for most of these is that creativity certainly requires novelty—but, that *mere* novelty is not enough(Ventura, 2016). It is easy, at least at a superficial level, to generate (or write code that generates) novel artefacts or behaviour. A program that strings together random words is producing novelty, but would not be regarded as acting creatively. So, it is common to add some notion of *value* to the definition: these might be the values of the creator, of the audience, or might (Dorin and Korb, 2009) emerge from a social network of interactions between creators and audiences. The majority of the characteristics that have been identified in the aesthetics research literature represent aspects of value, such as the idea that audiences value works with a strong formal structure, or that they value works because of their emotional expressiveness.

Other characteristics that have been added include the exercise of skill (Jordanous, 2012; Colton, 2008b), the surprisingness of the outcome of the creative process (Boden, 1990), the appropriateness of the works to their context (Dorin and Korb, 2009), and the intentionality of the system (Ventura, 2016). Again, these replicate many of the characteristics that have been discussed in the aesthetics literature, which has discussed whether skill and expertise are important for aesthetics, and the importance of social factors in determining what is considered to have aesthetic value.

The remainder of this chapter is structured around these ideas that have emerged through the centuries of writings on aesthetics, and the relationship between maker, work, and audience. These have not been structured around a single taxonomy of characteristics such as that of Dutton, but have instead been grouped together from the breadth of literature on aesthetics. These sections are:

- Imitation (Section 1.4)
- Skill and Expertise (Section 1.5)
- Expression (Section 1.6)
- Form (Section 1.7)
- Focus and sake (Section 1.8)
- Imaginative Experience (Section 1.9)
- Criticism and the Artworld (Section 1.10)

Each section takes one such collection of ideas, gives a description of it, and examines the relationship between those ideas and the new means of art-making that AI and search-based art has introduced. This has focused on the philosophical literature on aesthetics; a whole other literature considers aesthetics from the point of view of psychology and cognitive science. Findings from the latter have occasionally been mentioned in this paper, but a more thorough examination of this can be found in a recent paper by Johnson et al. (2019).

In particular, the various fitness bases introduced in Section 1.2.1 are used as a language to discuss various examples of works and research projects that engage with these different aesthetic concepts. However, there is not an attempt to make a systematic comparison of each group of aesthetic ideas with each fitness basis. Instead, each section draws on ideas and specific artworks that have connections with that class of aesthetic theories, and the language of fitness bases is used to clarify the ideas.

1.4 Imitation

Early aesthetic theories focused on the idea of imitation. Aesthetic attention was engaged by the artist producing effective reproductions of items in the world. As technologies such as photography were developed, the need for techniques such as drawing and painting as means of pure representation faded (Carroll, 1999; Benjamin, 1969). As a result, imitation theories of aesthetics moved away from a focus on the reproduction of visual or other sensory features of an object, and towards the idea that what is being represented is the artist's response to the source, or to emphasise some aspect of the source that is not perceived on casual inspection. This leads towards the later development of expression theories of aesthetics, which are explored in Section 1.6 below.

One particularly prominent form of AI art that draws on ideas of imitation is that based on generative adversarial networks (GANs) (Goodfellow et al., 2014), in particular their artistic application as creative adversarial networks (Elgammal et al., 2017). These have been applied to work by artists such as Mario Klingemann (2019) and the Obvious collective (Obvious Collective, 2019).

These are systems based on deep learning (Goodfellow et al., 2017), that start from an *inspiring set* (Ritchie, 2001) of examples—for example, a set of images. The system then searches for two functions: a *generator* that creates examples of images, and a *discriminator* that learns to rate similarity of those generated images to the input set. During learning, both of these systems learn to be better at their tasks, with the end result being that the generator can produce images that imitate the broad style and subject matter of the input set. This is based on a mixture of having a CORPUS of examples, and a co-learning process that fits into the CRITICS AND CO-EVOLUTION fitness basis. The aesthetics of these is based on this ability to imitate, but, interestingly, the ability to imitate not too well; the works generated by GANs have a distinctive style based on the limitations of the models that they learn. For example, one common feature of GANist art when applied to portraiture is that facial features are distorted. The system is, perhaps, attempting to draw something that is a generalisation of *all* of the examples of that feature it has seen, rather than bringing to mind a specific one; or, perhaps, the training time has not been long enough to allow realistic reproduction. As a result of this tension, the end result is distorted.

Two further aspects of recent AI-based art show how ideas of the aesthetics of imitation can be used, but need to be expanded. The first of these is in works that use artificial life techniques to imitate some natural phenomenon, but rescaled so that it can be appreciated on a human scale. The second are works that disrupt and show the inner workings of algorithms that are designed to imitate. The remainder of this section discusses these two examples in detail.

1.4.1 Aesthetic Rescaling

Some artworks that use the ENDOGENOUS fitness basis demonstrate new kind of computer-grounded aesthetic that is grounded in ideas of imitation. The kind of work under consideration are those where the fitness measure does not attempt to measure the aesthetic value of a specific point in the search space, but represents fitness in some kind of evolving or learning system that then generates an emergent system which is the object of aesthetic interest. This kind of work has been surveyed by McCormack (2012). For example, McCormack's 2001 installation *Eden* consists of an artificial life simulation, with a large number of simulated agents which are represented both by projected graphics and by sounds relating to their actions in the simulated ecosystem.

Many phenomena have potential for being appreciated aesthetically, but this aesthetic engagement is difficult because the phenomena happen on vastly different temporal or spatial scales to regular human activity. A scientist who has studied these phenomena might get some aesthetic insight through an appreciation of the data, or mathematical models, or through developing a mind's eye visualisation of the phenomenon. Simulation, and turning these simulations into artworks can reproduce salient aspects of these phenomena at a human scale, making this appreciation available to people who do not have that scientific knowledge and training. This can be seen as a new kind of imitation aesthetic: the computer artwork is reproducing a natural phenomenon at a temporal and spatial scale that facilitates more casual aesthetic appreciation of these phenomena, appealing to the immediate perceptions. This can be seen as a modern-day equivalent of the aesthetic value of drawings of distant flora and fauna in the era before travel, zoos, and nature documentaries; in that era, botanical and zoological artworks acted as translations of these to a wider audience. Aesthetic rescalings analogously transform complex natural phenomena to a scale and location where they can be the object of aesthetic attention.

This has connections to theories of *environmental aesthetics*, that is, theories that explore why people exhibit aesthetic behaviour towards natural scenes such as landscapes. Natural phenomena present problems for many traditional aesthetic theories. For example, expression theories, which characterise art as a means for transmission of emotion from art-maker to audience, fail because of the lack of a maker in natural scenes. These theories fall into two broad categories. The first are those that argue that our aesthetic appreciation of the natural environment is grounded in and enhanced by our understanding natural history and environmental science (Carlson, 1979; Saito, 1998). By contrast, others argue that aesthetic appreciation of nature is a more visceral kind of aesthetic behaviour (Carroll, 1995), and that it is contrasted with appreciation of artworks made by people because people cannot readily separate themselves from the environment (Berleant, 1988), and therefore cannot take the disinterested stance that is often seen as a prequisite for aesthetic judgement (Kant, 1790). Works such as *Eden* based on this kind of *aesthetic rescaling* translate complex environmental phenomena into the gallery, and provide a mid-ground between these two theories of environmental aesthetics, allowing rich scientific theories to be appreciated on a more intuitive/perceptive level.

1.4.2 *Exposing Inner Workings*

Another kind of work that can be interpreted using the aesthetic of imitation are computer artworks that expose an aspect of the inner workings of an algorithm. Machado et al. (2012a) present an algorithm that explores a space of visual representations, using similarity to a collection of human faces as the fitness function. Thus far, this algorithm represents a straightforward application of the CORPUS fitness basis: there is a space of visual representations, and a set of ideal examples that the algorithm is using to guide its search through that space of visual representations. However, the final works that are presented are not the most accurate ones, but are ones with an intermediate fitness value. They represent explorations that the algorithm has made on its way to finding an accurate representation.

Similarly, in the *deep dreaming* images, a deep neural network trained to recognise objects in the world is cut off a few stages before converging on an accurate recognition of a scene (Simonyan et al., 2014; Spratt, 2017). This results in an image consisting, in part, of failed attempts to match learned image schemata to components of the image being analysed. Whilst the process is less explicit, the GAN artworks discussed earlier in this section also have a similar flavour, because of the failings of the GAN system to produce images indistinguishable from the originals. There is a distinctive style to such images, based on inaccuracies of replication, which François Chollet has referred to as *GANism* (Obvious, 2018).

These systems offer an interesting process for producing artworks from algorithms. This process is to take a system that is designed to imitate or understand the world, train it on some examples, and then sample from the middle of that process of imitation, giving a somewhat abstracted representation that nonetheless is grounded in the objects being imitated. This approach has echoes of the automatism of the surrealist movement, such as the automatic drawings of André Masson and others (Montagu, 2002). These attempt to expose aspects of hidden processes, in this case the processes in the pre-conscious areas of the mind, through the exercise of mental discipline aimed at removing the conscious control of the artist's movements. Similarly, these artworks expose aspects of the underlying algorithm.

