Chapter 2
Creative Ecosystems

Jon McCormack

Abstract Traditional evolutionary approaches to computer creativity focus on optimisation, that is they define some criteria that allows the ranking of individuals in a population in terms of their suitability for a particular task. The problem for creative applications is that creativity is rarely thought of as a single optimisation. For example, could you come up with an algorithm for ranking music or painting? The difficulty is that these broad categories are shifting and subjective: I might argue that Mozart is more musically creative than Lady Gaga, but others may disagree. Objective, fine-grained ranking of all possible music is impossible, even for humans. I will show how reconceptualising the exploration of a creative space using an "ecosystemic" approach can lead to more open and potentially creative possibilities. For explanatory purposes, I will use some successful examples that are simple enough to explain succinctly, yet still exhibit the features necessary to demonstrate the advantages of this approach.

2.1 Creative Systems

In this book you will find a broad range of definitions of creativity. Dorin and Korb (Chap. 13), for example, emphasise a system's propensity to generate novelty irrespective of its perceived value, similarly Schmidhuber (Chap. 12) views creativity as a problem of learning information compression. Nake (Chap. 3) is more sceptical about formal computer models of creativity, seeing the popular concept of creativity today as "a US-American invention," one that may be considered as a *means* for activity, or as its *goal*. Pachet (Chap. 5) prefers to focus on "virtuosity", emphasising the thousands of hours that human artists must spend to master a discipline or instrument. Each of these views place a different emphasis on which qualities, properties or functions are important to understanding creativity precisely, and hence appreciating its worth or relevance in any given domain.

If we take Boden's popular definition—that creativity involves the generation of ideas or artefacts that are *new, surprising,* and *valuable* (Boden 2010)—then an

J. McCormack (✉)
Centre for Electronic Media Art, Monash University, Caulfield East, Victoria 3145, Australia
e-mail: Jon.McCormack@monash.edu

J. McCormack, M. d'Inverno (eds.), *Computers and Creativity*,
DOI 10.1007/978-3-642-31727-9_2,

interesting question to ask is: what are the mechanisms that enable this creativity? It appears likely that any such mechanisms are numerous and diverse. While creativity is commonly associated with the human individual, clearly societies and nature invent, too.

The psychologist David Perkins (1996) talks about "creative systems"; recognising that there are different mechanisms or classes of underlying systems that are all capable of producing creative artefacts. A creative system, in this view, is simultaneously capable of the production of novelty and adaptation in a given context. This suggests natural selection is a creative system, generating things like prokaryotes, multicellularity, eusociality and language, all through a non-teleological process of hereditary replication and selection. Social interaction is another creative system, having given rise to cultural customs such as shaking hands and a variety of grammatical forms in different human languages.

A number of authors have offered explanations of fundamental creative mechanisms based on evolution or evolutionary metaphors, e.g. Martindale (1999), Lumsden (1999), Dawkins (1999), Aunger (2002). George Basalla's *The Evolution of Technology* detailed a theory of technological evolution, offering an explanation for the creative diversity of human made artefacts: "*novelty* is an integral part of the made world; and a *selection* process operates to choose novel artifacts for replication and addition to the stock of made things" (Basalla 1998). Evolution has also played an important role in computer-based and computer-assisted creative systems (Bentley and Corne 2002), being able to discover, for instance, seemingly counterintuitive designs that significantly exceed any human designs in performance (Keane and Brown 1996, Eiben and Smith 2003, p. 10). Such results illustrate the potential of evolutionary systems to devise unconventional yet useful artefacts that lie outside the capabilities of current human creative thinking.

Defining a class of phenomena in formal, systemic terms allows for a transition to the computer. The purpose of this chapter is to look at what kinds of computational processes might qualify as "creative systems" in their own right. Here I draw my inspiration from natural systems, in particular evolutionary ecosystems. Biological evolution is readily accepted as a creative system, as it is capable of discovering "appropriate novelty". The computer science adaptation of evolution, a field known as *Evolutionary Computing* (EC), selectively abstracts from the processes of biological evolution to solve problems in search, optimisation and learning (Eiben and Smith 2003). It is important to emphasise *selectively abstracts* here, as only certain components of the natural evolutionary process are used, and these are necessarily highly abstracted from their physical, chemical and biological origins, for both practical and conceptual reasons. In the case of designing a creative system, the challenge is somewhat different than that of standard EC: understanding how a process that is creative in one domain (biology) can be transformed to be creative in another (e.g. the creation of art) requires different selective abstractions.

Generating the adaptive novelty exhibited in creative systems can be conceptualised as a process of exploration through a space of possibilities, searching for regions of high creative reward. Perkins (1996) uses the metaphor of the "Klondike space"—*Gold is where you find it.* Perkins identified four basic problem types in

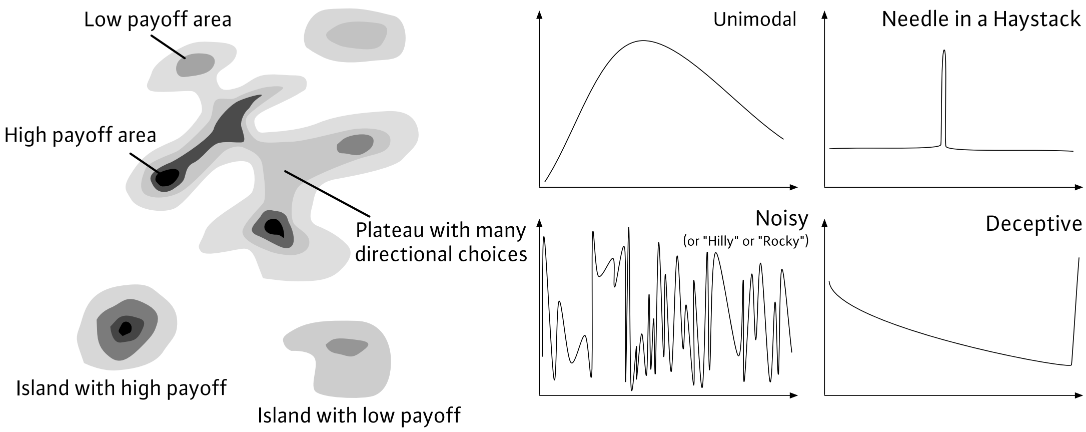

Fig. 2.1 Illustrative diagram of "Klondike spaces" (*left*, after Bell 1999) and, characterisation of archetypical search spaces in Evolutionary Computing (*right*, after Luke 2009)

the creative search of a conceptual space (Fig. 2.1, left): (i) *rarity:* viable solutions are sparsely distributed in a vast space of non-viable possibilities; (ii) *isolation*: places of high creative value in the conceptual space are widely separated and disconnected, making them difficult to find; (iii) *oasis*: existing solutions offer an oasis that is hard to leave, even though better solutions might exist elsewhere; (iv) *plateau*: many parts of the conceptual space are similar, giving no clues as to how to proceed to areas of greater creative reward.

