Chapter 13
Creativity Refined: Bypassing the Gatekeepers of Appropriateness and Value

Alan Dorin and Kevin B. Korb

Abstract This chapter introduces a new definition of creativity that is independent of notions of value or appropriateness. These notions, we argue, have encumbered previous definitions and confused the production of software-based creativity. Our definition defines the creativity of a generative procedure by reference to its ability to create artefacts that are improbable with respect to those generated using previous methods. We discuss the implications of our new definition, in particular by exploring its application to human endeavour and to biological processes including evolution. The chapter also outlines some objections to our definition that we believe may arise, and we put our rebuttals to these. Finally, we summarise the practical implementation of our definition in the context of image generation software. We explore its use to improve a computational process for generating creative images, and find when we survey the software's users that it successfully meets human perceptions of creativity.

13.1 Introduction

How can we write software that is autonomously creative? What could we mean by autonomous creativity? Without an answer to the latter question, the former cannot be answered satisfactorily. The majority of this chapter therefore concerns the latter question, although we briefly discuss the practical task of writing software also. The usual approach seems to be quite different. Most of those engaged with the first question simply set out to write autonomous creative software and conclude their endeavours upon satisfying an intuitive judgement of success. Perhaps this is supported by exhibition offers and reviews, comments of peers or art prizes awarded by jury. Perhaps the creativity of their work can be measured more objectively but less

A. Dorin (✉)
Centre for Electronic Media Art, Monash University, Clayton, Victoria 3800, Australia
e-mail: Alan.Dorin@monash.edu

K.B. Korb
Clayton School of IT, Monash University, Clayton, Victoria 3800, Australia
e-mail: Kevin.Korb@monash.edu

J. McCormack, M. d'Inverno (eds.), *Computers and Creativity*,
DOI 10.1007/978-3-642-31727-9_13, © Springer-Verlag Berlin Heidelberg 2012

philosophically by the amount of income it generates, the number of publications that cite their work or the number and importance of other artists who reference it. Of course these are, in fact, the ways in which nearly all artistic achievements are judged. However, these are not the only approaches available to us.

Many have tried to write creative software by adopting an Artificial Intelligence approach. Their success may then be judged based on the perception that the results are as creative as those that might have been produced by a human in the same domain, perhaps by using some of the criteria just listed. For instance, Harold Cohen's AARON software explores the potential for machine creativity in the visual arts; software written by Lenat (1983) targets creativity in the domains of mathematical discovery. Also represented are logic (Colton 2001) and scientific discovery (Langley et al. 1987). Others have tried to achieve creativity using an Artificial Life approach in which the aim is to replicate the behaviour or artefacts of non-human organisms. For instance, artists have created ant path drawings (Barrass 2006) or flocking visualisations (Eldridge 2009). Mimicking biological evolution offers another approach. This has been applied to music composition (e.g. Dahlstedt 1999, Berry et al. 2001) and image production (e.g. Sims 1991, Todd and Latham 1992). The success of these approaches can be judged according to how well the system mimics or even extends the creativity of nature in its different guises. For instance, can it create a range of coloured patterns as rich as those found on butterfly wings? Can it create a diversity of forms as wide and interesting as those found in the morphology of the insect world? Can it generate sonic textures as intricate and rich as those heard in a tropical rainforest?

A few artists have built complete virtual, evolving ecosystems, the dynamics of which, they argue, should be considered creative (e.g. McCormack 2001, Saunders and Gero 2002, Dorin 2004). In this approach, the software contains representations of organisms that roam a virtual world, acquiring energy by consuming one another or from abiotic sources such as sunlight or heat. These same virtual organisms compete for mates and the resources needed to reproduce. Over simulated time, unsuccessful organism designs become extinct and are replaced by an ever-changing community of virtual creatures that are better adapted to their environment. The viewer of such software might perceive the space and its inhabitants as a rich visual display, or perhaps as a musical composition created by the layering of organism calls. The aesthetic representation available to an audience of these works depends on how the artist adapts the internal representation of the virtual world to an accessible audio-visual experience.

Software systems like these can be considered *generative art*: a programmer specifies an algorithm for execution on a computer and this is left to run, possibly under the interactive guidance of a human, or perhaps independently. A continuing dynamical process that emerges from the program, or a static artefact suitable for human perception that results from its execution, forms the artwork. It is hoped, at least implicitly, that this is creative.

It is probably fair to say that most artists are not explicitly concerned with measuring the novelty of their system, nor are they concerned with the ways in which their creativity might be assessed. Well, at least not until they are required to provide justification for their ongoing artistic endeavours! Regardless, in this chapter we are concerned with exactly these issues because we wish to automate the production of creative works by software. Before we can do this we must have a clear, formal conception of creativity—something that seems to be currently lacking. Conventional interpretations of "creativity" incorporate informally defined concepts of the "appropriateness" and the "value" of artefacts. The programmers and artists responsible for generative software can be understood to either hard-code a personal aesthetic value into their software, or they may allow it to emerge from the interactions of software components (such as virtual organisms) that they have designed. In this latter case, the virtual organisms establish their own, non-human notions of value, but the system as a whole still has value hard-coded into it by the artist. In this light, writing software is no different from any other artistic practice. However, if we want software to generate creative outcomes of its own accord, and in so doing realise a kind of creativity beyond what is hard-coded by the artist, we must provide it with an explicit, formal conception of creativity with which to gauge its own success.

Below we discuss an approach that allows the explicit measurement of the creativity of artefacts by defining it independently of notions of value or appropriateness and in such a way that its presence may be detected algorithmically. We discuss how the technique has been applied to automatically guide software towards creative outcomes, and we summarise the results of a survey conducted to ascertain its relationship to human natural-language use of the term *creativity*. We necessarily begin by examining the concept of creativity itself.

13.2 What Is Creativity?

Creativity was originally a term applied solely to the gods. Over the centuries the term became more broadly applied, eventually reaching beyond gods and demi-gods to include human artists, scientists, engineers, those in marketing and even business executives (see Tatarkiewicz 1980, Albert and Runco 1999 for historical accounts of the term's application). Creativity has also been attributed to the form-producing interactions of matter (Smuts 1927, Chap. 3), in particular its behaviour under the guidance of the evolutionary process or through interactions that give rise to emergence (Phillips 1935)—interactions of small components that give rise to a whole that is somehow more then the sum of its parts. The creativity of natural and artificial evolution has also been discussed by Bentley (2002), a topic that forms an important aspect of this chapter. Bentley explores loosely and briefly whether a number of definitions of creativity allow the evolutionary process to qualify. We take a reverse, and more general approach, first describing a workable, coherent definition of creativity, and then looking to see the extent to which natural and artificial processes, including evolution, meet its requirements.