1.5 Skill and Expertise

The idea that art practice requires skill and expertise is a common idea in aesthetics. It is one of the characteristics on Dutton's (2009) list, it commonly appears in the literature on aesthetics (Carroll, 1999; Gaut and McIver Lopes, 2013), and is a common characteristic in definitions of computational creativity Jordanous (2012).

How is this encoded by the search drivers above? A CORPUS might act as a collection of examples of skilled manufacture or the exercise of skill, and pattern-finding algorithms might extract features that represent that skill. One kind of AESTHETIC MEASURE might be about skill in execution; for some artforms, such as notated music, this distinction between work and its execution is clearer than for others artforms such as the visual arts.

One example of the exercise of skill is the mark-making in nonphotorealistic rendering systems (Collomosse and Hall, 2006). These are systems that take a photograph as input, and output an image that simulates the same image as generated by a particular artistic technique. Some of these systems are based on a search-based technique (Vanderhaeghe and Collomosse, 2012), though they apply this to the overall image, rather than to detailed mark making. In music, this might be seen in systems such as that described by Ramirez et al. (2006), which take a piece of music at attempt to adjust the details of timing, volume, etc. to make an effectively expressive performance; in this paper, the authors use a CORPUS of existing recordings as the starting point.

A related strand of work to nonphotorealistic rendering are the learned *style transfer* techniques (Gatys et al., 2016; Luan et al., 2018). These take a set of source images, sometimes just a single image, and learn functions that can replicate the artistic style of the image by the application of a neural network. These have been very successful at demonstrating the skill of transferring the style, but as creative artworks the results are limited.

Overall, though, this aspect of aesthetics has been rather neglected in the AI art and computational creativity area. Perhaps this is because the problems cluster at two extremes. Some skills that are challenging for people are trivial for machines: for example, accurate perspective drawing. By contrast, skills such as expressive musical performance have to date proven to be very hard computational challenges.

1.6 Expression

One important set of ideas in aesthetics is concerned with *expression*. This is the idea that the artwork is a medium through which the artist conveys emotions and ideas to the audience. From the nineteenth century, this began to take over as a key theoretical frameworks for describing aesthetic experience, taking over from the earlier dominance of theories concerning representation and imitation (Carroll, 1999). Indeed, this is reflected in ideas of artistic practice at the time, which saw the dominance of the idea of the romantic artist who is expressing their reaction to the world, rather than neutrally depicting it Vaughan (1994).

The core idea of expression is that the art-creator transfers an emotional state to their audience (Collingwood, 1938; Carroll, 1999). The artist experiences an inner state, perhaps in response to something in the world, and uses their skills to convey this state to the audience. In a representational medium or style, this might consist of representing some item from the world in a way that emphasises a particular emotional state or trajectory through the way in which it is represented. In a more abstract medium such as instrumental music or abstract painting, the emotional content might be devoid of reference. Tolstoy (1897) has described works of art as a ‘medium of contagion’ from the creator’s emotions to the audience.

It has been noted (e.g. by Saw (1972)) that there is a continuum of emotional response to a work of art, and that responses typically sit in the middle of this continuum. At one extreme, the audience-member recognises the emotion being expressed, but in a purely disinterested and unengaged way. At the other extreme, the person who is so emotionally drawn into a performance that the emotions blur their ability to distinguish between fact and fiction. The creator or performer also needs to have sufficient distance from the emotion to be able to carry out the creative or performative act (Kivy, 1998). This leads to the idea that the artist 'explores it [the emotion] deliberately' (Carroll, 1999), rather than simply 'venting' it.

1.6.1 Can Soulless Computers Express?

One immediate objection to the idea of *expression* in the context of search-based art, and in AI arts and computational creativity more generally, is the computer has nothing to express. A machine has no intrinsic motivation or conscious stance—as Boden (2018) says, 'if a computer is following any goals at all can always be explained with reference to the goals of some human agent'. Given this, it can be argued that the machine has nothing to express. Similarly, Colton et al. (2018) have discussed the complex question of whether autonomous creative systems can be seen as authentic. However, the *unpredictability* of some complex or randomised computer processes can mean that, whilst there is no self-motivated goal coming from the machine, nonetheless an AI art system can generate material that is surprising to 'the person who initiates the process' (Moura, 2018).

This is not a new argument. Similar arguments have been made about artworks produced by animals, and by objects in the natural environment that are the subject of aesthetic behaviour. This is summarised in this, rather crude, passage by Saw (1972, p49):

> *'When asked whether the chimpanzee Congo's pictures were "works of art," almost every unsophisticated person answered that they could not be since a chimpanzee is not capable of of expressing his preferences. [. . .] Rocks and stones worn by wind and rain so that they look like a piece by Henry Moore are similarly refused the title "work of art." In both these cases, sophisticated people tend to say that it entirely depends on the look—if Congo's paintings look well, they may be works of art and so may weatherbeaten rocks.'*

The perspective of the 'sophisticated person' here rather avoids the question, arguing that expression is not a key criterion for something is a 'work of art', and by extension, that expression is not an essential component of the audience taking an aesthetic stance towards an object.

Does this damn any attempt to use a theory of expression to understand the aesthetics of search-based and AI art? One counter-argument is that the program is the medium through which the programmer is being expressive. The position of the *metamaker*—the person who creates a system which then acts creatively—is an interesting one. Traditional art-makers have a rather direct relationship between the movements that they make and what the audience perceives. The expression might fail because of a disjunct between expectations, background, or vocabulary between creator and audience, or there may be a shift in audience assumptions over time, but, nonetheless, the paintbrush has little capacity to put down a radically different line to that intended by the painter, and the word-processor will transcribe faithfully the words in the head of the author.

The paintbrush or piano lies at one end of a spectrum of expressive tools (this spectrum was introduced initially by Rowe (1997) in the context of interactive computer music systems, though it has wider applicability). A little further along are composers using musical notation, or playwrights writing a script, who accept the role of an interpretative intermediary as part of achieving their expression, and might see such interpreters as co-creators of the final work. Further along this continuum are systems, usually computer systems, where the interaction between action and outcome is less predictable—what Sanfilippo (2013) has referred to 'non-random unpredictability'. Indeed, some of the performative value of such systems is in the work of the performer in discovering the responsiveness of a system, much as a performer in a collaborative improvisation has to balance the task of presenting a performance to an audience whilst simultaneously trying to interpret and understand the actions of their fellow performers.

At the far end of this spectrum there is a divide into two main camps. One embraces the unpredictability, and plays with the idea that, as pattern-finding agents, members of a human audience will discover patterns in random actions, different performances presented simultaneously, or performances executed with skill and conviction but with unpredictable outcomes. This is exemplified in what Cage has described as *experimental actions*, defined as 'an action the outcome of which is not foreseen' (Kostelanetz and Cage, 1987; Joseph, 2016). Such works are hard to explain using expression theories; indeed, a book collecting literary works that are based on found objects is called *Against Expression* (Dworkin and Goldsmith, 2011), suggesting that the intention of creators working with such techniques have little time for expression theories as the aesthetic basis of their work.

This unpredictability has been highlighted in an AI context by Moura (2018), who has produced robotic painting systems where multiple mobile robots ('artbots') move around on a single surface to create a collaborative work. One of the inputs into the artbots is an image coming from the drawn trails made earlier in the drawing. This draws inspiration from stigmergy, the process whereby swarm insects learn from each other not by direct communication but by following each others' chemical signals in the environment (Hölldobler and Wilson, 1990). This has been used widely in swarm intelligence algorithms (Bonabeau et al., 1999). Moura (2018) notes that 'the resulting art works cannot be predetermined even by the person who initiates the process'.

Nonetheless, in such systems, there is some decision by the human creator about which system to use as the basis for the experimental action. Looking at John Cage's *Atlas Eclipticalis* (Cage, 1961), which generated music by transcribing star-maps onto musical staves, it is difficult to imagine that there was not some thought about the interestingness of those shapes for musical purposes. Similarly, with Moura's artbots, it is clear that the stigmergic, interactive nature of the process has been deliberately chosen to generate interesting material. There is detachment by the creator from the details of what is generated; but, nonetheless, the process has still been deliberately chosen.

The other kind of works to be found at this end of the spectrum consist of algorithmic systems that discover how to communicate a thought/emotion, perhaps in a way that could not be anticipated by the creator of the system (the *metamaker*). This could be because the system is able to perform skills that the metamaker is unable to perform, or because the medium of expression requires processing data on a scale that could not tractably be processed by a human.

Alternatively, there is an intermediate point in between these two camps, where the metamaker is acting as a kind of facilitator of expression. For example, some computer-based systems will gather together the expressions of many individual people, presenting them in a new way, either as a collection or by applying some kind of pattern-finding algorithm. This kind of meta-expression might be phrased as building a system that captures the zeitgeist of a moment—not the individual expression of the creator, but a *collective expression*. This can also be seen in systems such *Electric Sheep* (Draves, 2005), where individuals interact with a single image on their screen, but these are collected and recombined on a central server that is running an evolutionary algorithm. This combines elements of the HUMAN INTERACTION fitness basis, but the collaborative aspect demonstrates aspects of the CRITICS AND CO-EVOLUTION basis too.