This classification is similar to archetypical search and optimisation problems encountered in EC (Fig. 2.1, right), where algorithms search for optima in what are often difficult phenotypic spaces (Luke 2009). For example, "rarity" corresponds to "Needle in a haystack", "oasis" to "Deceptive". Noisy landscapes are particularly problematic, where evolutionary methods may do no better than random search.

Knowing as much as possible about the structure of the space you are searching is immensely important, as it allows you to strategically search using the most efficient methods. Additionally, being able to restructure the space can make it more intuitive for creative exploration. Hence the design of any creative system should take the structural design of the creative space very seriously. It is also important to emphasise that the search process is an explorative one. For most creative systems, this search space is *Vast* (McCormack 2008b), and there may be many isolated "Klondike spaces" of rich creative reward. The challenge is to efficiently and effectively find and explore them.

2.1.1 Spaces of Possibility

We should make further distinctions about creative spaces and spaces of possibility. As I have previously discussed (McCormack 2008b), in many domains there are large and crucial differences between the possible and actual. For example, consider a digital image defined by executing an arbitrary Lisp expression over some domain (x, y), where x and y are the co-ordinates of a rectangular grid of pixels that comprise the image. Iterating through each co-ordinate, the expression returns the corresponding pixel's colour. Different expressions will usually generate different images (although many different expressions will also generate the same image). In theory, this system is capable of generating *any* possible image, provided you have the appropriate Lisp expression to generate it.

This represents a space of possibilities than encompasses every possible image that can be represented by coloured pixels over (x, y). For any reasonable image dimensions, the size of this space is Vast, far beyond comparisons with astronomical maximums such as the age of the universe, or the number of basic sub-atomic particles estimated to exist in the universe.

However, the *actual* space of images that can be practically created with a Lisp expression is considerably smaller, limited by physical constraints. From the perspective of evolutionary creativity, if we evolve a Lisp expressions using, for example, an Interactive Genetic Algorithm (IGA, see Sect. 2.2), the actual images produced are all relatively similar and represent an infinitesimally small fraction relative to the possible space of which the system is theoretically capable.[1]

So while a representational system may theoretically cover a large range of possibilities, searching them—even with evolutionary methods—will only permit examination of insignificantly small regions. Furthermore, transformation or modification of the underlying generative mechanism[2] may open up new spaces not so easily found by the original, e.g. the addition of symmetry functions for the Lisp expression example would make it easier to generate images with symmetric elements. Of course we need some way of finding the "right" transformations or modifications to make. This is a kind of "meta-search" (a search of the different types of generative mechanisms that define a representational space). Further, this opens a hierarchy (meta-meta-search, meta-meta-meta-search, etc.), which effectively amounts to the same problem of the possible and actual in our original "flat" search.

What this means in practical terms is that there must be some human-defined generative mechanism as the basis for any computational creative system,[3] which will require serious human ingenuity and creativity if it's design is to be effective. I will return to this point in Sect. 2.4.3. While much research effort and discussion has focused on evaluation and judgement in computational creative systems, representation has received far less attention.

A somewhat analogous situation exists in biology. The space of possible DNA sequences is far greater than the space of viable, or possible, phenotypes.[4] The space of possible phenotypes (those which could exist) is again larger than the space of

[1] By my estimates, about 5×10^{-1444925} % for images of modest dimensions, far beyond astronomically small.

[2] By "generative mechanism" I am technically referring to the genotype and the mechanism that expresses it into a phenotype.

[3] The mechanism can include the ability to self-modify, change, or learn.

[4] We might think of "viable" as meaning being able to effectively express a living organism from a zygote or through mitosis of a parent cell. But this is problematic for many reasons, most of which are too tangential to the argument to list here.

actual phenotypes (those which have existed, or currently exist). In nature, what can be successfully expressed by DNA is limited materially by physical constraints and processes. In contrast to our Lisp expression example, once RNA and DNA were established evolution has not really experimented with different self-replication mechanisms. We think of DNA as being a highly successful self-replicating molecule, which might be true, but we have little to compare it with. Many factors affect the *variety* of life that has evolved on Earth. As evolution involves successful adaptations, the changing environment of the Earth is an important factor in determining evolutionary variety. In addition to geological events, environments change due to presence of species and their interactions, a point that I will return to later in this chapter.

2.2 Evolutionary Computing and Creativity

As noted in the previous section, EC methods (which include techniques such as Genetic Algorithms, Evolutionary Strategies and Genetic Programming) have demonstrated success in assisting users of complex creative systems to better locate regions of high creative reward (Bentley and Corne 2002, Romero et al. 2008). In broad terms they are "generate and test" algorithms that evolve a population of candidate solutions or artefacts. New, child artefacts are generated through random mutation and/or recombination with selected parents. Populations are tested or ranked by some measure, with the most highly valued individuals and their offspring more likely to survive in subsequent generations. Incrementally, the overall "quality" of the population *should* improve according to the fitness measure used. How well the method does depends on many factors, including the nature of the fitness landscape (determined in part by the representational scheme) and the evaluation of solution fitness in artefacts. Success or otherwise is dependent on (i) the structure of the phenotype space, and (ii) the effectiveness of the fitness evaluation in determining the quality of the artefacts produced.[5]

Evolutionary approaches and aesthetic evaluation are reviewed extensively in the chapter by Galanter (Chap. 10). So it is pertinent here to make just a few points. Firstly, it is important to differentiate between an evolutionary system that gives *creative* results and one that generates *aesthetically pleasing* results. The former does not preclude the latter, but they are in general, independent (i.e. it is possible for a machine or algorithm to generate aesthetically pleasing images without that system being creative). This distinction is often overlooked.

Some evolutionary systems use learnt or predefined measures of "creative" features in their generated artefacts (Baluja et al. 1994, Machado and Cardoso 2002), or rely on some form of aesthetic measure to evaluate an individual's fitness (Birkhoff 1933, Staudek 2002, Ramachandran 2003, Svangåard and Nordin 2004, Machado et al. 2008). Others use iterative human selection to rank individuals as part of the

[5] This issue is a topic of discussion in Chap. 4.

evolutionary process (Takagi (2001) provides a comprehensive survey). These approaches suffer from difficulties, however. Pre-defined measures of aesthetic properties, for example, risk implicit judgements as to which specific properties are of value (thus determining *what* will be measured). While a number of researchers describe "aesthetic universals" of evolutionary origin (Brown 1991, Dissanayake 1995, Martindale 1999, Ramachandran and Hirstein 1999, Dutton 2002), it is long proposed that aesthetic values also shift according to individual taste, time and culture. Moreover, aesthetics has many interpretations (Koren 2010), and in contemporary art surface aesthetic qualities are often downplayed or given little significance in appreciating the creativity of the work. Evolving artefacts exclusively for aesthetic value does not necessarily make them creative.