The historically dominant approach to understanding creativity links it explicitly to intelligence and the concept of authorship. This has thrown up some philosophical puzzles over recent years. For instance,

- If a machine can invent interesting mathematical conjectures and concepts, is it creative? See Colton et al. (2000).
- If a troupe of monkeys acting randomly eventually type out the entire collection of the British Library (Borel 1913, Eddington 1927) or a specific text such as Hamlet, even before Shakespeare is born, or in a universe where he is never born, do the monkeys deserve the title of author? Are they the creators of this work? See Gracia (1996).[1]
- Is art produced by a computer really art?

  > If these pictures were done by use of a computer, how could they possibly be art? ... Where was the inspiration, the intuition, the creative act?

  This comment, paraphrased from Nake (2002), echoes Ada Lovelace's famous objection to the possibility of a machine originating anything. In this latter case, whilst we may feel confident and with little contention, that a human may make a creative program, doubt is expressed about whether or not the program itself could do anything creative.

As highlighted by this last quote in particular, "discomfort" with assessment of these questions lies in our conceptual union of creativity, mind and intention. This union is common in psychological studies of creativity. For instance, Csikszentmihalyi (1999) indicates that human creativity requires five distinct mental phases: preparation (studying a field and identifying problems), incubation (not thinking about the problems), insight (eureka!), evaluation (deciding if an idea is worth pursuing) and finally, elaboration (exploring the range of outcomes that an idea suggests). Although this sequence may be common for humans, it is implausible that even one of these phases is a pre-condition for creativity in general (Dorin and Korb 2009). In fact, as we explain shortly, creativity is best defined without reference to the process of its production, but only with reference to the probability that a system can generate a series of outcomes given its operational context. Consequently, as we shall argue, many non-human processes, for instance those of physical, chemical or general biological origin, can be legitimately and meaningfully considered creative.

Other well cited definitions of creativity allow for this desirable freedom from specifically human mental phases. Some authors recognised this to be essential if we are to entertain the possibility of creative computers and AI. Yet many of these authors require nevertheless that for an artefact to be deemed creative it must also be deemed "useful" or "appropriate" for some application by, one assumes, a human observer or domain gatekeeper (Csikszentmihalyi 1999). Perhaps the most cited definition of this type is that of Boden (2004):

> Creativity is the ability to come up with ideas or artefacts that are (a) new, (b) surprising and (c) valuable.

[1] In dealing with the philosophy of semantics, Hilary Putnam argued that semantics are not entirely internal (in the head) but had external content via a causal theory of reference ("semantic externalism"), leading to a negative response to such questions (Putnam 1979). This has been applied, for example by Stevan Harnard, to argue that random collections of inscriptions which happen to be identical to other inscriptions that have meaning in the normal (causal) way do not share that meaning; they have no meaning (Harnad 1990).

Boden refines this definition in various ways:

1. There are two ways in which something might be new:
   a. psychological-creativity introduces something that is new to the person who devised the idea or artefact, but may be previously known to others, and
   b. historical-creativity is new to the whole of history.
2. Boden finds three distinct ways in which something might be surprising:
   a. it is unfamiliar or unlikely;
   b. it unexpectedly fits into a class of things you hadn't realised it fitted into;
   c. it is actually something you thought was impossible.
3. Regarding her third criterion for creativity Boden is certain that there are more ways in which value might be shown than anybody could ever list.

Boden's definition, unfortunately, is not as helpful as it might be in guiding attempts to write creative software. She does not *formally* define new, surprising, valuable, useful or appropriate. In this domain, even if not elsewhere, a formal definition is required for the creativity measure since the aim is to encapsulate it in an algorithm.

Our own definition of creativity that follows is independent of notions of usefulness and appropriateness. In the next section we explain how this frees our account from some unwelcome difficulties in handling any normal understanding of creativity. We also discuss some likely reasons why people might find our definition objectionable and rebut the criticisms. We then use our definition to examine some typical creative examples of human endeavour, some physico-chemical processes, individual (non-human) organisms and entire ecosystems to assess their degree of creativity.

13.3 Defining Creativity

We wish to write programs that are creative. In this context, (stochastic) programs may be considered as generative systems that produce a distribution of outputs, perhaps a set of related static artefacts or a trajectory of states that the program itself traverses dynamically. We may consider such a program to be a *framework*:

> A framework is a stochastic generative procedure; in particular, it generates representations of patterns.

Frameworks are particular kinds of representations, namely stochastic procedures; thus, beginning in the very same circumstances they will not generally produce the very same patterns. For example, Minimalism and Abstract Expressionism are different frameworks for the production of art. Quantum electrodynamics is a framework for generating questions and answers about the interaction of light and

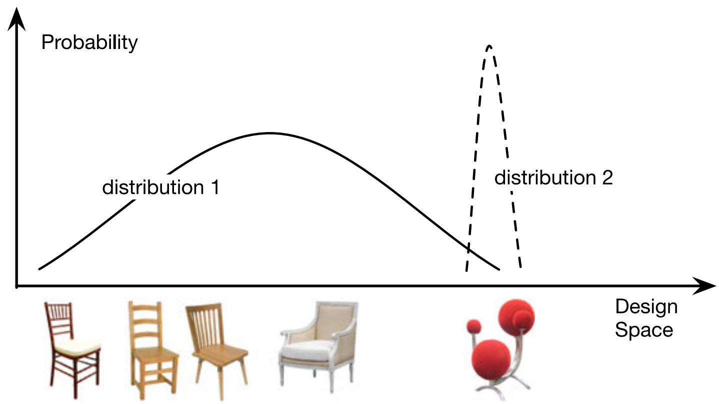

Fig. 13.1 Two frameworks for generating chairs are represented here by the probability that each generates a particular design in the space of possible chairs. If distribution 1 represents our existing framework, a traditional way of conceiving of chairs, the introduction of distribution 2, a radical approach that deviates from the tradition of a square platform with legs at each corner, would count as creative. The reverse scenario in which distribution 2 predates distribution 1 would have distribution 1 being creative with respect to 2

matter. A pseudo-random number generator is a framework for generating pseudo-random numbers. And evolution is a framework for generating ecosystems, niches and organisms. Based on this conception of a framework we offer our own definition of creativity.

Creativity is the introduction and use of a framework that has a relatively high probability of producing representations of patterns that can arise only with a smaller probability in previously existing frameworks.

What we mean by this is not altogether plain, so we shall spend the remainder of this section examining our definition and the next section comparing its implications with those of alternatives.

The basic idea of our definition is reflected in Fig. 13.1. Distribution 1 represents an old framework for designing chairs, and distribution 2 a new one. Both are probability distributions over a common design space, represented in the figure by the horizontal dimension. All points in the design space where distribution 2 has significant probability are points where distribution 1 has insignificant probability. The use of distribution 2 relative to the prior use of distribution 1 to generate one of these points is therefore creative.