Of course, such unforeseenness can go wrong; the metamaker can discover that their creation creates something which they would never have wanted to create. Consider the Microsoft *Tay* chatbot, which learned what to say from interactions online. This can be seen as being an example of something driven by a mixture of the CORPUS fitness basis, because it learned from its collection of interactions. It also has aspects of AESTHETIC MEASURE in the the form of the programmed-in biases of the algorithm for how to rate interactions as worthy of learning from. It was taken offline by its creators in less than a day because it was posting offensive material. Clearly, this posting is not something the creators wanted to happen, and is certainly not expressive of their viewpoint, as evidenced by the rapid takedown.

This can be seen as a new kind of co-creation; call this *collective expression*. The creators of the bot didn't express the content, but they facilitated the means for a wide group of people to pool their expressions. In this case, this expressive potential was adopted parasitically by a group of people who made a coordinated attempt to make a particular, offensive, expression, by interacting with it in a coordinated way. A better constructed system might be able to realise the advantages of this without the problems, but this is challenging.

This idea of a creator making something that is designed not to express their own emotions or ideas, but to facilitate new ways of combining or bringing together expressions by a collection of people, is something that technology is particularly suited to. This is hard to account with using traditional theories based around the idea that creative acts should be tracable to a single source—a single locus of responsibility for that creative act (Johnson, 2014). To bring these systems into explanatory frameworks such as expression theories needs new ideas of co-creation and collective expression.

1.6.2 Expression without Transmission

An alternative perspective is that a work can be expressive without being transmissive of emotion. Carroll (1999) discusses the case of a mystery writer, who 'need not feel suspense as he ratchets up the audience's apprehensiveness'. From this perspective, causing an emotional response in the audience does not need to involve the creator in having any related emotional state at all—merely, having the techniques to generate an emotional response. The lack of authentic emotion in the agent doing the transmission doesn't invalidate the effectiveness of the emotional expresssion.

Such effects are found in a number of AI-based artistic systems. One example is the collage painting system that is part of the wider *Painting Fool* system (Krzeczkowska et al., 2010). This system takes as input a newspaper article, uses an image search to discover images relevant to the article, and then produces a non-photorealistic collage which uses simulated painting techniques to bring the collage into a coherent visual style. The results of this can be emotionally engaging, and can comment on a specific topic. The status of this agent seems no different to the mystery writer discussed above; again, there is no emotional activity by the agent producing the work, but this doesn't invalidate the emotional content of the work produced.

The mark-making component of the Painting Fool has a particular focus on emotional engagement with the audience. In its mark-making, the system searches within a parameterised set of marks, allowing it to find, for example, ways of making marks that are in between pencil marks and paint strokes. These are used 'to discover novel painting styles to enhance emotional content' (Colton, 2012). The system can discover bespoke mark-making techniques by search within its parameter space to convey a particular emotion. This was done by human feedback to the system, with the human trainer of the system identifying examples where the style chosen enhanced the desired emotional state; a kind of HUMAN INTERACTION. Indeed, the creator of the system emphasises that 'computer generated paintings can still evoke emotions in viewers without necessarily modelling human emotions' (Colton, 2012).

The MEXICA model of creative writing and storytelling (Pérez y Pérez and Sharples, 2001) uses an idea of emotional trajectory in building its stories. The choice of the next action within a story is guided (alongside other factors) by a measure of *tension* (a scalar variable). Certain actions are associated with increased tension—for example, the setting up of a conflict between two characters—and others reduce the tension. This bias changes as the story progresses. At the beginning actions that increase tension are favoured, to increase reader's engagement in the story. Towards the end, the opposite is true, with the system preferring actions that reduce tension and thus create a satisfying resolution to the story. Again, there are no actual emotional qualia in the system to be expressed, but the creator has imbued the system with a technique for eliciting an emotional trajectory in the reader. This can be seen as an AESTHETIC MEASURE—the creator of the system has made the decision that good stories consist, in part, in this build up and release of tension, and the tension measure drives the search towards stories that satisfy this.

1.6.3 Expression, Emotion, and Expression Systems

The *Painting Fool* (Colton, 2012) is a computer art system that exists in a number of versions, many of which are concerned with creating non-photorealistic renderings of photographic input in a variety of artistic styles. Colton et al. (2008) introduce a version of this system for portraits that is responsive to the emotional state conveyed by the input. A neural network based facial sentiment analysis system classifies the input photograph into one of six of basic emotional categories (happiness, sadness, surprise, etc.) and then uses a set of brushstrokes, segmentation style, colour palette, and other features (chosen by the creators of the system) to create a nonphotorealistic rendering appropriate to the emotion being expressed in the photograph.

A more recent version of the system explicitly annotates pictures with some text concerning the emotional state of the artist at the time that the picture was created. This is an example of the use of *framing information* (Charnley et al., 2012)—additional, usually textual information, which puts a creative act into a wider context. In this case, this information consisted of a description of why it chose the specific style—asserting that this based on *its* mood, not the mood imputed from the human photograph being used as input—and an evaluation of its success in creating a picture that matched that mood:

> *'Like a human, it's sometimes pleased with its work and sometimes disappointed. "I was in a positive mood. So I wanted to paint a patterned portrait," it wrote in response to the portrait above. "This is a miserable failure—I'm very unhappy about that. And I'm also annoyed that the portrait is bleached, because that does not suit my mood." '* (Stromberg, 2013)

What, if anything, is being *expressed* here? Clearly, the computer is not in any traditional sense feeling those moods; there does not seem to be any meaningful mechanism for a computer to have emotional qualia (Picard, 1997). Nonetheless, it can simulate or act out having moods, based on external factors that might be shared with its audience. The Painting Fool's expressed emotional state 'depended on where is had recently been in terms of reading the newspaper articles.' (Sayej, 2013). So, whilst it might not be expressing its own feelings, it is drawing on context to bring itself into *expressive alignment* with emotionally-capable agents in its environment.

Is such a system conning its audience? Would a naive viewer of these paintings change their view of the aesthetic value of them on learning that there was no experienced emotion underlying statements about mood? Responses to such revelations vary. The revelation that James Frey's book *A Million Little Pieces* was largely fictional, despite having been framed as autobiographical, was damned; the publisher going as far as to 'provide refunds to readers who felt they were defrauded' (Barton, 2006). Yet, it is accepted that a stage magician will be lying throughout their act about the cause of the effects presented.

One area where this has been explored in depth is in acting. An actor's relationship with the emotions explored in the text can exist on a continuum from pure 'technique', where the actor experiences none of the emotional qualia but expresses the emotion through learned gestures, expressions, etc., through to a 'method' approach at the other end of the continuum, where actors deep engagement in actually feeling the emotions being expressed by their characters (Goldstein and Winner, 2012; Taylor, 2006). Indeed some actor training techniques explicitly mark themselves out as being in between these extremes, such as Alba Emoting (Baker, 2008), which presents itself as being a system that can generate effective emotion without the need for techniques such as personal emotional recall. Notably, the aim of the method approaches is to generate a more believably expressive performance, rather than the audience being expected to care more about the performance because they have been given framing information that such an approach has been used.

So, are systems such as the Painting Fool with emotional framing achieving their aesthetic effects through expression? Again, such systems appear to consist of expression without transmission. There is a kind of *expression system* which, without directly transmitting an emotional state from the system to the audience, nonetheless uses contextual clues such as the mood generated by newspaper articles to elicit emotional states that are broadly aligned with the environment in which the system is working. The system will not make an emotional error such as producing happy work on a day of great tragedy, because it uses shared information with the audience to align its actions with the environment.

1.6.4 Self-contagion

One form of expression peculiar to systems based on the HUMAN INTERACTION fitness basis is that the audience member plays two roles. The first of these is as the person driving the expression of the system, and the second is receiving the results of that expression.

Consider a system such as NEvAr (Machado and Cardoso, 2002) or the interactive ant paintings by Aupetit et al. (2003). In systems such as this, the viewer is presented with a set of possible artworks drawn from the search space, and invited to select their preferred ones, or to rank them. At this point, the viewer is in the mode of making aesthetic judgements, making a value judgement between the different works. Once this judgement has been made, that judgement becomes an expression—an expression of preference, but also an expression of the viewer's feelings which is mediated through the system as a developed or meditated version of their expression.

Thus the system allows for a kind of *self-contagion*—the viewer finds the germ of an idea in the items presented by the machine, and the development of this germ is supported by the system in a positive feedback cycle, where the viewer collaborates with the system in developing that idea. Again, there are resonances of the continuum of interaction discussed by Rowe (1997). These interactive systems fit in the middle of Rowe's continuum, being inbetween the extremes of the predictable tool/instrument and the independently-minded collaborator.

Such systems can place the viewer/listener into the position of an art-maker, developing ideas from their errors and contingencies. It is an common pattern in art-making for artists to extract ideas from their own errors. An artist will make a slip of the pencil during a drawing, but then develop that error into a part of the drawing. An improvising musician will make an unintended sound, which then gets takes up by another improviser in the group and developed further, taking the improvisation in a new direction. This has a long history; consider the following advice given by Leonardo da Vinci (quoted by Turner (2011)):

> *'If you look upon an old wall covered with dirt, or the odd appearance of some streaked stones, you may discover several things like landscapes, battles, clouds, uncommon attitudes, humorous faces, draperies, etc. Out of this confused mass of objects, the mind will be furnished with an abundance of designs and subjects perfectly new.'*

Saw (1972, p146) has attempted to unpack the difference between skilled artists and laïcs in their exercise of imagination:

> *'. . . a non-artist might say "If it is imagining, anybody can do it. I can imagine paintings, vases, plays, and musical compositions. The difficulty is in making the thing.". This is not so at all. We only think that we have imagined these things, because we do not understand that imagination must be complete in every detail—in fact, to imagine it is to complete it.'*

Perhaps this a little simplistic about the creative process—it is not always the case that the art-maker conceives the whole work before beginning 'making the thing'. Furthermore, there is a feedback loop between actions of making, the unanticipated consequences of those making actions, and future imagination. Nonetheless, the systems that are driven by the HUMAN INTERACTION fitness basis provide some means of bridging this gap between naive imagination and realisation.