Some attempts have been made to expressly minimise or remove the aesthetic judgement of a particular individual. This is what is referred to as removing "the signature" of the artist (Boden 2010, Chaps. 9 & 10). The *Drawbots* system described by Bird et al. (2008) attempted to create a line-drawing robot using evolutionary robotics. Researchers defined "implicit" fitness measures that did not restrict the type of marks the robot drawer should make, including an "ecological model" involving interaction between environment resource acquisition and expenditure through drawing. However, the results demonstrated only minimal creativity, and the authors concluded that fitness functions which embodied "artistic knowledge about 'aesthetically pleasing' line patterns" would be necessary if the robot were to make drawings worthy of exhibition to humans.

Using human selection (known as the *Interactive Genetic Algorithm,* IGA) suffers from a "fitness evaluation bottleneck" that reduces the human operator role to that of a "pigeon breeder" who quickly fatigues (Takagi 2001, Dorin 2001). IGAs are generally more suited to explorations by a non-expert user, who is unfamiliar with the generative mechanism being evolved. Here the IGA allows limited navigation through a space of possibilities without necessarily understanding the underlying mechanisms that generate them.[6]

These standard evolutionary approaches, while historically important and capable of significant results, are not able to consistently generate convincingly creative results in many domains. Can we do better? Biology seemingly can. A useful insight is in recognising that finding the creative "Klondike spaces" is not simply an optimisation problem (i.e. finding a global optima using some fitness criteria). Indeed, for most creative domains the idea of evolving towards a single optimum is counterintuitive, as an artist or designer normally produces many new artefacts over their professional lifetime. New designs or techniques often "evolve" from previous ones, offspring of both the originating artist and his or her peers (Basalla 1998). As Basalla (1998) and others have pointed out using the example of technological evolution, the Western emphasis on individual creativity (reinforced socially through patents and other awards) obscures the important roles played in the evolutionary

[6]Although there are exceptions where the IGA has proved useful to expert users as well, e.g. Dahlstedt (2006), McCormack (2008a).

ecosystem of interactions between environment and prior work of many individuals. Hence:

> The trajectory through a creative space is not one of incrementally optimising towards a single goal or fitness measure, rather it is a complex pathway through a series of intermediate and changing goals, each of which may determine the pathway of the next, and may be creative in its own right.

If we are interested in discovering new creative spaces through the synergetic combination of human intelligence and intuitive structuring and representation of the conceptual space, then there are other possibilities. The evolution of species on earth involves a complex set of interrelated processes and events. For example, species do not exist in isolation from their environment or from other species; together they form a complex network of interdependencies that may impact on the evolutionary process significantly. Let us see what happens if we re-conceptualise the search of a creative space using insights from the structure and function of evolutionary biological ecosystems.

2.3 Ecosystems

Ecosystems are a popular yet somewhat nebulous concept increasingly adopted in contemporary culture. Environmental groups want to preserve them, businesses want to successfully strategise and exploit them, and the media is part of them. With recent sales of Nokia mobile smartphones on the decline, Nokia CEO Stephen Elop bemoaned that fact that his company, unlike its rivals, had failed to create an "ecosystem": one that encompassed smartphones, the operating system, services and users (Shapshak 2011). Media theorists speak of "media ecologies"—the "dynamic interrelation of processes and objects, beings and things, patterns and matter" (Fuller 2005). Philosopher Manuel De Landa emphasises the flows of energy and nutrients through ecosystems manifesting themselves as animals and plants, stating that bodies are "nothing but temporary coagulations in these flows: we capture in our bodies a certain portion of the flow at birth, then release it again when we die and micro-organisms transform us into a new batch of raw materials" (De Landa 2000).

In the broadest terms, the modern concept of an ecosystem suggests a community of connected, but disparate components interacting within an environment. This interaction involves dependency relationships leading to feedback loops of causality. The ecosystem has the ability to self-organise, to dynamically change and adapt in the face of perturbation. It has redundancy and the ability to self-repair. Its mechanisms evoke symbiosis, mutualism and co-dependency, in contrast to pop-cultural interpretations of evolution as exclusively a battle amongst individuals for fitness supremacy. Yet we also speak of "fragile ecosystems", implying a delicate balance or harmony between elements that can easily be broken by external interference. Any anthropomorphic projection of harmony or stability to ecosystems is naïve however. The history of evolution is the history of change: species, their diversity, morphology and physical distribution, the chemical composition of the biosphere, the geography of the earth—all have changed significantly over evolutionary time. The ecosystem's stability is seemingly transitory then, tied to the shifts in species distribution and environment.

2.3.1 Biological Ecosystems

Of course, ecosystems and Ecology are the domain of Biology, where we find a formal understanding, along with many inspirational ideas on the functional relationships found in real biological ecosystems. Modern Ecology is the study of species and their relations to each other and their environment. The term "Ecology" originated with the German Biologist and Naturalist, Ernst Haeckel,[7] who, in 1866, defined it as the "science of the relationship of the organism to the environment", signifying the importance of different species embedded in specific environments. The term "Ecosystem", from the Greek ($o\iota\kappa o\varsigma$, household; $\lambda o\gamma o\varsigma$, knowledge) is attributed to the British Ecologist, Sir Arthur Tansley, who coined it from fellow Botanist Arthur Clapham. It grew out of debates at the time about the similarity of interdependent communities of species to "complex organisms". Importantly, Tansley's use of the term ecosystem encompassed "the inorganic as well as the living components" (Tansley 1939), recognising that the organism cannot be separated from the environment of the biome, and that ecosystems form "basic units of nature" (Willis 1997).

Contemporary definitions of ecosystems begin with the work of American Ecologists Eugene and Howard Odum. Eugene wrote the first detailed Ecology text, *Fundamentals of Ecology*, published in 1953. Odum recognised energy flows, trophic levels,[8] functional, and causal relationships that comprised the ecosystem. Willis defines the modern concept of an ecosystem as "a unit comprising a community (or communities) of organisms and their physical and chemical environment, at any scale, desirably specified, in which there are continuous fluxes of matter and energy in an interactive open system" (Willis 1997).

In more modern terms, Scheiner and Willig (2008) nominate seven fundamental principles of ecosystems:

1. Organisms are distributed in space and time in a heterogeneous manner (inclusionary rule).

[7] Danish biologist Eugen Warming is also attributed as the founder of the science of Ecology.

[8] *Autotrophs*, such as plants, produce organic substances from simpler inorganic substances, such as carbon dioxide; *heterotrophs* unable to perform such conversions, require organic substances as a source of energy.