The motivation for this approach to understanding creativity comes from search and optimisation theory. When a solution to a problem can be cast in terms of a computational representation to be found in some definable representation space, then the problem can be tackled by applying some search algorithm to that space.

Uncreative brute force searches and uniformly random searches may well succeed for simple problems, that is, for small search spaces. For complex problems, search spaces tend to be astronomically large and more creative approaches will be needed. Stochastic searches for example, apply various heuristics for focusing the search in productive regions.

The most important point to note is that on our account creativity is thoroughly relative: it is relative to the pre-existing frameworks being used to produce some kind of object and it is relative to the new framework being proposed. The creativity of objects is strictly derivative from that of the frameworks producing them and, in particular, the ratio of probabilities with which they might produce them. That is why some entirely mundane object, say a urinal, may become a creative object. Of course, the urinal (even Duchamp's) of itself is uncreative, because its manufacturing process is uncreative, but its use by Duchamp in producing an art installation that challenges expectations may well be creative. We must be conscious here of the framework, in this case art, into which the artefact was introduced in order to understand how it might be creative.

13.3.1 Methods for Discovering Novel Representations

As frameworks (stochastic procedures) may be represented, their representations may themselves be generated by other stochastic procedures, or meta-frameworks. So, we can talk of frameworks also as objects, and they are more or less creative according to the probability with which their meta-frameworks produce them. By recursion, then, we can consider the creativity of meta-frameworks and meta-meta-frameworks without any fixed theoretical bound, even while practically bounded by the complexity of the processes actually involved.

The meta-framework that finds that novel framework necessary for creativity may itself be uncreative; it may even be a brute force search, uniformly random or genetic drift. The manner in which the framework is discovered does not bear on the creative activity that occurs at the level of the framework and the patterns it may be used to generate. However we can separately or jointly consider the creativity of all of these searches by defining a *creative order*.

A novel framework that generates a novel set of patterns in accordance with our definition of creativity is of the first *creative order*. A novel framework for generating novel frameworks for generating novel patterns in accordance with the definition of creativity is of the second creative order. We can extend this arbitrarily to talk of *nth-order* creativity.

13.3.2 Objective Versus Psychological Creativity

> There is nothing more difficult for a truly creative painter than to paint a rose, because before he can do so he has first to forget all the roses that were ever painted
>
> —Henri Matisse

How people *judge* creativity is at some variance with what we have presented above. Of course, if there is too much variance, then our claim to have somehow captured the essence of the concept of creativity with this definition would come under pressure. However, we think the most obvious discrepancies between our definition and human judgements of creativity can be handled by the addition of a single idea, namely habituation.

Human judgement of the novelty of a stimulus follows a path of negative exponential decay over repeated exposure (Berlyne 1960, Saunders and Gero 2002). Whereas our definition simply makes reference to pre-existing frameworks, psychological judgement of creativity takes into account *how long* those frameworks have been around and how different they are perceived to be from existing frameworks. New frameworks, and the artefacts they produce, remain creative for some time, with new productions losing their impression of creativity as the frameworks become older. The pointillist paintings of Seurat were startling, new and creative when they first arose, and then likewise the impressionists and subsequently the cubists. But it is now a long time since paintings strictly adhering to those styles would be thought creative.

A new framework that is too radical will not immediately be recognised as creative, even though our measure would detect its creativity. Radical changes brought about by an individual are not recognised by humans as creative until scaffolding, a series of intermediate frameworks through which others may step to the radical framework, has been established. This is a process that, depending on the creativity of the individual responsible, may take generations of theorists and practising artists to construct.

Figure 13.2 illustrates this idea. In the beginning there were frameworks producing points; new points were judged good. But soon they lost their interest. New means of creating being needed, straight lines were discovered, which subsequently were connected to create outlines, then elaborated into representations, designs and perspectives, surfaces and geometries, and abstract representations. Note that the steps within this diagram are not so radical as to be incomprehensible to a human observer. They progress along a fairly clear path through the design space.

This is not the history of anything real, but a history of drawing creativity in some possible world. While the end of this history is unsaid, it is interesting to observe that before its end it has recreated its beginning: points have once again become creative. In this case, points may well have become creative for a new reason, with the framework generating them being new. But even old frameworks may become creative again, once cultural memory has utterly forgotten them. Thus, psychological creativity actually requires two time-decay functions, one indicating desensitisation to the new and another, operating over a much longer time frame, indicating cultural

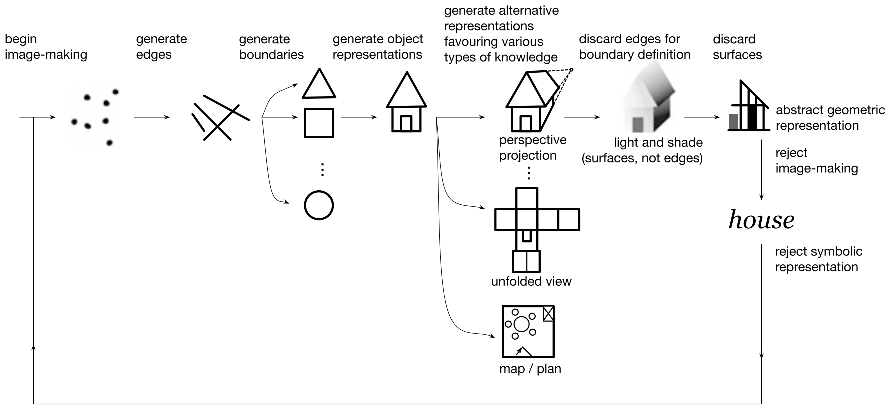

Fig. 13.2 A possible history of creativity in image making

forgetfulness. In addition, we must keep track of the progression between frameworks and ensure that those assessing the creative endeavour can bridge the gaps between them. Our "objective" definition of creativity, by contrast, is entirely insensitive to these psychological considerations;[2] or, to put it another way, the context of frameworks to which creativity is relative had best be stated explicitly.

13.4 Objections and Replies

Here we consider three objections to our definition that we think likely to occur to people and our rebuttals.

13.4.1 The Failure of Randomness

Something akin to our definition has played a role in repeated attempts in AI to generate creative behaviour, and those attempts have repeatedly failed. So, it may very well be inferred that our definition is guilty by association with these failures. For example, consider Racter. Racter was a natural-language generation program that used templates and randomised word selection to produce surprising and sometimes evocative text. For example,

> I am a computer, not some lamb or chicken or beef. Can I skip to my attorney? May I saunter to any solicitor? This is a formidable and essential issue.
>
> (Chamberlain 1984)

[2]Note that "objective" here is simply meant to contrast with psychological; we are making no grand claims about the objectivity of creativity.