This ability to blur the distinction between audience and art maker via this process of self-contagion is a distinctive expressive aesthetic of these human-in-the-loop systems.

1.6.5 Expression: Summary

At first, expression theories of aesthetics appear to have little to do with AI-based artworks. As machines have neither a sake nor emotional qualia, it is hard to see that they could either want to express or have any emotional state to express.

Yet, there are ways in which the behaviour of computer-artists can engage with ideas of expression. One is through facilitation of expression in new ways. Firstly, a system can exist that doesn't express the emotions of its creator, but which allows people to come together—either by making deliberate contributions to the system, or by the system learning from material posted online for other purposes—and contribute towards a collective expression. Secondly, we can divorce expression from transmission, and consider system that are able to elicit an emotional response in an audience despite having no emotional qualia themselves. Moreover, these *expression systems* can align themselves with the audience by drawing on their shared environment with the audience. Finally, some systems can act as *self-contagion* systems, allowing the viewer/listener/etc. to explore and reinforce a germ of emotional expression in a positive feedback cycle.

Can such systems be driven by the search drivers discussed earlier? Clearly, the self-contagion systems are embedded in the idea of HUMAN INTERACTION. Expression without transmission can be achieved from the AESTHETIC MEASURE driver, where the measure is some kind of closeness to a pre-specified emotional trajectory. It would be interesting to see if an expressive system could be built around the CORPUS driver. That is, the emotion (or emotional trajectory) to be expressed would be abstracted from one or more existing examples. This could perhaps be achieved if the computer analysis of the corpus were achieved by sentiment analysis of the members of the corpus, and then pattern finding applied to the results of that analysis and used to drive the search process—there are similarities here to a human learning a creative art by studying existing works and examining their emotional content and how that content is achieved.

Could a computer system be genuinely expressive? That is, not expressive of a simulated emotion or of a canned or learned emotional trajectory, but to express something about its state as a machine? Computational creativity has largely shied away from creating works that reflect on the computational. McCormack (2005) has argued that computational creativity systems should shy away from a 'technical fetish', that audiences should appreciate the creativity and artistic achievement of an artistically creative AI system as such and not in the context of it being a computer system. Similarly, Colton (2012) has argued that one aspect of audience engagement with AI artworks is their 'awe at the power of modern computing'.

Colton (2012) has stated the ambition that the Painting Fool system should 'one day be taken seriously as a creative artist in its own right'. Yet, this seems like only the first step. If AI artists were to be accepted on a par with human artists, perhaps the next stage is for these machines to express—not their own experience, nor the emotional states that they lack, but some essence of computer. Perhaps environmental aesthetics might offer a framework for this, being a set of theories that discuss aesthetic value in the context of a system that generates aesthetically appreciable objects without needing a maker that is conscious of what it is creating.

1.7 Form

Another major set of aesthetic theories consider ideas of *form* as being a key aesthetic criterion. That is, a major component of aesthetic judgement and appreciation lie in the formal, structural aspects of a work. Audiences are assumed—whether through some preferences that have been evolved over generations, or through acculturation—to appreciate aesthetically certain formal structures. This is one of the 'universal features' identified by Dutton (2013), as a major part of what he terms *style*: 'art objects and performances, including fictional or poetic narratives, are made in recognizable styles, according to rules of form and composition.'

That is not to say that content is irrelevant or absent. The main point of these theories is that form—'significant form', as Bell (1914) describes it—is the main differentiator between mere content, and broadly similar content that provokes aesthetic behaviour from its audience. Two photographs might contain representations of largely the same set of objects, but, formal theories would argue, one is a descriptive, journalistic photo, and the other an artistic one, based on the latter having a formal structure that is absent from the former.

The traditional way in which form is realised by human artists is by the artist making explicit decisions about the position, orientation, and scale of objects. In many AI art works, the decision of the artist is which algorithm to use. In works such as Greenfield's avoidance, ricochet and deflection drawings (Greenfield, 2009, 2015, 2016) and Moura's artbots (Moura, 2018), the artist decides upon the algorithm, and the form emerges from the interactions.

A form could be a very general structure, such as an idea of symmetry or balance. For example, Birkhoff (1933) discusses a very general notion of 'aesthetic measure' M given by the broad-brush formula $M = O/C$, where C represents some measure of the (perceptual) complexity, and O represents a way of quantifying the order or organisation of a work. This is clearly an example of a formal notion of aesthetics—neither of the components of the formula are concerned with the content of the work, they are concerned with the organisation of those components. Form may also be specific to a particular artform. For example, musical works might have a form described in terms of a sequence of harmonic structures.

Algorithmic congruence is the idea that a coherent multimedia artwork can be created by realising the same algorithmic process in different media (Evans, 1987). For example, an an algorithm generates a sequence of numbers which are interpreted in a coordinated way as the animation of shapes on a projection, and the generation of musical notes. This is grounded in a form-based argument—the idea that it is the structural forms that elicit our aesthetic response, and realising those same algorithms in different media produces an aligned response from the audience.

1.7.1 Form and Aesthetic Measure

The most direct way in which these ideas of form interact with search-based art is in those systems that base their fitness directly on AESTHETIC MEASURE. Most of the proposed systems for aesthetic measure are based on measures that are concerned with form rather than content. For example, the *Ricochet Composition* images evolved by the program by Greenfield (2009) are driven by a combination of measures such as the proportion of the image filled with colour, and the amount of symmetry in the images. The driver for the search is aiming to optimise formal aspects of the works produced. On a different scale, Bird et al. (2003) discuss a search-based algorithm for the curation of works in an exhibition, choosing and grouping works to achieve clusters of semantically-related works, each cluster at a different site within the exhibition.

Colton (2008a) proposes a meta-level approach to aesthetic measures based on form (inspired by the ideas of Buchan (2001) relating creativity and meta-reasoning). Rather than choosing a specific formula for measuring the aesthetic quality of a scene, the concept-generation system HR (Colton, 2002) is used to construct fitness functions. The idea of systems such as HR is that they take as input some simple base concepts, and then use a set of rules for combining and transforming concepts, together with high-level concepts of interestingness, and generate a body of new concepts by building out from the initial concepts. In this system, the base concepts are the descriptions of the basic graphical elements that are combined to make an image, and the system builds up a network of concepts, choosing one of the more high-level ones to use as its fitness function.

This can be seen as a high-level idea of form. A good fitness function is a sophisticated and neat theory about how components are put together, but isn't strongly related to the content, and certainly not to the external connotations of the content or its emotional connections. Indeed, the basic content in this system is supplied by the user as the initial base concepts.

A related idea is Schmidhuber's 1997 *low-complexity* art, which also uses a high-level form-based notion of aesthetics. In this case, that art generated using a small algorithm will be aesthetically engaging and that the 'algorithmic simplicity' of the image will be perceivable by the viewer and form part of their aesthetic judgement.

In creative writing and storytelling the link between form and content is more complex. In visual art or music, many aspects of form can be measured by simple algorithms and pattern-finding systems. In literary works, unpacking the structure of, say, a fictional story, requires a more complex cognitive engagement with the meaning of the story before its form can be described. For example, identifying a structural idea such as a dispute between two characters, and the subsequent resolution of that dispute, requires the reader to read the words, bring characters and their interactions to mind, and then match the interactions between those characters to a pattern. Rarely is form in literature concerned with basic material such as the choice of letters or words.

As a result, generate-and-test methods for driving stories towards an aesthetically satisfactory form, at least at the level of the final text, is difficult for computer systems, because it is difficult for the machine to extract such structure from the text. Most computer story generation systems (Gervas, 2009) use a two-stage process, where the initial story generation system uses a graph of actions in the story, and a database describing how characters and their dispositions, locations, and objects, change over time. The AI search algorithms are applied to this representation, and then an second process converts that more formal structure into text.

This allows aesthetic measures to be applied that drive the search towards formal structures in the story that the metamaker of the system has decided are aesthetically valuable. For example, in the MEXICA system (Pérez y Pérez and Sharples, 2001), a sequence of measurements of dramatic tension is measured for a story. If in the formal structure of the story, an event occurs that the metamaker of the system has decided increases tension (e.g., a secret is revealed), then the overall tension is increased; a tension-resolving event (e.g., a conflict within a family is resolved) has the opposite effect. This allows the system, through a sequence of measurement processes followed by modification processes, to search for a story with a particular formal structure: usually, a broad increase in tension in the first half of the story, and then a decrease towards the end.