2. Organisms interact with their abiotic and biotic environments (inclusionary rule).
3. The distributions of organisms and their interactions depend on contingencies (exclusionary rule).
4. Environmental conditions are heterogeneous in space and time (causal rule).
5. Resource are finite and heterogeneous in space and time (causal rule).
6. All organisms are mortal (causal rule).
7. The ecological properties of species are the result of evolution (causal rule).

For those wanting to know more details on the contemporary science, a text such as that by Begon et al. (2006) provides a useful overview of Ecology science.

2.3.2 Ecosystem Models in the Creative Arts

A number of different "ecosystemic" approaches exist in the arts. Examination finds that they are quite diverse and only loosely drawn from biological concepts, probably due to multiplicitous and nebulous understandings of Ecology outside Biology, and various metaphoric interpretations of the ecosystem concept.

Design and Architecture. Given the state of human impact on the environment, much theory in landscape and architectural design has sought to bring ideas from Ecology and ecosystems into the design lexicon (see, e.g. Bell 1999). Through a greater understanding of nature's process and function, it is believed that designers can better integrate human interventions within the landscape, minimising their detritus impact, or at least appreciate how design decisions will effect change to the environment over the life of a project, and beyond. In architecture, *Design Ecologies* seeks connections between biological Ecology, human communication, instruction and aesthetics, with an emphasis on "novel concepts of ecologically informed methodologies of communication through design practice" (Murray 2011).

Generative design uses processes adopted from evolution as a source of design variation and customisation. It brings a number of desirable features to the design of artefacts, including a means to generate and manage complexity; self-maintenance and self-repair; design novelty and variation (McCormack et al. 2004). As discussed (Sect. 2.2), evolutionary methods such as the IGA are useful for generative design when the designer has only a rudimentary grasp of the underlying generative mechanism that is being evolved. They permit design changes without the need to understand in detail the configuration or parameter settings that generated the design. The application of generative design to customised manufacture has become feasible in recent years due to the availability of automated, programmable fabrication devices, such as 3D printers, laser cutters, etc. that can inexpensively translate computer representations into one-off physical objects. This allows physical generative designs to be customised to individual constraints or desires on commercial manufacturing scales.

Design associations with Ecology and ecological principles often suggest the superiority of natural over human design, and ecosystems embracing harmony and stable configurations, "in tune" with nature and natural surroundings. Ecological processes provide a certain cachet, appeal and authority that conveniently lend both a design and moral credibility to a project. Such views have been rightly criticised (Kaplinsky 2006). Evolution needs only to offer adequate solutions—ones that are sufficient for growth, survival and reproduction—not necessarily the best or globally optimal ones. "Optimality" for evolution is dependent on environment (obviously polar bears don't do well in deserts). But it is not that nature has nothing useful to teach us. Moving beyond mimicry, a better understanding of the function and behaviour of real biological ecosystems offers new and rewarding possibilities for design, along with a greater awareness of how our activities ripple out through the environment and affect other species.

Music and Performance. Waters (2007) uses the concept of a "performance ecosystem"—one that encompasses composition, performance, performers, instruments and environment. Here music and music making are seen as part of a multi-layered, complex dynamical system, operating from the acoustic to the social. Emphasis is placed on the dynamical interactions and, importantly, feedback processes between components of the ecosystem. For example, the feedback between a performer and their instrument encompasses the body, tactility, vibrating materials, physical and acoustic properties of the room in which the instrument is played, along with the "psychological adaptations and adjustments" in the body of the performer, who is deeply connected to, and part of these interacting elements.

Such connections evoke the cybernetic: instruments can be considered part of a continuum that originates from the body, extending through instrument and environment. Italian composer, Agostino Di Scipio (2003) seeks a reformulation of what is meant by "interaction" in a technological performance context and invokes the cybernetic concept of ecosystems and feedback dependencies as a sonic interaction paradigm. This is indicative of a more general sense of failure, in creative contexts, of standard technical approaches to human-computer interaction. These traditional approaches emphasise the functional over the explorative and connected. An alternate view, advocated by Di Scipio and many others, sees interaction as "a by-product of lower level interdependencies among system components" (Di Scipio 2003). Components are *adaptive* to their surrounding external conditions and able to *manipulate* them. In the case of sound, this involves a sound ecosystem of sound-generating, sound-listening and sound-modifying components, connected in feedback loops with their acoustic environment. In this configuration sound itself is the medium in which the ecosystem exists. The coupling of components with their environment allows them to change and reconfigure in response to environmental variation: an environment that the components themselves may be modifying.

Visual and Installation Art. My own interactive installation, *Eden* (McCormack 2001), is a complex artificial ecosystem running in real-time on a two-dimensional lattice of cells, projected into a three-dimensional environment (Fig. 2.2). The simulation includes seasonal variation, planetary albedo modified by biomass composition (Lenton and Lovelock 2001), and a simulation of sound propagation and at-

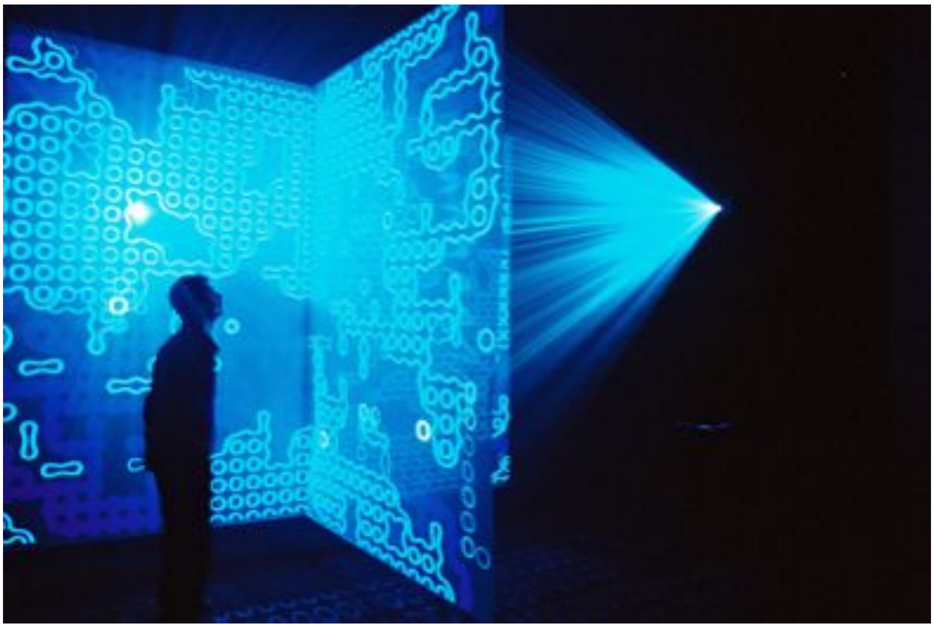

Fig. 2.2 The author's *Eden* installation: an evolving ecosystem of virtual creatures learn new behaviours based on interaction with their environment and with their human audience

tenuation. Evolving, learning agents modify and adapt to their surroundings. Interestingly, the agents learn a number of behaviours not explicitly programmed into the system, including hibernation during winter months when food resources are scarce, predation, and primitive signalling using sound. A computer vision system links human visitor presence to the generation of biomass (food for the agents), and over time agents learn to make interesting sequences of sound in order to keep visitors attracted near the work, thus increasing their supply of food and chances of reproductive success (McCormack 2005).