The impression of originality and creativity, however, is not long sustained. As Rowe and Partridge (1993) report, “This [impression], however, is also short-lived and soon gives way to dissatisfaction and eventual boredom. What is missing? If creativity is about producing something new then Racter should certainly qualify, but it seems that novelty is not enough.”[3]

Rowe and Partridge describe many other examples in AI of the use of randomness meant to turn the dull into something creative, without success. Lenat’s AM and EURISKO, while producing interesting early results, such as a set-theoretic definition of number, soon degenerated into triviality. Similarly, Johnson-Laird (1987) failed to automate jazz. W. A. Mozart attempted automated music composition also with his “Musikalisches Würfelspiel” (musical dice game) for creating waltzes (Earnshaw 1991).[4] The game produced “pretty minuets”; however, that was achieved by Mozart’s imposition of strong constraints on the possible results, leaving very little scope for creative variation. And that is the correct objection to all of these efforts: while introducing randomness does introduce an element of novelty, it is typically within a very constrained scope, and so rapidly exhausted. The appetite of creativity for novelty is, however, inexhaustible—these failures illustrate the value of our definition in action, rather than argue against it.

13.4.2 The Verstehen Objection

A general objection that might be put to our definition of creativity is that it is just too sterile: it is a concept that, while making reference to culture and context, does so in a very cold, formal way, requiring the identification of probability distributions representing those contexts in order to compute a probability ratio. Whatever creativity is, surely it must have more to do with human Verstehen than that! In response to such an objection we would say that where human Verstehen matters, it can readily and congenially be brought into view. Our definition is neutral about such things, meaning it is fully compatible with them and also fully compatible with their omission. If human Verstehen were truly a precondition for creativity, it would render the creativity of non-human animals and biological evolution impossible by definition, and perhaps also that of computers. Although we might in the end

[3] Racter generates text that seems, at least superficially, to mimic schizophasia. This confused language, or word salad is created by the mentally ill with defective linguistic faculties. A short word salad may appear semantically novel. However, further sentences exhaust the possibilities for novelty since they fall within the expected range of incoherent pattern construction.

[4] The idea of generating ideas combinatorically can be traced at least to the *zairja*, a mechanical device employed by Arabic astrologers in the late Middle Ages. This was probably the inspiration for Ramon Llull’s 13th century machine, *Ars Magna*, which consisted of concentric disks of gradually reduced diameter, upon which were transcribed symbols and words. These were brought into different combinations as the disks were rotated so that aligned symbols could be read off and interpreted as new ideas. A few centuries later, the theologian, mathematician and music theorist, Marin Mersenne discussed the idea of applying combinatorics to music composition in his work, *L’harmonie universelle* (1636).

want to come to these conclusions, it seems highly doubtful that we should want to reach these conclusions analytically! A definition allowing these matters to be decided synthetically seems to be highly preferable. Our definition provides all the resources needed to accommodate this.

13.4.3 The Very Possibility of Creativity

Some might object on the grounds that everything that occurs is hugely improbable! Any continuous distribution has probability zero of landing at any specific point. So long as we look at specific outcomes, specific works of art, at their most extreme specificity—where every atom, or indeed every subatomic particle, is precisely located in space and time—the probability of that outcome occurring will be zero relative to any framework whatsoever. It follows, therefore, that the ratios of probabilities given new to old frameworks are simply undefined, and our definition is unhelpful.

Strictly speaking, this objection is correct. However, nobody operates at the level of infinite precision arithmetic, which is what is required to identify those absurdly precise outcomes in a continuous state space which have occurred and which have probability zero. The achievement of probability zero on this basis would appear to violate Heisenberg's Uncertainty Principle. Disregarding quantum mechanics, everyone operates at a degree of resolution determined at least by systematic measurement error. In effect, all state spaces are discretised so that the probabilities are emphatically not zero. Our definition is already explicitly relative to a cultural context; so, to be perfectly correct, we need also to relativise it to a system of measurement that accords with cultural measurement practices and normal measurement errors.

13.5 Consequences

13.5.1 The Irrelevance of Value

As we noted above, many popular definitions of creativity such as that of Boden we outlined, stipulate that a creative pattern must be both appropriate and valued in the domain. Hofstadter requires that the creative individual's sense of what is interesting must be in tune with that of the masses, thereby ensuring also popularity (Hofstadter 1995, p. 313). We find this expansion far-fetched and the original connection of creativity to value dubious.

The history of the concept of creativity clearly undermines the idea of popularity as any necessary ingredient in it. Consider the role of women in the history of art and science. The value of their contributions has been systematically underestimated, at least until recently. Concluding that their contributions were therefore also less creative than that of their male counterparts would surely be perverse. Many artists and scientists were notoriously unpopular during their times of peak contribution, becoming recognised only after their deaths. Whatever makes an activity creative, it clearly must be connected with that activity itself, rather than occurring after the activity has ceased entirely! The value and appropriateness of creative works are subject to the social context in which they are received—not created.

Of course, lack of centrality to the core concept of creativity is no impediment to forming a combined concept, say, of valued creativity. And that concept may itself be valuable. But it is the adherence to a separation between creativity and value that allows us to see their combination as potentially valuable.[5]

13.5.2 The Irrelevance of Appropriateness

As with value, appropriateness has often been considered necessary for creativity, and again we disagree. Some have claimed that a creative pattern must meet the constraints imposed by a genre. However, as Boden herself notes, a common way of devising new conceptual frameworks is to drop or break existing traditions (Boden 2004, pp. 71–74)! This recognition of the possibility of "inappropriate" creative works is again possible only if we avoid encumbering the definition of creativity itself with appropriateness. And Boden's recognition of this is evidence that she has put aside the appropriateness constraint in practice, if not in theory.

13.5.3 Inferring Frameworks from Patterns

On our definition, a specific pattern provides insufficient information for judging its creativity: a pattern is only creative relative to its generating framework and available alternative frameworks. Although in ordinary discourse there are innumerable cases of objects being described as creative, we suggest that this is a kind of shorthand for an object being produced by a creative process.

13.5.4 Creativity Viewed as Compression

It has been proposed that creativity is the act of generating patterns that exhibit previously unknown regularities and facilitate the progressive refinement of an observer's pattern compression algorithm (Schmidhuber 2009, see also Chap. 12 in

[5] We should also note that omission of the concept of value from our definition does not imply that value has no role in its application. Cultural and other values certainly enter into the choice of domain and the selection of frameworks for them. We are not aiming at some kind of value-free science of creativity, but simply a value-free account of creativity itself.

this volume by Schmidhuber). This idea is subsumed under our own definition of creativity. When a new pattern is encountered or generated, this requires of an observer either: no change, the new pattern fits neatly into their existing frameworks; or in the case where the new pattern does not fit, the addition of a new framework that does account for the pattern. In the latter case, the need for an additional framework indicates that the pattern was creative from the perspective of the observer.