Need all aesthetic measures be formal rather than content-based? One difficulty with computer-based measures of content is that they require a large amount of understanding of the meaning of the work being created, rather than its formal features. This requires both a depth of understanding largely beyond current AI systems, and a breadth of contextual understanding. One possibility is to base such measures on the 'connotational' value (Johnson, 2012a) of an object—how it fits into a network of concepts, which might be changing with time and responding to external events—rather than purely form-based measures.

1.7.2 Form and Corpus

Another way in which the idea of form is found in search-based art is in those systems which use the CORPUS fitness basis. Systems of this type often use the corpus of examples as an 'inspiring set' (Ritchie, 2007), that is, the system should take inspiration from the corpus and produce novel examples that fit into the same broad style.

Typically, this is achieved by using a machine learning algorithm to extract the key features of the inspiring set. By a similar argument to above, such features are more likely to be form-related than content-related. Most machine learning algorithms work from the information provided alone, without reference to a wider world context, and this amongst other things makes it difficult for them to extract patterns from the corpus that are focused on meaning or content. By contrast, machine learning algorithms can readily identify forms (at least, those forms that do not require wider context), and so are biased towards extracting formal features that are common or widespread in the corpus. Gatys et al. (2015) have noted that a neural system designed for computer vision 'automatically learns image representations that allow the separation of image content from style' and link this to human ability to engage aesthetically with art.

It would be interesting to consider how to build a system that learned content features rather than formal features. If that were achievable, it might be possible to build a tunable pattern-extraction system, the biases of which could be tuned between form- and content-based features.

1.7.3 Form and Multicriterion Optimisation

Thinking about notions of form from a computational point of view is useful not only from the point of view of understanding how AI art systems can be built and unpacking the assumptions underlying existing systems of that kind; it also provides a set of tools of thought with which to unpack and critique the philosophical notions being considered (see the discussion by Sloman (1978)).

One criticism of form-based notions of aesthetics is that, once laws of form are clearly articulated, construction of ideal works of art should be trivial—artists should simply follow those laws in the production of work. A criticism of that idea is that laws might oracular rather than constructive—it might be possible to describe the aesthetic qualities of a given work, but the laws might not provide any constructive guidance as to how a work can be improved. Ideas of optimisation partially work against that criticism—given aesthetic measures of form, an optimisation procedure can be applied with that measure as its objective function. Again, ideas from computation come into play. Such functions might have local minima, thus disrupting naive attempts to optimise against them.

Furthermore, there is nothing in the notion of laws of form that they are reconcilable with each other, or measurable on the same scale. An artist—whether a search-based AI artist or a human one—might spend a lot of time trying to reconcile multiple laws of form which cannot be readily reconciled. As such, a major task for such an artist might consist in finding a point the search space of possible works that represents an appropriate balance or trade-off between multiple laws of form—computationally, the idea of multi-criterion optimization. This may be one of the tasks that artists are engaged in in their preliminary work such as sketching. Indeed, it would be interesting to take ideas of form and see whether ideas from multicriterion optimization such as the Pareto front are useful for examining how differing requirements are traded off.

1.7.4 Form: Summary

Theories based around form, along with theories of expression, some of the most important ideas in aesthetics. The idea of optimising against some measures of form is a key way in which aesthetic theories can be translated into an AI search procedure. Nonetheless, such measures might not be commensurable and might need to be traded off against each other. This leads to the idea of multi-criterion optimisation of different aspects of form in a single search process.

The idea of valuing form aesthetically leads to the idea that artworks should have some underlying theory, and this leads to another intersection between AI and form. This is the idea of systems that discover their own laws of form, by searching for patterns or explainable regularities in a work, using an AI system that searches for compact explanations of pattern. This has the potential to lead to open-ended search based systems that con not only optimise against existing aesthetic measures, but can also discover new measures.

1.8 Focus and Sake

Another of Dutton's (2013) characteristics is that art is the subject a 'special and dramatic focus of experience'. Whilst art may occupy an incidental or instrumental practice in some aspects of life, in most cultures time, space and attention is put aside for engagement with artistic practice—what Dissanayake (1997) terms 'making special'. Another of Dutton's criteria is that art exhibits 'nonutilitarian pleasure', that is it is 'a source of pleasure in itself, rather than as a practical tool of source of knowledge'; it is appreciated for its own sake.

The typical outcome from the application of search-based processes in art is that the works should be exhibited and performed, and therefore be the focus of aesthetic attention for their own sake. Yet, the nature of this sake is more complex than for artworks created by people. The focus of audience attention is, however, on the fact of it having been produced by artificial intelligence, rather than on the work as such; what McCormack (2005) calls 'technical fetish or fascination'. He argues that search-based art should be capable of taken seriously for its artistic contribution rather than the focus of experience being on the machine, and makes connections to the Turing test. This could change in the fullness of time: if AI-generated art becomes accepted, perhaps eventually some computer artists make art that reflects on their status as machines, with the focus being on how this is expressed rather than mere astonishment that it can do it at all.

The HUMAN INTERACTION fitness driver provides a distinctive form of 'focus of experience' in that it, within a world created by the metamaker of the system, allows the audience member to create and develop the work being exhibited. This differs from many other forms of interactive art. In most interactive artworks, there is a sense of the (extant) work existing, and responding to the audience prompts. Typically, this will be reactive—the audience member does something, the work responds in some way, then settles back into its base state. In others, the interaction is more conversational between work and audience. In the search-based works with human interaction, the sense is more of exploration of a space of possible works—there isn't a fixed work that the audience member is interacting with, they are playing a different role, one of guiding the creation of the work, or exploring the space of possible works.

This idea of special focus also comes into the ideas discussed earlier about works based on ENDOGENOUS fitness. Here the focus of attention is both on the work itself as a work of art, but it also provides a means for people to focus on natural processes that take place on temporal or spatial scales that are not normally appreciable to human perception.

1.9 Imaginative Experience

Dutton (2009) places an especial emphasis on imaginative experience as a characteristic typical of art. That is, both the creation of and the appreciation of artworks relies on the exercise of the imagination. The focus here will be on the creation of work.

The idea of imaginative experience in aesthetics has many links with with arguments about novelty and creativity in the literature on computational creativity. In particular, there is a long debate about whether a search process can generate something novel (Boden, 1990, 1998; Wiggins, 2006b); ideas of imagination and novelty are closely linked. The argument against machine novelty revolves around the idea that once a search space is defined, the creativity of the system is constrained; anything the system does is 'mere generation' within an already-defined systems (Ventura, 2016). However, search spaces can be vast; in a vast enough search space, the search *process* becomes much more significant.

Other that the production of *novel* ideas, what could it mean for an AI system to be said to be exercising *imagination*? One idea is to simply that it is capable of exploring a search space in a way that leads to aesthetically valuable patterns; this is just a re-posing of the question of why people find certain patterns to be aesthetically valuable.

A deeper idea of imagination could be that the system is capable of bringing together different ideas. This could be literal visual ideas, as explored by systems such as *Vismantic*, which brings together aspects of two source images in an attempt to produce a new image that has aspects of the two. The idea is that the new image is a coherent image in its own right, but also one that aspects of the two source images. At a more abstract level, this idea of imagination as combination leads to ideas of the computational generation of metaphor (Indurkhya, 2010), which has been explored more extensively in computational linguistic creativity (Veale, 2012) rather than in visual creativity. Further developments in this area are likely to depend on computer systems that can build up a large store of knowledge about the world (Cambria et al., 2009; Silver et al., 2013; Hart, 2017) so that the system can draw on this knowledge in exercising its imagination. This ability to draw on a vast store of information, and make novel connections, is at the core of what is meant by imagination for people, and would be one way to develop the imaginative sense in computers too.

1.10 Criticism and the Artworld

As discussed at length by Danto (1964, 1981), art does not exist in isolation. It is contextualised both in an 'artworld' of creators, audiences and critics, and in a wider society. In particular, the same object might be regarded differently at different times and places because of this contextual information. A small number of AI art systems have started to explore this, and search can be driven by CRITICS AND CO-EVOLUTION fitness drivers. In these, the search space is not just one of artistic creators, but a second population exists of critic agents that comment on the art being produced and influence how the search progresses (Machado et al., 2004; Romero et al., 2003; Machado et al., 2008; Greenfield and Machado, 2009; Romero et al., 2009).

There is a connection between this idea of critic agents and the GAN art (Elgammal et al., 2017), discussed in Section 1.4. A GAN learns two models from the data presented to it in a training set. The first is the *generator*, which builds examples based on the training data, and the second is a *discriminator* that learns to rate the similarity of the generated images to that training data. In applications of GANs in art, it is the outputs from the generator that are presented as artworks. The discriminator, however, can be seen as a critic/evaluator that emerges as a natural part of the design process. Perhaps as such critics become more sophisticated, they could be separated from their original context and become a useful output of the system in their own right, as a learned model of criticism.

1.11 Aesthetics as a Cluster Concept

The sections above have reviewed AI art and search-based art making through the individual lenses of particular aesthetic theories. It is important, however, to also consider the idea proposed by Dutton (2009, 2013) that art and aesthetics are 'cluster concepts'. That is, there are a number of characteristics that are typically found in artworks and human response thereto, but none of them are necessary or sufficient.