Over the last twenty years, Dutch artists Erwin Driessens and Maria Verstappen[9] have been experimenting with generative "processes of production" in their art practice. This has extensively encompassed the use of ecosystem metaphors in a number of their works. For example, *E-volver* is a generative visual artwork where a small collection of agents roam a gridded landscape of coloured pixels, choosing to modify the pixel underneath them based on it's colour, and those of the neighbouring pixels. Each agent has a set of rules that determine how to change the colour and where to move next (Driessens and Verstappen 2008). Through the interaction of these pixel-modifying agents and their environment (the pixels which comprise the image), E-volver is able to generate a fascinating myriad of complex and detailed images (Fig. 2.3 shows one example), all of which begin from a uniformly grey canvas. The images, while abstract, remind the viewer of landscape viewed from high altitude, or an alien mould overwhelming a surface, or electron micrographs of some unidentified organic structure. Importantly, they exhibit details on a variety of scales, with coherent structures extending far beyond the one pixel sensory radius of

[9] See their website at: http://www.xs4all.nl/~notnot/index.html.

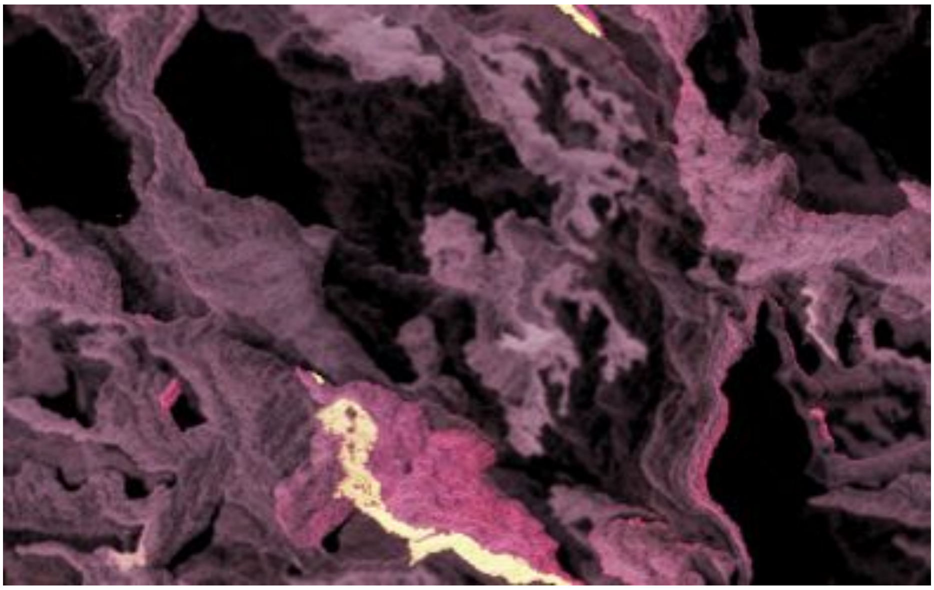

Fig. 2.3 An image produced by Driessens and Verstappen's E-volver. Eight pixel modifying agents build the image by modifying pixels. Notice the image contains coherent structures over multiple levels of detail

the agents that created them. This suggests a collective self-organisation achieved through agent-environment interaction, with the environment acting as a "memory" that assists agents in building coherent structures within the image.

Like Di Scipio's sonic ecosystems, E-volver's "environment" is *the medium itself* (an image comprised of coloured pixels). For Eden, the real and virtual environments are causally connected through sound, human presence and the production of resources. In both E-volver and Eden, agents modify their environment which, in part, determines their behaviour. Causally coupling agent to environment allows for feedback processes to be established, and the system thus becomes self-modifying. This iterative self-modification process facilitates the emergence of heterogeneous order and fractaline complexity from an environment of relative disorder and simplicity. For Eden this is further expanded by the use of an evolutionary learning system (based on a variant of Wilson's XCS (Wilson 1999)) that introduces new learning behaviours into the system. Learnt behaviours that have been beneficial over an agent's lifetime are passed onto their offspring.

Unlike Eden's learning agents, E-volver's agents are not evolutionary over the life of the ecosystem, yet they are evolved: a variation on the IGA allows the user of the system to evolve ecosystem behaviours through aesthetic rejection ("death of the unfittest"). The entire ecosystem (a set of eight agents and their environment) is evolved, not individual agents within a single image. Selection is based on the subjective qualities of the images produced by an individual ecosystem.

There are numerous other examples of successful artworks based on ecosystem metaphors and processes. To return to the central questions of this chapter: *how* and *why* do they work successfully?

Table 2.1 General properties of creative ecosystem models

Property	Features
Components & their environment	Together these constitute the ecosystem
Dynamical system	Enables the ecosystem to temporally adapt and change in response to internal and external conditions
Self-observation	Provides a link between component action and environment
Self-modification	Allows a component to adjust its behaviour within the system
Interaction	Components must interact with each other and their environment to give rise to emergent behaviours of the system as a whole
Feedback loops	Provide pathways of control, regulation and modification of the ecosystem
Evolution	Allows long term change, learning and adaptation

2.4 Ecosystem Design Patterns

Within our research group[10] at the Centre for Electronic Media Art we have investigated ecosystemic processes as a basis for designing or enhancing generative artworks (see e.g. McCormack (2001, 2007b, 2007a), Eldridge et al. (2008), Eldridge and Dorin (2009), Bown and McCormack (2010)). Our long-term aim has been to develop a catalogue of ecosystemic "design patterns" in the spirit of Gamma et al. (1995), which facilitate the building of creative evolutionary systems. Developing these patterns does not imply a "plug-and-play" approach where one just selects the appropriate patterns, connects them together, and then sits back to watch the creativity evolve. Rather, the patterns serve as starting points in conceptualising a specific creative system, documenting intermediate mechanisms and the typical behaviours they produce. Choosing *which* pattern to use and *how* to apply them remains a matter of significant creative judgement.

Di Scipio sees the artistic system as a "gathering of connected components", and it is these components and their interdependencies that must be carefully designed if successful system-level results are to ensue. Components must additionally be adaptive to surrounding external conditions and be able to manipulate them.

Table 2.1 summarises the basic properties we think are important to creative ecosystem models. The key to developing a successful ecosystem model is in the design of the system's components, their meaning, interpretation and interaction. In the following sections, I will explore some of these features in more detail, using completed ecosystem artworks as examples.