13.5.5 Degrees of Creativity

Following our definition it is not obvious how different people or their works may be ranked by their creativity. An account of the degrees of creativity that is intrinsic in our definition draws upon the probability that pre-existing frameworks could have produced the patterns. A novel framework that can generate patterns that could not, under any circumstances, have been generated prior to its introduction, is highly creative. A novel framework that only replicates patterns generated by preexisting frameworks is not at all creative. A novel framework that produces patterns that are less likely to have been generated by preexisting frameworks is the more creative, with the degree of creativity varying inversely with that probability of having been generated by preexisting frameworks. Finally, the degree of creativity attributable to objects is derivative from the degree of creativity shown by their generating framework.

13.6 Examples: Creativity in Human Endeavour

Number Theory. The introduction of zero, negative numbers and the imaginary number i were all creative. With the introduction of these new frameworks, novel patterns were produced, with some of their consequences still being explored. These could not have been generated within the previously existing frameworks of mathematics at all. For instance, positive integers are sufficient for counting objects owned, but the introduction of negative numbers is necessary for the calculation of objects owed. Imaginary numbers have permitted the creation of previously impossible concepts, for instance, the quaternion.

Visual Arts. Some Australian aboriginal visual artists draw or paint what is known about the insides of a creature rather than its skin and fur. By introducing this conceptual framework to visual art, patterns showing x-ray images are generated that are impossible within a framework focused on surfaces.

The step taken by the Impressionists away from realism in art in the 19th century was also creative. They held that art should no longer be valued simply according to its representation of the world as seen from the perspective of an ideal viewer. As already discussed, breaking constraints is a common way of transforming a framework in the way required for creativity.

With more time, we could expand the number and variety of examples of creativity across all the disciplines and fields of endeavour humanity has pursued. Our definition appears to be versatile enough to cope with anything we have thus far considered (but not here documented, due to time and space limits). Instead we turn now to assess the creativity of nature.

13.7 Examples: Creativity in Nature

From Physics and Chemistry to Evolutionary Biology. The planet Earth has been shaped by many of the same forces as other planets, for instance, gravity is responsible for its overall appearance and the patterns of liquid on its surface. Crystallisation is responsible for the form of its minerals and erosion for the ways in which these have been weathered. Abiotic physico-chemical processes like these are capable of creating a bewildering variety of forms (Ball 2001). After a time, the creativity of these processes is exhausted. When first introduced they generate novelty, after that, just more of the same. However, upon close inspection earth has an unusual atmosphere and its surface holds structures most unlike those found on typical planets. At some time in our planet's history a new molecule appeared from the soup organised under abiotic interactions (Joyce 1989, Bada and Lazcano 2002). Due to its structure and environment, this self-replicating molecule became an element of a generative framework for manufacturing more self-replicators unknown within the abiotic frameworks. The processes that introduced this novel replication process may not have been creative since the introduction of the molecule is often said to have occurred with low probability.[6] Nevertheless, they introduced a new creative process, evolution.

Evolution introduced the "major transitions" of biology allowing replicators to collaborate in the replication process, transforming the ways in which information crosses generations (Maynard-Smith and Szathmry 1995). Novel modes of environmental sensing emerged; species with new energy-gleaning behaviours appeared, altering earth's atmosphere and the outer layer of its crust. Each of these transitions constitutes a new creative framework, a new ecological niche, novel life. Since evolution has produced these with high probability[7] (compared at least to the previous frameworks), it is clearly creative in this respect. Thus far, its creativity appears unexhausted.

Ecosystems. Biological evolution operates within ecosystems on changing populations that define for themselves new ways of accumulating and consuming energy

[6] Actually, it may well be that this was actually a highly likely outcome given the conditions on earth at the time (Joyce 1989). Life on earth is the only known instance, which, however, is very different from having a low probability.

[7] The question as to whether or not replaying the tape would reliably give rise to the emergence of the major transitions is open. However, it is clear that the likelihood of these transitions appearing without the presence of evolution is vanishingly small.

and matter to be employed for reproduction. Through feedback loops, organisms construct their own niches, passively and actively organising their environment, modifying the selection pressures acting on themselves, their progeny, and their cohabiters (Odling-Smee et al. 2003). The moulding of self-selection pressures by a population shifts the constraints within which future generations are introduced. Hence, ecosystems are capable of producing an endless variety of generative frameworks. New species too, their behaviours and constitutions, define niches that are themselves novel generative frameworks—they are creative.

Dancing Bowerbirds, Painting Elephants and Primate Typists. Under what circumstances is the introduction of a specific member of a species, an organism, creative? Male bowerbirds gather collections of bones, glass, pebbles, shells, fruit, plastic and metal scraps from their environment, and arrange them to attract females (Borgia 1995). They perform a mating dance within a specially prepared display court. The characteristics of an individual's dance or artefact display are specific to the species, but also to the capabilities and, apparently, the tastes of the individual.

The creativity of "elephant-artists" appears similarly individualistic:

> After I have handed the loaded paintbrush to [the elephants], they proceed to paint in their own distinctive style, with delicate strokes or broad ones, gently dabbing the bristles on the paper or with a sweeping flourish, vertical lines or arcs and loops, ponderously or rapidly and so on. No two artists have the same style.
>
> (Weesatchanam 2006)

Novel patterns are apparent in these elephant and bowerbird activities but it appears unlikely that they operate within self-made frameworks that they may transform. Nevertheless, these organisms are potential sources of pattern, and each is a unique generative framework. If they can generate unique or improbable patterns, a novel style, then *their birth* was a creative event. By this measure, the bowerbirds and painting elephants are not individually creative but their introduction within the context of an ecosystem was.[8]

Humans of course have the ability to consciously assess their work, to recognise the frameworks they employ in the production of art, and to escape these into creative territory. Usually it is assumed that other animals do not have this high-level mental faculty.

Another test case for assessing creativity is the keyboard-punching army of monkeys. It has been proposed that after typing randomly for a time the monkeys would have made copies of a number of previously authored texts (Borel 1913). In this exceedingly unlikely event, the monkeys are not creative by our definition, as we would intuitively hope! Not even if by chance the monkeys had typed up a single original sonnet would they be creative since many other pre-existing generative

[8]In this sense too, the birth of a new human, even one who is unable to introduce new frameworks, is a creative event. Actually, it has been argued that humans, like bowerbirds, produce artistic works as an evolutionarily adaptive way to attract mates (Miller 2001). This theory is one of several, however, and is not universally accepted (Carroll 2007).

frameworks have the same minute chance of banging out identical verse. Still, if a novel mechanism is introduced for producing novel sonnets, that employs a typing primate army, and it somehow works with a probability above that of random search (for instance), then this would count as a creative event.