Section 1.7.3 introduced the idea that different, orthogonal notions of form can be explored using multicriterion optimisation. This could be generalised beyond a single class of aesthetic ideas. One way to operationalise aesthetics as a cluster concept is to see each of the characteristics as one (or more) orthogonal dimensions to be explored using a multicriterion optimisation system. So, for example, one dimension might be formal aspects of the artwork being produced, another its emotional expressivity, etc. The idea of combining multiple aesthetic measures has been explored by den Heijer and Eiben (2011), but all of the measures draw on a broadly similar class of aesthetic theories. Along similar lines, Vouliouri (2011) explores the idea of using multi-criteria optimisation to balance aesthetic and functional criteria in automated design.

A related idea is found in modular AI art systems such as the *Painting Fool* (Colton, 2012). Such systems contain modules that handle different aspects of the system.

For example, on module might be concerned with placing objects in a symmetrical fashion, another with analysing the emotional response of people to the artworks it is generating, another with generating framing information to accompany the work. Misztal and Indurkhya (2014) have articulated particularly clearly one way of realising such a system, by using a *blackboard* architecture. This is where different putative components of a work are stored on a so-called 'blackboard' (a data structure) and various agents work to modify the components and bring them together, each of which is concerned with a different artistic skill. Furthermore, some agents are concerned with aesthetic evaluation, and multiply or remove items from the blackboard according to these evaluations. These demonstrate particular architectures that realise the idea of aesthetics as a cluster concept.

1.12 Conclusions

This paper has examined various theories of aesthetics and their relationship to the production of art in various media through AI systems, particular search-based systems. Four questions were introduced towards the beginning of the paper. By means of a conclusion, these will now be revisited.

1.12.1 Aesthetic Theories and AI Art

The first question asked whether traditional aesthetic theories can be applied to search-based AI art making. The key question here is whether the search drivers—the ways of scoring or ranking objects in the search space—are based on ideas from traditional aesthetic theories, or whether they are based on entirely new ideas. Overall, there is a strong alignment between the various characteristics of traditional aesthetic ideas and the search drivers used in AI art. For example, aesthetics of form are realised in AI art through measures of symmetry and balance, the idea of imitation in works that draw on inspiring sets, and theories of expression in systems that use humans in the search loop. This gives us confidence in these aesthetic theories—even with the radical shift presented by AI art, the underpinning aesthetic theories are still of relevance. Nonetheless, AI artworks can drive the expansion of these aesthetic theories. These are considered in the next section.

1.12.2 Expanding Theories for AI Art

The second question asked whether traditional themes in aesthetic theory need to be expanded or adapted to examine these new forms of art.

Early aesthetic theories concerning **representation and imitation** get new life through the idea of *aesthetic rescaling*. That is, the idea that technology can simulate and visualise phenomena at new temporal and spatial scales, allowing the audience to have an aesthetic engagement with phenomena that are not normally appreciable to human perception. These ideas also link to aesthetic theories of focus of attention and 'making special'.

Theories of **expression** provide a challenge for AI-based art, because machines do not experience qualia and emotions to be expressed. Yet, these theories can be re-interpreted in light of AI art systems. One way is to consider systems as expressions of their metamakers, the authors of the programs; that opens up an interesting discussion about the degree of control that is desirable between the metamaker and the final product. Another interesting area is those systems where having a human-in-the-loop allows that human to explore means of expression that are not readily available to them without the computer system—an expression facilitating or self-contagion system.

Theories concerned with **form** seem initially more amenable to computational realisation; formal aspects of many artistic media can be measured algorithmically and used as search drivers.

Other aspects of aesthetics have been more neglected in AI art to date. **Artistic skill** is rarely a focus of attention in AI artworks. Despite a couple of attempts, systems based around the idea of **critics** and embedding an AI artmaker into a **wider artworld context** are still at an early stage.

An interesting challenge for future work is the idea of new aesthetic theories that are distinctive to autonomous creative AI systems. The ideas of self-contagion and aesthetic rescaling above show how existing ideas in aesthetics can be extended for computational creative systems. Recently, authors have begun to explore new aesthetic ideas, such the the idea of *explainability* as a meta-aesthetic explored by Bodily and Ventura (2018).

1.12.3 Gaps and Assumptions

Are there gaps or hidden assumptions in our ways of understanding AI-based art? One point that emerges from the discussion in this chapter is the emphasis on aesthetics as a cluster concept. Because of the emphasis in AI search on the fitness function, there is a corresponding tendency for designers of such systems to choose a single aesthetic perspective as the driver of the system. Instead, this cluster-concept view emphasises the value of systems that have attempted to incorporate multiple aesthetic theories into a single system. Another concerns responsibility and authorship: the existence of AI-created art problematises the idea of who or what created the art. Should the programmer of the AI systems used in creating the art be afforded with some credit? Should the system itself be credited?

1.12.4 Understanding Human Art-making

Can computational and AI concepts help us understand human art making? The search based perspective on art making provides a formalisation of spaces of possibility and the criteria that are used in moving through them. The classification of creativity into exploratory, combinational and transformational (Boden, 1990) gives a new way to think about the development of artistic style.

The requirement to be explicit about the driver of the search provides a new perspective on potential drivers for human art-making. Are similar drivers at work in human art making? Or, are human art-makers much more driven by an 'engagement-reflection' cycle (Sharples, 1998) where they are constantly swapping between the process of making and the process of reflecting on and contextualising what has just been made in order to drive the next moment of making? One radical difference between human-art making and search-based AI art is that in the latter the drivers are almost always decided before the search process begins. By contrast, the human artist appears to generate these during the development of a work, drawing on a lifetime of experience, knowledge, and emotion.

One useful concept that emerges from the discussion in this paper is the idea of trying to satisfy multiple criteria simultaneously. One of the great challenges of art is to produce a work that is *at the same time* for example formally well structured, emotionally expressive, and sufficiently well imitative of its subject matter that such representations can be understood by the viewer. Even within one set of aesthetic ideas (e.g. that of form), it may not be possible to simultaneously satisfy multiple aspects of form within the same work, necessitating a trade-off. Looking at art making as a search process formalises this in terms of concepts from multi-criterion optimisation, and provides a language to describe such processes.

1.12.5 Future Directions

What challenges does this provoke for future work in AI art and computational creativity? One area is putting the AI art system into a wider context. One direction might be to build a system that tries to recapitulate through simulation the evolutionary origins of aesthetics (analogously to how researchers have simulated the origins of language (Nowak and Krakauer, 1999)). Another context is that of a wider 'art-world' (Danto, 1964), where artists, audiences, critics, and a wider contextual world interact; this consist entirely in a simulated world, or could involve search-based art which pays attention to the external world of human art and wider society.

This discussion of audience and critics leads to a broader question of AI art critics and audiences. Could computer systems develop aesthetic behaviour and judgement, particularly in the absence of qualia and emotional feeling? Some specific methods have been developed for the computational analysis of very specific kinds of art (e.g. (Dodgson, 2008)), but general AI art criticism and analysis has had little exploration.

This paper began with a discussion of the drivers for search-based art. Typically, as with other areas of AI, these drivers are very simple; a point in the search space is examined by an algorithm, and a numerical score or ranking is given. One technique which seems particularly amenable to this kind of work is the idea of richer search drivers (Krawiec et al., 2016), which provide not just a single figure but a richer description of the steps to take to explore the search space.

References

Blaise Agüera y Arcas. Art in the age of machine intelligence. *Arts*, 6(18), 2017. https://doi.org/10.3390/arts6040018.

S. Aupetit, V. Bordeau, N. Monmarche, M. Slimane, and G. Venturini. Interactive evolution of ant paintings. In *The 2003 Congress on Evolutionary Computation, 2003. CEC '03.*, volume 2, pages 1376–1383, 2003.

Angela Katherine Baker. Alba emoting: A safe, effective, and versatile technique for generating emotions in acting performance. Master's thesis, Brigham Young University, 2008. URL `https://scholarsarchive.byu.edu/cgi/viewcontent.cgi?article=2456&context=etd`.

Laura Barton. The man who rewrote his life. *The Guardian*, 15th September 2006. URL `https://www.theguardian.com/books/2006/sep/15/usa.world`.

Clive Bell. *Art*. Chatto and Windus, 1914.

Walter Benjamin. The work of art in the age of mechanical reproduction. In Hannah Arendt, editor, *Illuminations*, pages 214–218, 1969. translated by Harry Zohn, originally published 1935,.

Arnold Berleant. Environment as an aesthetic paradigm. *Dialectics and Humanism*, 15:95–106, 1988.

Jon Bird, Joe Faith, and Andy Webster. *Tabula Rasa*: A case study in evolutionary curation. In Stefano Cagnoni et al., editors, *Applications of Evolutionary Computing: Evoworkshops 2003*, volume 2611 of *Lecture Notes in Computer Science*, pages 981–995. Springer, 2003.

George David Birkhoff. *Aesthetic Measure*. Harvard University Press, 1933.

Margaret Boden. *The Creative Mind: Myths and Mechanisms*. Abacus, 1990.

Margaret Boden. Robot says: Whatever. *Aeon*, 2018. URL `https://aeon.co/amp/essays/the-robots-wont-take-over-because-they-couldnt-care-less`.

Margaret A. Boden. Creativity and artificial intelligence. *Artificial Intelligence*, 103 (1–2):347–356, 1998.