2.4.1 Environments: Conditions and Resources

In broad terms, biological environments have two main properties that determine the distribution and abundance of organisms: *conditions* and *resources*. Conditions are

[10]Which has included over the last few years: Oliver Bown, Palle Dahlstedt, Alan Dorin, Alice Eldridge, Taras Kowaliw, Aidan Lane, Gordon Monro, Ben Porter and Mitchell Whitelaw.

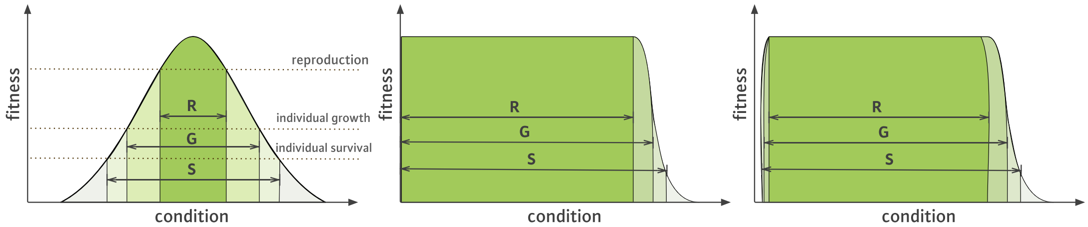

Fig. 2.4 Example organism viability curves for reproduction, growth and survival, from Begon et al. (2006)

physiochemical features of the environment (e.g. temperature, pH, wind speed). An organism's presence may change the conditions of its local environment (e.g. one species of plant may modify local light levels to that which another species is adapted for). Conditions may vary in cyclic patterns or be subject to the uncertainty of prevailing environmental events. Conditions can also serve as stimuli for other organisms. Resources, on the other hand, are consumed by organisms in the course of their growth and reproduction. One organism may become or produce a resource for another through grazing, predation, parasitism or symbiosis, for example.

For any particular condition or resource, an organism may have a preferred value or set of values that favour its survival, growth and reproduction. Begon et al. (2006) define three characteristic curves, which show different "viability zones" for survival, growth and reproduction (Fig. 2.4).

In developing artworks, we can abstract these concepts significantly as long as we are clear about the functional relationships between conditions, resources and organism. From here on we will consider the organism as a "component" of an ecosystem, this more genetic term useful to remind us of the abstractions in play. Components may often be called "agents" in a computer simulation, typically representing autonomous entities with parameterised, possibly evolving, behaviours.

2.4.2 *Self-observation and Feedback*

Self-observation gives rise a type of feedback process, similar to a governor or more simply "rein control" (Harvey 2004). Here "observation" means the system monitoring of environmental conditions or resources that are necessary for reproduction, growth and survival and shifting its configuration in response. A component is causally coupled to the environment through relevant conditions or resources within its environment. Observation may be implicit or explicit, local or global. Observation forms a critical connection between a component's effect on the environment and its ability to modify its behaviour in response, typically to retain homeostasis in local conditions or resources. The use of the term "observation" is deliberately a loaded one. It is used in the cybernetic sense and does not imply a necessary concept of agency (although it does not preclude it). It might be considered the most simple precursor to more complex observational intelligence. It also suggests a system-level (as opposed to an individual-level) ontology that emerges through the interaction of system components.

The well-known model of planetary homeostasis, *Daisyworld*, uses a simple form of system level self-observation (Lenton and Lovelock 2001). Planetary albedo is affected by proportions of black and white daisies, whose relative proportions change according to surface temperature. What is fascinating about Daisyworld is its ability to maintain a homeostatic surface temperature while the incoming radiant heat energy increases.

In the ecosystem artwork *Colourfield* (McCormack 2007a), individual components ("agents") are bands of colour occupying a 1D lattice of cells. Genetic information controls the colour the agent produces, along with its preference to adapt to the colour of its neighbours and its propensity to occupy vacant neighbouring cells (thus making a larger contribution to the overall colour distribution). A feedback mechanism uses a colour histogram of the overall colour distribution to allocate resources to each individual agent on a per-time step basis (Fig. 2.5). Here the observation mechanism—resource allocation based on the image histogram—is implicit and global (the system as a whole is observing itself). An individual agent's contribution to the overall image influences the production of its own resources and those of others. The more cells an individual occupies, the greater the reliance of other individuals to it. Here feedback is an environmental reward function that favours symbiotic adaptations because of its global nature (resources are equally divided between cells). As the system is evolutionary, as a whole it has the ability to modify its colour composition and distribution in response to the "self-observation" provided by this feedback mechanism.

A different self-observation mechanism is in operation in the ecosystem artwork *Niche Constructions* (McCormack 2010). Niche construction is the process by which organisms, through their activities, modify their heritable environment (and potentially the environments of others). Advocates of niche construction theory in biology argue that it is an initiator of evolutionary change, rather than simply an evolutionary outcome (Odling-Smee et al. 2003). The complete set of conditions and resources affecting an organism represent its *niche*, which can be conceptualised as a hypervolume in n-dimensional space.

In the Niche Constructions artwork, evolutionary line drawing agents draw on an initially blank canvas as they move around. A set of normalised scalar values forms an agent's genome, which directs its behaviour over its lifetime. Individual alleles control rate of drawing curvature, "irrationality" (Fig. 2.6), fecundity and mortality. Agents die if they intersect with any previously drawn line or run off the page. The canvas is seeded with a small initial population of *founder agents*—initialised with uniformly distributed random genomes and positions—that proceed to move, draw and reproduce. There is no limit to the number of offspring an agent may have, but in general the lifespan of agents decreases as the density of lines becomes greater, because it is increasingly difficult to avoid intersection with existing lines. Eventually the entire population dies out and the image is complete. This finished drawing represents the "fossil record" of all the generations of lines that were able to live over the lifetime of the simulation.