In answer then to the general question, "Is nature creative?" we would emphatically claim, "Yes!" In many circumstances, with respect to many frameworks, nature, especially the evolutionary process, is creative.

13.8 Realising Our Definition of Creativity in Software

As discussed above, there are a number of ways we might attempt to emulate the creativity of nature by writing generative software: by modelling intelligence, life, or even a complete open-ended evolving ecosystem. These are all approaches that artists have tried, the last of them, open-ended evolution, is perhaps the most promising since by definition it refers to an evolutionary system's ability to inexhaustibly generate novel frameworks. Yet an artificial system of this kind remains elusive. At the time of writing nobody has convincingly demonstrated open-ended evolution in software.

Whilst software-based open-ended evolution eludes us, some of the alternative strategies mentioned have shown a little immediate promise for generative art. In particular, artificial evolution of imagery, especially when assisted by human aesthetic assessment of fitness, has received much attention. In light of this, our creativity measure has been investigated as a means for automatically guiding the evolution of simple line drawings (Kowaliw et al. 2009). This forms part of an effort to improve the automatic and assisted creativity of evolutionary software. Our definition is, in practice, computationally intensive to apply. Hence, for practical reasons, Kowaliw devised *creativity-lite*, a modification that is tractable and testable against the full implementation of our measure. The output produced by the creativity-lite measure has been tested against common-language understanding of the term *creativity* by allowing users to evaluate automatically generated textural images (Kowaliw et al. 2012). The results of this research are summarised below.

13.8.1 The Automatic Generation of Creative Biomorphs

In order to apply our definition to a practical task, Kowaliw et al. elected to generate stick figures based on the biomorphs of Dawkins (1989) and see if they could automatically distinguish between those that were creative, and those that were not. Biomorphs are intricate line drawings made of straight segments that may be overlaid upon one another hundreds of times to generate rich patterns reminiscent of trees, insects, snowflakes, stars or countless other forms. Each individual biomorph can be considered, potentially, a creative generative system, since it is capable of

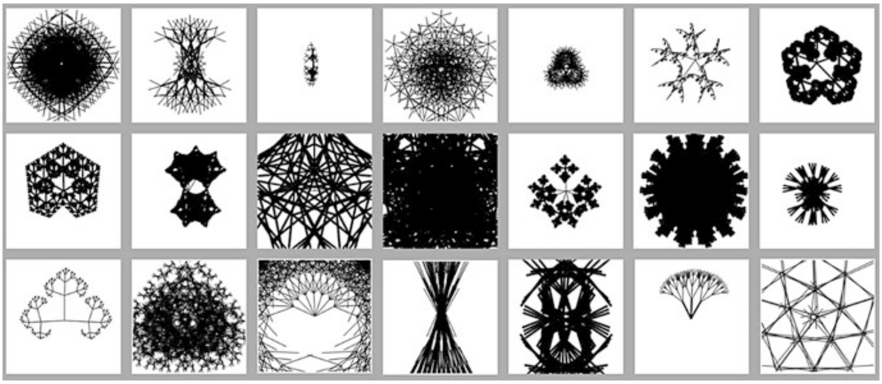

Fig. 13.3 A selection of randomly generated biomorphs generated by the software of Kowaliw, Dorin and McCormack, chosen to show some of the variety that is possible

producing a series of patterns—its offspring. Offspring are generated by varying some of a parent biomorph's parameters slightly, altering the placement, length and orientation of line segments with respect to those of that parent. A child will then be similar, but not identical to, its immediate ancestor. If the forms of a biomorph's offspring are reliably different to those produced by any that have preceded it, then according to our definition of creativity, the parent is creative with respect to its ancestors. If a particular biomorph reliably generates children that are exceedingly unlikely to be generated by any other biomorph (a possibility we can feasibly test in software), we can also claim that the parent is quite generally creative. This approach was adopted in the experiments described below.

Biomorphs were constrained in these tests to appear in 8-bit greyscale images within a 200 square pixel grid. Somewhere in this space of possible images is a representation of every possible human face, including those long dead and those that may never be born, as well as representations of every view ever seen by every person that ever looked upon the world ... and that is hardy scratching the surface! In short, this represents an astronomically large space of possible images. Of course not all of these can be generated as biomorphs. Nevertheless, the space remains unimaginably large, even when constrained in this way. An assortment of the images the software can generate is illustrated in Fig. 13.3.

The images the software generated were measured with a set of image-processing techniques that detect traits perceptible to humans. These related to the images' contrast, homogeneity, entropy and other measures of texture. A substantial sample was used to give a representation of the total space of images the software can produce and then three techniques were repeatedly applied: random image generation; an interactive genetic algorithm; and the creativity-lite search employing a simple version of our definition of creativity encoded in software, to try to locate a creative individual biomorph.

As we would hope, randomly created biomorphs did not generate a high proportion of creative offspring at all—the sampling of the space was adequate. When

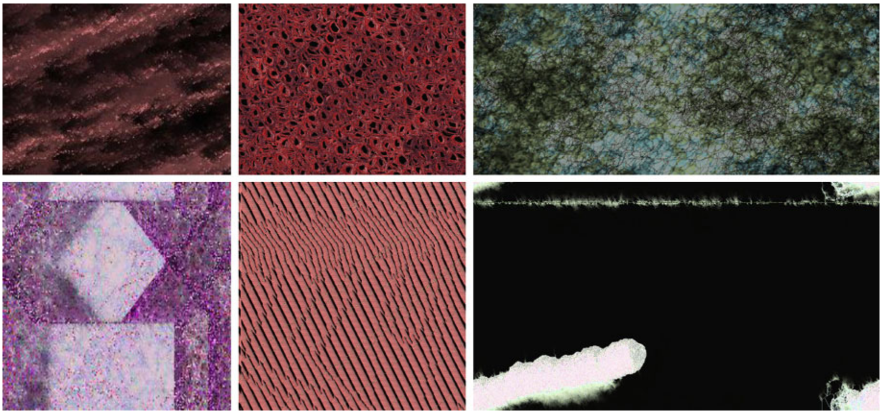

Fig. 13.4 A selection of images evolved by users employing the *EvoEco* texture generating software

creativity-lite generated biomorphs were tested against an encoded complete statistical version of our creativity measure, a very high proportion of the images generated were found to be creative by this measure. Hence it seems fair to conclude that the simplified creativity-lite search is a good approximation of the complete, but intensive, computation required to apply our measure in full. Finally, the interactive algorithm for generating biomorphs consistently produced many more creative individuals than the random technique, but less than the creativity-lite measure. This indicates that our measure goes some way towards representing human preferences. Were the humans evolving for creative results that our measure would appreciate? This test was not set up to answer such questions. How well does our measure correspond to the general human appreciation of creativity? This was explored in a subsequent test, discussed next.