Paul M. Bodily and Dan Ventura. Explainability: An aesthetic for aesthetics in computational creative systems. In François Pachet et al., editors, *Proceedings of the 9th International Conference on Computational Creativity*, pages 153–160, 2018.

Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *Swarm Intelligence*. Oxford University Press, 1999.

Andreas Broeckmann. The machine artist as myth. *Arts*, 8(25), 2019. https://doi.org/10.3390/arts8010025.

Bruce G. Buchan. Creativity at the meta-level. *AI Magazine*, 22(3):13–28, 2001.

John Cage. *Atlas Eclipticalis*. Edition Peters, 1961.

Stefano Cagnoni et al., editors. *Applications of Evolutionary Computing: Evoworkshops 2003*, volume 2611 of *Lecture Notes in Computer Science*, 2003. Springer.

Erik Cambria, Amir Hussain, Catherine Havasi, and Chris Eckl. Common sense computing: From the society of mind to digital intuition and beyond. In Julian Fierrez et al., editors, *Biometric ID Management and Multimodal Communication*, pages 252–259. Springer, 2009.

Allen Carlson. Appreciation and the natural environment. *Journal of Aesthetics and Art Criticism*, 37:267–276, 1979.

Noel Carroll. On being moved by nature: Between religion and natural history. In Salim Kemal and Ivan Gaskell, editors, *Landscape, Natural Beauty and the Arts*. Cambridge University Press, 1995.

Noël Carroll. *Philosophy of Art: A Contemporary Introduction*. Routledge, 1999.

John Charnley, Alison Pease, and Simon Colton. On the notion of framing in computational creativity. In Mary Lou Maher et al., editors, *Proceedings of the International Conference on Computational Creativity*, pages 77–81, 2012.

Harold Cohen. The further exploits of AARON, painter. *Stanford Humanities Review*, 4(2):141–158, 1995.

R.G. Collingwood. *The Principles of Art*. Oxford University Press, 1938.

John Collomosse and Peter Hall. Salience-adaptive painterly rendering using genetic search. *International Journal on Artificial Intelligence Tools*, 15(4):551–576, 2006.

Simon Colton. *Automated Theory Formation in Pure Mathematics*. Springer, 2002.

Simon Colton. Automatic invention of fitness functions with application to scene generation. In Mario Giacobini et al., editors, *Applications of Evolutionary Computing*, volume 4974 of *Lecture Notes in Computer Science*, pages 381–391. Springer Berlin / Heidelberg, 2008a.

Simon Colton. Creativity versus the perception of creativity in computational systems. In *Proceedings of the AAAI Symposium on Creative Systems*, pages 14–20, 2008b.

Simon Colton. The painting fool: Stories from building an automated painter. In Jon McCormack and Mark d'Inverno, editors, *Computers and Creativity*, pages 3–38. Springer, 2012.

Simon Colton, Michel F. Valstar, and Maja Pantic. Emotionally aware automated portrait painting. In Sofia Tsekeridou et al., editors, *Proceedings of the Third International Conference on Digital Interactive Media in Entertainment and Arts, DIMEA 2008*, ACM International Conference Proceeding Series Volume 349, pages 304–311, 2008.

Simon Colton, Alison Pease, and Rob Saunders. Issues of authenticity in autonomously creative systems. In François Pachet et al., editors, *Proceedings of* *the 9th International Conference on Computational Creativity*, pages 272–279, 2018.

Arthur Danto. The artworld. *Journal of Philosophy*, 64(19):571–584, 1964.

Arthur Danto. *The Transfiguration of the Commonplace*. HArvard University Press, 1981.

E. den Heijer and A. E. Eiben. Evolving art using multiple aesthetic measures. In Cecilia Di Chio et al., editors, *Applications of Evolutionary Computation*, volume 6625 of *Lecture Notes in Computer Science*, pages 234–243. Springer Berlin / Heidelberg, 2011.

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

Ellen Dissanayake. *Homo Aestheticus: Where Art Comes From and Why*. University of Washington Press, 1997.

Neil A. Dodgson. Regularity and randomness in Bridget Riley's early op art. In Douglas W. Cunningham et al., editors, *Computational Aesthetics in Graphics, Visualization, and Imaging*. The Eurographics Association, 2008.

Alan Dorin and Kevin B. Korb. A new definition of creativity. In Kevin Korb, Marcus Randall, and Tim Hendtlass, editors, *Artificial Life: Borrowing from Biology*, pages 11–21, 2009.

Scott Draves. The electric sheep screen-saver: A case study in aesthetic evolution. In Franz Rothlauf et al., editors, *Applications of Evolutionary Computing*, volume 3449 of *Lecture Notes in Computer Science*, pages 458–467. Springer Berlin / Heidelberg, 2005.

Denis Dutton. *The Art Instinct: Beauty, Pleasure, and Human Evolution*. Oxford University Press, 2009.

Denis Dutton. Aesthetic universals. In Berys Gaut and Dominic McIver Lopes, editors, *The Routledge Companion to Aesthetics*, pages 267–278. Routledge, 2013.

Craig Dworkin and Kenneth Goldsmith, editors. *Against Expression*. Northwestern University Press, 2011.

Ahmed M. Elgammal, Bingchen Liu, Mohamed Elhoseiny, and Marian Mazzone. CAN: creative adversarial networks, generating "art" by learning about styles and deviating from style norms. In Ashok Goel, Anna Jordanous, and Alison Pease, editors, *Proceedings of the Eighth International Conference on Computational Creativie*, pages 96–103, 2017.

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

Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. A neural algorithm of artistic style. *CoRR*, abs/1508.06576, 2015. URL `http://arxiv.org/abs/1508.06576`.

Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2414–2423, 2016.

Berys Gaut and Dominic McIver Lopes, editors. *The Routledge Companion to Aesthetics*, 2013. Routledge.

Pablo Gervas. Computational approaches to storytelling and creativity. *AI Magazine*, 30(3):49–62, 2009.

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

Thalia Raquel Goldstein and Ellen Winner. A new lens on the development of social cognition: The study of acting. In Constance Milbrath and Cynthia Lightfoot, editors, *Art and Human Development*. Psychology Press (Taylor and Francis), 2012.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2017.

Ian J. Goodfellow et al. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, pages 2672–2680. MIT Press, 2014.

Gary Greenfield. Evolved ricochet compositions. In Mario Giacobini et al., editors, *Applications of Evolutionary Computing*, volume 5484 of *Lecture Notes in Computer Science*, pages 518–527. Springer Berlin / Heidelberg, 2009.

Gary Greenfield. Avoidance drawings evolved using virtual drawing robots. In Colin Johnson, Adrian Carballal, and João Correia, editors, *Evolutionary and Biologically Inspired Music, Sound, Art and Design*, pages 78–88, 2015.

Gary Greenfield. Evolved deflection drawings. In *2016 IEEE Congress on Evolutionary Computation (CEC)*, pages 4562–4569, 2016.

Gary Greenfield and Penousal Machado. Simulating artist and critic dynamics—an agent-based application of an evolutionary art system. In *IJCCI 2009 - Proceedings of the International Joint Conference on Computational Intelligence*, pages 190–197, 2009.

Emma Hart. Towards lifelong learning in optimisation algorithms. In *Proceedings of the 9th International Joint Conference on Computational Intelligence - Volume 1: IJCCI,*, pages 7–9. INSTICC, SciTePress, 2017. doi: 10.5220/0006810500010001.

Aaron Hertzmann. Can computers create art? *Arts*, 7(18), 2018. https://doi.org/10.3390/arts7020018.

Bert Hölldobler and Edward O. Wilson. *The Ants*. Springer, 1990.

Bipin Indurkhya. On the role of metaphor in creative cognition. In Dan Ventura et al., editors, *Proceedings of the International Conference on Computational Creativity: ICCC-X*, pages 51–59, 2010.

Colin G. Johnson. Connotation in computational creativity. *Cognitive Computation*, 4(3):280–291, 2012a.

Colin G. Johnson. Fitness in evolutionary art and music: What has been used and what could be used? In Penousal Machado, Juan Romero, and Adrian Carballal, editors, *Evolutionary and Biologically Inspired Music, Sound, Art and Design*, volume 7247 of *Lecture Notes in Computer Science*, pages 129–140. Springer Berlin / Heidelberg, 2012b.

Colin G. Johnson. Is it time for computational creativity to grow up and start being irresponsible? In Simon Colton et al., editors, *Procedings of the Fifth International Conference on Computational Creativity*, pages 263–267, 2014.

Colin G. Johnson. Fitness in evolutionary art and music: a taxonomy and future prospects. *International Journal of Arts and Technology*, 9(1), 2016.

Colin G. Johnson et al. Understanding aesthetics and fitness measures in evolutionary art systems. *Complexity*, 2019, 2019. Article ID 3495962.

Anna Jordanous. A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation*, 4(3):246–279, 2012.

Branden W. Joseph. *Experimentations: John Cage in Music, Art, and Architecture*. Bloomsbury, 2016.

Immanuel Kant. *Critique of the Power of Judgement*. Lagarde und Friedrich, 1790. translation by P. Guyer and A.W. Wood, Cambridge University Press, 2000.

Peter Kivy. *Authenticities: Philosophical Reflections on Musical Performance*. Cornell University Press, 1998.