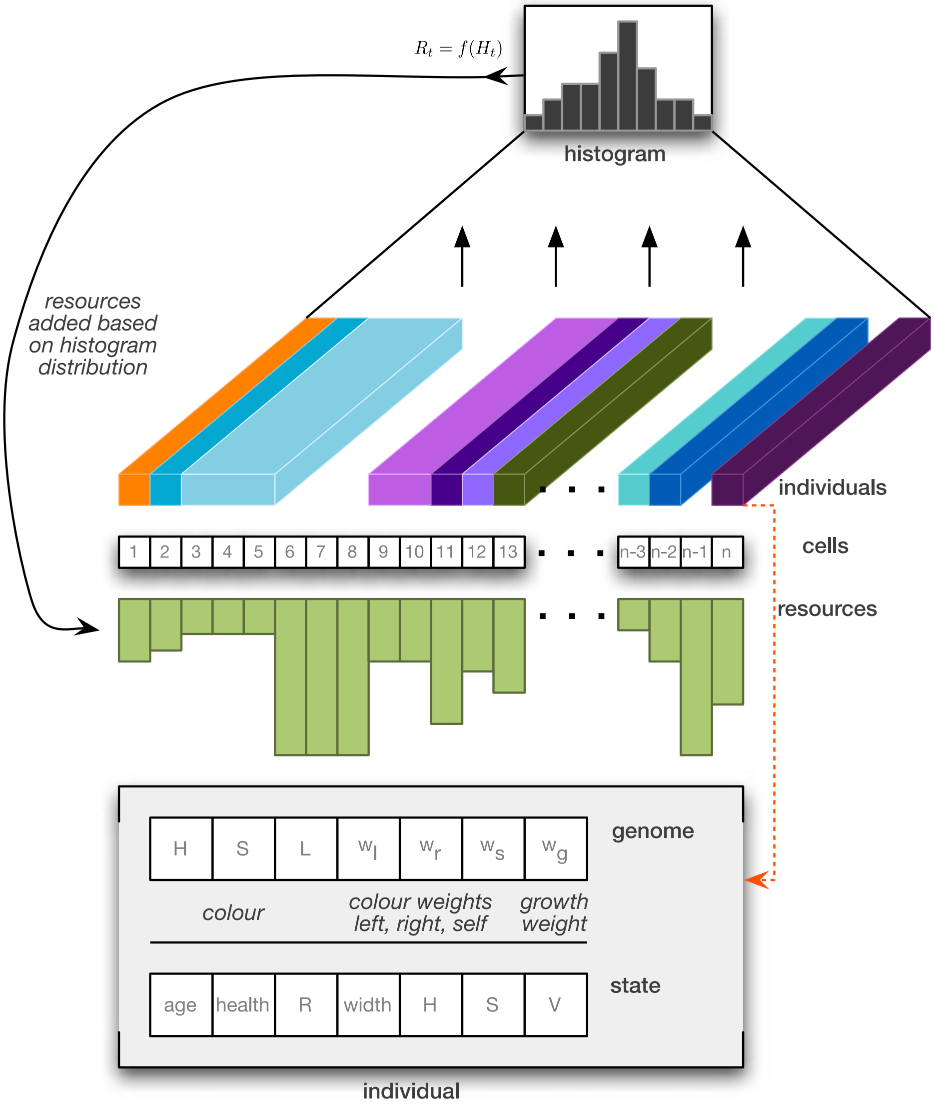

Fig. 2.5 Feedback relationships between component and environment creates a self-observation in the ecosystemic artwork "Colourfield"

Niche construction is enabled in this work through the addition of a self-observation mechanism that genetically links drawing behaviour to local conditions. As an individual agent draws on the canvas, the local density around it is measured. Each agent has an allele that represents its ideal density preference, i.e. the local line density that is most conducive to its survival, growth and reproduction. As the actual density shifts away from this ideal value, the agent finds it harder to reproduce, grow and survive. If the preferred density and actually density differ too greatly, the agent will die (see Fig. 2.7). Of course the actual value of this density preference is subject to evolutionary change and over the life of the drawing, average density preference increases in the population (McCormack 2010). The niche construction process influences agent behaviour: low density liking agents try and draw large, closed spaces to prevent other lines from decreasing their local density. High density seeking lines

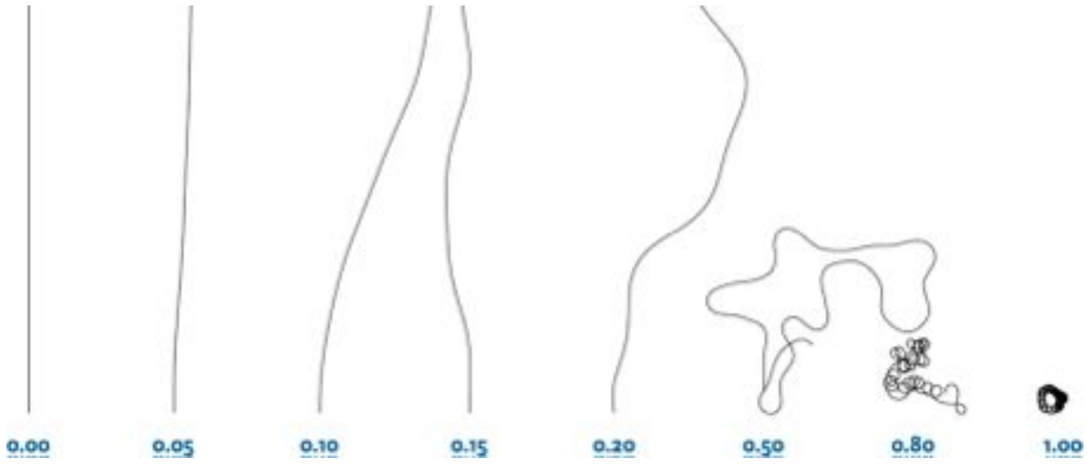

Fig. 2.6 Individual line drawing agents with different genetic values of irrationality. Note that the "die if intersect" rule has been turned off for these examples

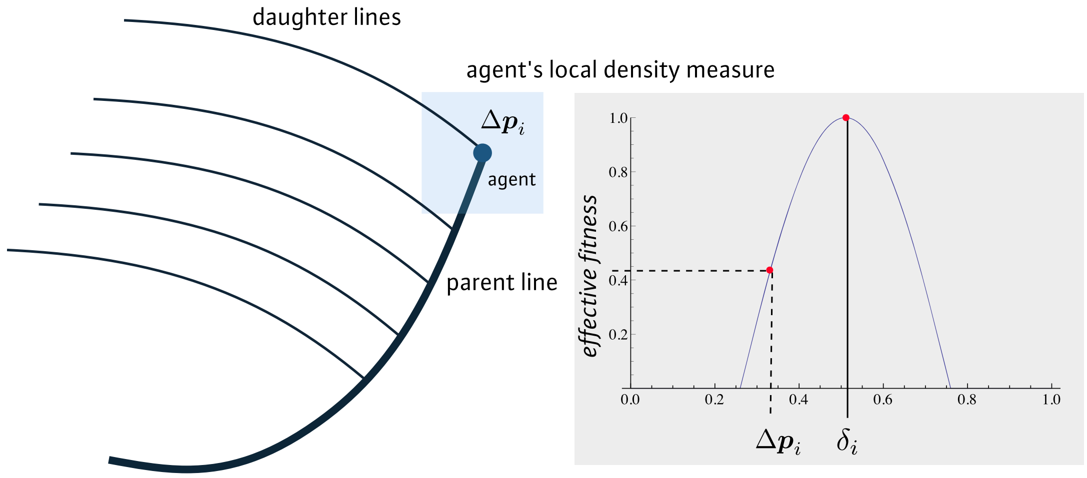

Fig. 2.7 The niche construction mechanism for drawing agents: a local line density measure, Δp_i, facilitates a self-observation mechanism. The agent's genome includes an allele that represents a preferred density (δ_i). The difference between preferred density and measured density affects the agent's effective fitness, hence its ability to survive, grow, and reproduce

give birth to large numbers of offspring, who quickly fill the canvas with lines of close proximity. Some examples are shown in Fig. 2.8.