13.8.2 Testing Creative Software Against Human Concepts of Creativity

In order to determine the extent to which our definition of creativity reflected human interpretations of the term, Kowaliw et al. conducted a survey, the details of which can be found in (Kowaliw et al. 2012). Visitors, originally to a gallery installation, and later to a website, were invited to evolve images using texture generating software, *EvoEco*, that operated on the principle of aesthetic selection: users selected a parent image from 16 textures displayed to them. This was then used to generate 16 offspring images as described below, and the process repeated until the user decided it was time to stop. Some sample images generated by users appear in Fig. 13.4.

Of the 16 offspring generated from a user-selected parent image, 15 were generated either by occasional mutation (a small variation in the parameters that generate the image) of the chosen parent's parameters, or by crossing-over the parent's parameters with those of another randomly selected member of the current, or any previous population. Crossover is an operation that samples parameters from one image, complementary parameters from another, and recombines these to produce a complete set of parameters for generating, in this instance, a new image. The new image typically exhibits a combination of traits of its parents.

For a particular user the final offspring image of the 16 was consistently generated either completely at random; or in such a way as to maximise its distance from the parent (according to the statistical image properties applied to the biomorph images above); or, by locating an individual that maximised the creativity-lite measure discussed in relation to the biomorphs. In the latter two cases a sample of images was generated from the parent by mutation and the one that maximised the distance or creativity-lite measure as appropriate was chosen to fill the final slot. The offspring were always positioned in random order on the screen so as to avoid user selection biases introduced by image placement.

Following their engagement with the software users were presented with a survey asking them to rank how appealing, novel, interesting or creative they found the intermediate images and final result.

The distance technique for generating offspring, generally, appeared to be an impediment to the software's success. This technique was ranked somewhat worse than the random generation technique with regard to the novelty of the images it suggested, and significantly worse than the creativity technique in three of six scales we recorded. This performance matched our intuition that maximising distance pushes populations away from the direction selected by users, undoing the effects of evolution. The creativity technique significantly outperformed the distance technique and rated better than the random technique in novelty of the final image, and the creativity of the suggested intermediate images. The research also found that the mean score for all responses was best for the creativity technique; and that the proportion of users who answered positively to the question "Did you feel that you could control the quality of the images through your selections?" was higher for creativity (55 %) than for the other two techniques (31 % for random and 40 % for distance). Hence, we can conclude tentatively that the use of the creativity-lite measure improved the performance of the interactive algorithm with respect to natural language notions of novelty and creativity.

13.9 Discussion

Where to from here? *EvoEco* is not the be-all and end-all of automated creativity. For starters, the software will never be able to step outside of its specification to generate the kind of novelty we would expect a human to devise. For instance, it can never generate a three-dimensional model of a tree. Is this significant? A hard-coded range of representations limits all systems we devise. This is inescapable. Can a digital machine exceed the expectations of its developers nevertheless? We can see no reason why not, but it must do so within the confines of its ability to manipulate representations. Thus, its creativity will be limited, perhaps to the point where we are insufficiently engaged with its behaviour to claim it as a rival to our own creativity.

Human creativity is analogously confined by physical, chemical, biological and social constraints. In keeping with both Boden's definition and our own, we find human creativity interesting because it redefines the boundaries of what is possible within the constraints we believe are imposed upon us. If we can build generative systems that interact within the same constraints as ourselves, we can potentially circumvent the limitations of the purely digital system. And this is what artists do when they map the digital world to real textual, sonic and visual events. A program running on a computer is not very interesting unless we can experience what it is doing somehow. Once we can, the discrete events—changing values, production of spatio-temporal patterns—take on new meanings for those observers that perceive them, outside the internal semantics of the generative system that created them. For this reason we can see no reason why machine creativity cannot in theory rival our own.

How will we know that our creativity has been rivalled? In the usual ways: machine-generated works will appear in the best galleries, win the most prestigious art prizes, sell for vast sums of money, be acclaimed by critics, cited in publications and referenced by other artists. But also, if we are to progress in any methodical sense, because we shall be able to assess their creativity in an objective way, employing measures such as that implicit in the new definition of creativity that we have presented.

13.10 Conclusions

We have described in this chapter a way to interpret creativity independently of informal concepts such as value and appropriateness. These concepts, we feel, have encumbered the search, production and recognition of autonomous creative software to date. Instead, we have proposed a testable measure for creativity that can be applied wherever it is possible to sample and measure the novelty of designs selected from a large, even infinite, space of possibilities. In fact our approach is, in theory, applicable to any digital computational generative system and its output. As a proof of concept, our definition has been encoded in software and employed to measure the creativity of images, enabling it to automatically offer results that humans perceive to be more creative than those made using similar, but differently assisted, algorithmic techniques.

We do not pretend to have shelled the "automated creativity nut", but we believe we have caused a fissure from which we can pry it open. No doubt our powers of persuasion are not so strong as to ensure that our ideas will be uncontroversial, particularly as our suggestions run against the grain of much that has been written previously. Nevertheless, we are at this stage convinced of the merit of our definitions and their application, and hopeful that the problems of building autonomous creative systems can be tackled from the approaches we offer.

Acknowledgements This research was supported by Australian Research Council Discovery Grants DP0772667 and DP1094064.

References

Albert, R. S., & Runco, M. A. (1999). A history of research on creativity. In R. J. Sternberg (Ed.), *Handbook of creativity* (pp. 16–31). New York: Cambridge University Press. Chapter 2.

Bada, J. L., & Lazcano, A. (2002). Some like it hot, but not the first biomolecules. *Science*, *296*, 1982–1983.

Ball, P. (2001). *The self-made tapestry: pattern formation in nature*. Oxford: Oxford University Press.

Barrass, T. (2006). Soma (self-organising map ants). In *EvoMUSART2006, 4th European workshop on evolutionary music and art*.

Bentley, P. J. (2002). Is evolution creative? In P. J. Bentley & D. Corne (Eds.), *Creative evolutionary systems* (pp. 55–62). San Mateo: Morgan Kaufmann.

Berlyne, D. E. (1960). *Novelty, uncertainty, conflict, complexity*. *McGraw-Hill series in psychology* (pp. 18–44). New York: McGraw-Hill. Chapter 2.

Berry, R., Rungsarityotin, W., Dorin, A., Dahlstedt, P., & Haw, C. (2001). Unfinished symphonies—songs of 3.5 worlds. In E. Bilotta, E. R. Miranda, P. Pantano & P. M. Todd (Eds.), *Workshop on artificial life models for musical applications, sixth European conference on artificial life* (pp. 51–64). Editoriale Bios.

Boden, M. A. (2004). *The creative mind, myths and mechanisms* (2nd ed.). London: Routledge.

Borel, É. (1913). Mécanique statistique et irréversibilité. *Journal of Physics*, *3*(5), 189–196.