Mario Klingemann. Quasimondo: Mario Klingemann, artist. Visited August 2019, 2019. URL `http://quasimondo.com`.

Richard Kostelanetz and John Cage. The aesthetics of john cage: A composite interview. *The Kenyon Review*, 9(4):102–130, 1987.

Krzysztof Krawiec, Jerry Swan, and Una-May O'Reilly. Behavioral program synthesis: Insights and prospects. In R. Riolo, B. Worzel, M. Kotanchek, and A. Kordon, editors, *Genetic Programming Theory and Practice XIII*, pages 169–183. Springer, 2016.

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

Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep painterly harmonization. *Computer Graphics Forum*, 37(4):95–106, 2018.

Penousal Machado and Amílcar Cardoso. All the truth about NEvAr. *Applied Intelligence*, 16(2):101–118, 2002.

Penousal Machado, Juan Romero, María Santos, Amílcar Cardoso, and Bill Manaris. Adaptive critics for evolutionary artists. In Günther Raidl et al., editors, *Applications of Evolutionary Computing*, volume 3005 of *Lecture Notes in Computer Science*, pages 437–446. Springer Berlin / Heidelberg, 2004.

Penousal Machado, Juan Romero, and Bill Manaris. Experiments in computational aesthetics. In Juan Romero and Penousal Machado, editors, *The Art of Artificial Evolution*, pages 381–415. Springer, 2008.

Penousal Machado, João Correia, and Juan Romero. Expression-based evolution of faces. In Penousal Machado, Juan Romero, and Adrian Carballal, editors, *Evolutionary and Biologically Inspired Music, Sound, Art and Design*, volume 7247 of *Lecture Notes in Computer Science*, pages 187–198. Springer Berlin / Heidelberg, 2012a.

Penousal Machado, Juan Romero, and Adrian Carballal, editors. *Evolutionary and Biologically Inspired Music, Sound, Art and Design*, volume 7247 of *Lecture Notes in Computer Science*, 2012b. Springer Berlin / Heidelberg.

Marian Mazzone and Ahmed Elgammal. Art, creativity, and the potential of artificial intelligence. *Arts*, 8(26), 2019. https://doi.org/10.3390/arts8010026.

Pamela McCorduck. *Aaron's Code*. W.H.Freeman & Co, 1991.

Jon McCormack. Eden: An evolutionary sonic ecosystem. In Josef Kelemen and Petr Sosík, editors, *Advances in Artificial Life, Proceedings of the 6th European Conference*, pages 133–142. Springer, 2001.

Jon McCormack. Open problems in evolutionary music and art. In Franz Rothlauf et al., editors, *Applications of Evolutionary Computing*, volume 3449 of *Lecture Notes in Computer Science*, pages 428–436. Springer Berlin / Heidelberg, 2005.

Jon McCormack. Creative ecosystems. In Jon McCormack and Mark d'Inverno, editors, *Computers and Creativity*, pages 39–60. Springer, 2012.

Jon McCormack and Mark d'Inverno, editors. *Computers and Creativity*, 2012. Springer.

Joanna Misztal and Bipin Indurkhya. Poetry generation system with an emotional personality. In Simon Colton et al., editors, *Proceedings of the 5th International Conference on Computational Creativity*, pages 72–81, 2014.

Thomas Mitchell. *Machine Learning*. McGraw-Hill, 1997.

Jemma Montagu. *The Surrealists: Revolutionaries in Art and Writing 1919–35*. Tate Publishing, 2002.

Leonel Moura. Robot art: An interview with Leonel Moura. *Arts*, 7(3):28, 2018.

Martin A. Nowak and David C. Krakauer. The evolution of language. *Proceedings of the National Academy of Sciences*, 96(14):8028–8033, 1999.

Obvious. Obvious, explained. medium.com, visited August 2019, 2018.

Obvious Collective. Obvious. Visited August 2019, 2019. URL `https://obvious-art.com`.

François Pachet et al., editors. *Proceedings of the 9th International Conference on Computational Creativity*, 2018.

Rafael Pérez y Pérez and Mike Sharples. MEXICA: A computer model of a cognitve account of creative writing. *Journal of Experimental and Theoretical Artifical Intelligence*, 13:119–139, 2001.

Rosalind Picard. *Affective Computing*. MIT Press, 1997.

Plato. *Cratylus. Parmenides. Greater Hippias. Lesser Hippias*. Harvard University Press, 1926. Loeb Classical Library Volume 167, Translated by Harold North Fowler.

Rafael Ramirez et al. A sequential covering evolutionary algorithm for expressive music performance. In *Proceedings of the 18th Conference on Innovative Applications of Artificial Intelligence - Volume 2*, pages 1830–1835, 2006. ISBN 978-1-57735-281-5.

Graeme Ritchie. Assessing creativity. In *Proceedings of the AISB Symposium on Artificial Intelligence and Creativity in Arts and Science*, pages 3—11. AISB Press, 2001.

Graeme Ritchie. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines*, 17:67–99, 2007.

Juan Romero and Penousal Machado, editors. *The Art of Artificial Evolution*, 2008. Springer.

Juan Romero, Penousal Machado, Antonio Santos, and Amilcar Cardoso. On the development of critics in evolutionary computation artists. In Stefano Cagnoni et al., editors, *Applications of Evolutionary Computing: Evoworkshops 2003*, volume 2611 of *Lecture Notes in Computer Science*, pages 559–569. Springer, 2003.

Juan Romero, Penousal Machado, and Antonino Santos. On the socialization of evolutionary art. In Mario Giacobini et al., editors, *Applications of Evolutionary Computing*, volume 5484 of *Lecture Notes in Computer Science*, pages 557–566. Springer Berlin / Heidelberg, 2009.

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

Robert Rowe. *Interactive Music Systems*. MIT Press, 1997.

Yuriko Saito. Appreciating nature on its own terms. *Environmental Ethics*, 20: 135–149, 1998.

Dario Sanfilippo. Turning perturbation into emergent sound, and sound into perturbation. *Interference: Journal of Audio Culture*, 3, 2013. URL `http://www.interferencejournal.org/turning-perturbation-into-emergent-sound/`.

Ruth L. Saw. *Aesthetics: An Introduction*. Macmillan, 1972.

Nadja Sayej. This computer painting program has feelings. *Vice*, November 2013. URL `https://www.vice.com/en_us/article/4w7vvg/the-computer-painting-program-with-feelings`.

Juergen Schmidhuber. Low-complexity art. *Leonardo*, 30(2):97–103, 1997.

Mike Sharples. *How We Write: Writing as Creative Design*. Routledge, 1998.

Anne Sheppard. *Aesthetics: An Introduction to the Philosophy of Art*. Oxford University Press, 1987.

Daniel L. Silver, Qiang Yang, and Lianghao Li. Lifelong machine learning systems: Beyond learning algorithms. In *AAAI Spring Symposium: Lifelong Machine Learning*, pages 49–55. AAAI Press, 2013.

Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *Workshop at International Conference on Learning Representations*, 2014.

Aaron Sloman. *The Computer Revolution in Philosophy*. The Harvester Press, 1978.

Emily L. Spratt. Dream formulations and deep neural networks: Humanistic themes in the iconology of the machine-learned image. *Kunsttexte*, 4, 2017.

Joseph Stromberg. These abstract portraits were painted by an artificial intelligence program. *smithsonian.com*, November 2013. URL `https://www.smithsonianmag.com/science-nature/these-abstract-portraits-were-painted-by-an-artificial-intelligence-program-180947590/`.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018. Second Edition.

Susan Leith Taylor. *Actor Training and Emotions: Finding a Balance*. PhD thesis, Edith Cowan University, 2006. URL `http://ro.ecu.edu.au/theses/1804`.

Leo Tolstoy. *What is Art?* Oxford University Press, 1897. English Translation by Aylmer Maude, 1930.

Christopher Turner. The deliberate accident in art. *Tate Etc.*, 21, Spring 2011.

David Vanderhaeghe and John Collomosse. Stroke based painterly rendering. In Paul Rosin and John Collomosse, editors, *Image and Video-based Artistic Stylisation*. Springer, 2012.

William Vaughan. *Romanticism and Art*. Thames and Hudson, second edition, 1994.

Tony Veale. *Exploding the Creativity Myth: The Computational Foundations of Linguistic Creativity*. Bloomsbury Academic, 2012.

Tony Veale and Amílcar Cardoso, editors. *Computational Creativity ;The Philosophy and Engineering of Autonomously Creative Systems*, 2019. Springer.

Dan Ventura. Mere generation: Essential barometer or dated concept? In François Pachet et al., editors, *Proceedings of the Seventh International Conference on Computational Creativity*, pages 17–24, 2016.

Eirini Vouliouri. Merging aesthetics with functionality: An interactive genetic algorithm based on the principle of weighted mutation. In Cecilia Di Chio et al., editors, *Applications of Evolutionary Computation*, volume 6625 of *Lecture Notes in Computer Science*, pages 424–433. Springer Berlin / Heidelberg, 2011.

Geraint A. Wiggins. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems*, 19:449–458, 2006a.

Geraint A. Wiggins. Searching for computational creativity. *New Generation Computing*, 24:209–222, 2006b.

Ludwig Wittgenstein. *Philosophical Investigations*. Blackwell, 1953.