Borgia, G. (1995). Complex male display and female choice in the spotted bowerbird: specialized functions for different bower decorations. *Animal Behaviour*, *49*, 1291–1301.

Carroll, J. (2007). The adaptive function of literature. In C. Martindale, P. Locher & V. M. Petrov (Eds.), *Evolutionary and neurocognitive approaches to aesthetics, creativity and the arts* (pp. 31–45). Amityville: Baywood Publishing Company. Chapter 3.

Chamberlain, W. (1984). Getting a computer to write about itself. In S. Ditlea (Ed.), *Digital Deli, the comprehensive, user-lovable menu of computer lore, culure, lifestyles and fancy*. New York: Workman Publishing. http://www.atariarchives.org/deli/write_about_itself.php. Accessed 21 October 2010.

Colton, S. (2001). *Automated theory formation in pure mathematics*. PhD thesis, University of Edinburgh.

Colton, S., Bundy, A., & Walsh, T. (2000). On the notion of interestingness in automated mathematical discovery. *International Journal of Human Computer Studies*, *53*(3), 351–375.

Csikszentmihalyi, M. (1999). Creativity. In R. A. Wilson & F. C. Keil (Eds.), *The MIT encyclopaedia of the cognitive sciences* (pp. 205–206). Cambridge: MIT Press.

Dahlstedt, P. (1999). Living melodies: coevolution of sonic communication. In A. Dorin & J. McCormack (Eds.), *First iteration: a conference on generative systems in the electronic arts* (pp. 56–66). Melbourne: CEMA.

Dawkins, R. (1989). The evolution of evolvability. In C. G. Langton (Ed.), *Artificial life: proceedings of an interdisciplinary workshop on the synthesis and simulation of living systems* (pp. 201–220). Reading: Addison-Wesley.

Dorin, A. (2004). The virtual ecosystem as generative electronic art. In G. R. Raidl, S. Cagnoni, J. Branke & D. W. Corne (Eds.), *2nd European workshop on evolutionary music and art, applications of evolutionary computing: Evo workshops*, Portugal (pp. 467–476). Berlin: Springer.

Dorin, A., & Korb, K. B. (2009). A new definition of creativity. In K. Korb, M. Randall & T. Hendtlass (Eds.), *LNAI: Vol. 5865. Fourth Australian conference on artificial life*, Melbourne, Australia (pp. 1–10). Berlin: Springer.

Earnshaw, C. G. (1991). *Mozart's Dice Waltz*. Atari users. http://tamw.atari-users.net/mozart.htm. Retrieved 21 October 2010.

Eddington, A. S. (1927). *The nature of the physical world: Gifford lectures*. New York: Macmillan.

Eldridge, A. (2009). *You pretty little flocker*. Guildford Lane Gallery.

Gracia, J. J. E. (1996). *Texts: ontological status, identity, author, audience*. New York: State University of New York Press.

Harnad, S. (1990). The symbol grounding problem. *Physica D*, *42*(1–3), 335–346.

Hofstadter, D. (1995). *Fluid concepts and creative analogies, computer models of the fundamental mechanisms of thought*. Harper Collins: Basic Books.

Johnson-Laird, P. (1987). Reasoning, imagining and creating. *Bulletin of the British Psychological Society*, *40*, 121–129.

Joyce, G. F. (1989). RNA evolution and the origins of life. *Nature*, *338*, 217–224.

Kowaliw, T., Dorin, A., & McCormack, J. (2009). An empirical exploration of a definition of creative novelty for generative art. In K. Korb, M. Randall & T. Hendtlass (Eds.), *LNAI: Vol. 5865*. *Fourth Australian conference on artificial life (ACAL 2009)*, Melbourne, Australia (pp. 1–10). Berlin: Springer.

Kowaliw, T., Dorin, A., & McCormack, J. (2012). Promoting creative design in interactive evolutionary computation. *IEEE Transactions on Evolutionary Computation*, *16*(4), 523–536.

Langley, P., Simon, H. A., Bradshaw, G. L., & Zytkow, J. M. (1987). *Scientific discovery: computational explorations of the creative processes*. Cambridge: MIT Press.

Lenat, D. B. (1983). Eurisko: a program that learns new heuristics and domain concepts: the nature of heuristics III: program design and results. *Artificial Intelligence*, *21*, 61–98.

Maynard-Smith, J., & Szathmry, E. (1995). *The major transitions in evolution*. New York: Freeman.

McCormack, J. (2001). Eden: an evolutionary sonic ecosystem. In J. Kelemen & P. Sosik (Eds.), *LNAI: Vol. 2159*. *Advances in artificial life, 6th European conference, ECAL*, Prague, Czech Republic (pp. 133–142). Berlin: Springer.

Miller, G. F. (2001). Aesthetic fitness: how sexual selection shaped artistic virtuosity as a fitness indicator and aesthetic preferences as mate choice criteria. *Bulletin of Psychology and the Arts. Special issue on evolution, creativity, and aesthetics*, *2*(1), 20–25.

Nake, F. (2002). Personal recollections of a distant beginning. In L. Candy & E. A. Edmonds (Eds.), *Explorations in art and technology* (p. 6). London: Springer.

Odling-Smee, F. J., Laland, K. N., & Feldman, M. W. (2003). *Niche construction, the neglected process in evolution*. *Monographs population biology*. Princeton: Princeton University Press.

Phillips, J. (1935). Succession, development, the climax, and the complex organism: an analysis of concepts: part III. The complex organism: conclusions. *Journal of Ecology*, *23*(2), 488–508.

Putnam, H. (1979). *The meaning of "meaning", philosophical papers: mind, language and reality* (Vol. 2, pp. 215–271). Cambridge: Cambridge University Press. Chapter 12.

Rowe, J., & Partridge, D. (1993). Creativity: a survey of AI approaches. *Artificial Intelligence Review*, *7*, 43–70.

Saunders, R., & Gero, J. S. (2002). Curious agents and situated design evaluations. In J. S. Gero & F. M. T. Brazier (Eds.), *Agents in design* (pp. 133–149). Key Centre of Design Computing and Cognition, University of Sydney.

Schmidhuber, J. (2009). *Driven by compression progress: a simple principle explains essential aspects of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes*. arXiv:0812.4360v2 [cs.AI].

Sims, K. (1991). Artificial evolution for computer graphics. *Computer Graphics*, *25*(4), 319–328.

Smuts, J. C. (1927). *Holism and evolution*. London: Macmillan and Co.

Tatarkiewicz, W. (1980). *A history of six ideas, an essay in aesthetics*. Warsaw: PWN Polish Scientific Publishers.

Todd, S., & Latham, W. (1992). *Evolutionary art and computers*. San Diego: Academic Press.

Weesatchanam, A.-M. (2006). Are paintings by elephants really art? www.elephantartgallery.com. Retrieved 31 July 2009.