This local, implicit self-observation plays a vital role in influencing the overall density variation and aesthetics of the images produced. We know this because turning the mechanism off produces images of significantly less density variation (statistically) and visual interest (subjectively).

2.4.3 Automation and the Creative Role of the Artist

> automation (noun): the use of largely automatic equipment in a system of manufacturing or other production process

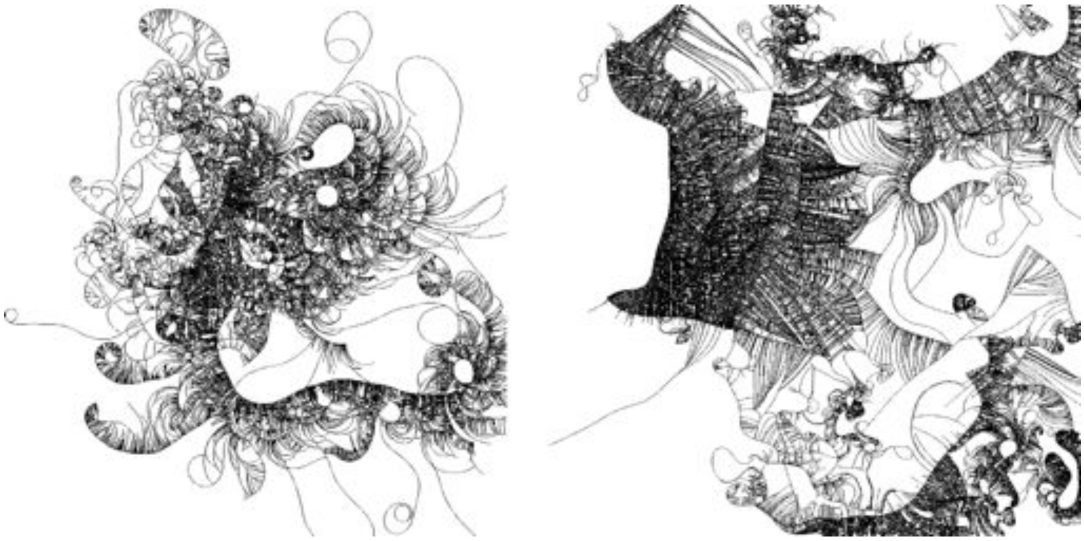

Fig. 2.8 Two sample outputs from the line drawing system with niche construction

The term “automation” originated in the USA, from the newly industrialised engineering of the 1940s, although similar concepts arose prior in different guises, both historically and geographically. The central idea was to create machines to perform tasks previously performed by humans. The rational was largely economic: machines that could replace and even out-perform their human counterparts will increase production efficiency. As a central driving force in US industrialisation and technologisation throughout the twentieth century, computers enabled the increasing sophistication and range of capabilities for automation within the capitalist economic system. The idea of machines automating human tasks still underpins many technology-driven approaches to “automating creativity”. Traditional AI or EC approaches seek the automation of aesthetic or creative optima finding. In contrast, the ecosystemic approach, as outlined here, does not seek to automate the human out of the creative process, nor claim to equal or better human creative evaluation and judgement. It views creative search and discovery as an *explorative* process, as opposed to an optimisation.

Ecosystemic processes recognise the importance of the link between structure and behaviour. Ecosystem components must be embedded in, and be part of, the medium in which they operate. The design of the system—components and their interdependencies—requires skill and creativity. This design forms the conceptual and aesthetic basis by which the outcomes can be understood. So rather than removing the artist by automating his or her role, the artist’s contribution is one of utmost creativity—creativity that is enhanced through interaction with the machine. As is also argued elsewhere in this book, forming an “ecosystem” that encompasses humans, technology and the socially/technologically mediated environment, opens up further ecosystemic possibilities for creative discovery.

There are of course, many reasons why we might seek some form of "automated creativity" or aesthetic judgement,[11] apart from replacing human labour. For example, automated creativity could lead to creative discovery that exceeds any human capability, or provides greater insights on the mechanisms of human creativity by attempting to model it. But these are "blue sky" speculations, and current technological advances in this area can just as easily homogenise and suffocate the creative decision-making process for human users, as they can expand or enhance it. A good example can be seen in recent digital camera technologies. Over the last ten years, as computational power has escalated, digital cameras have increasingly shifted creative decision making to the camera instead of the person taking the picture. We see modes with labels like "Intelligent Auto" or scene selection for particular scenarios ("Fireworks","Landscape","Sunset", "Beach"). These modes supposedly optimise many different parameters to achieve the "best" shot—all the photographer has to do is frame the image and press the button.[12] Recent advances even take over these decisions, choosing framing by high-level scene analysis and deciding when the picture should be taken based on smile detection, for example. Such functionality trends towards the removal of much human creative decision-making, subjugating the human photographer to an increasingly passive role.

As anyone who has used a entirely manual camera knows, hand-operated "slow technology" forces the user to think about all aspects of the photographic process and their implications for the final image. The user's role is highly active: experimentation, mistakes, and serendipitous events are all possible, even encouraged—well known stimuli for creativity. If the *design* of components and their interaction is good, then using such a device isn't marred by complexity or limited by inadequate functionality, which is often the rationalisation given in automation of creative functionality.

Shifting the thinking about the design of technology from one of "complexity automation" (where complexity is masked through "intelligent" simplicity) to one of "emergent complexity" (where interaction of well designed components generates new, higher-level functionality) allows the human user to potentially expand their creativity rather than have it subsumed and homogenised.

2.5 Conclusions

Ecosystemics represents an alternative, biologically-inspired approach to creative discovery over more traditional methods such as genetic algorithms or genetic programming. It offers an interesting conceptual basis for developing new creative systems and processes, even in non-computational settings. Incorporating an "environment", and allowing interactions between dynamic components and that environment, permits a rich complexity of creative possibilities for the artist wishing to

[11] Chapter 4 discusses this issue in more detail.

[12] Reminiscent of Kodak founder George Eastman's famous tag line of 1888 for the Kodak No. 1 camera: "You press the button, we do the rest".

exploit the generative nature of ecosystem processes. While ecosystemic methods don't offer a "magic bullet" in terms of searching the creative Klondike spaces of any generative system, they do make it easier to at least begin to conceptualise and design systems capable of high creative reward. As the complexity and sophistication of ecosystem artworks develop, we are likely to see further advances in the new creatively made possible with computers that use this approach.

Acknowledgements This research was supported by Australian Research Council Discovery Grants DP0877320 and DP1094064.

References

Aunger, R. (2002). *The electric meme: a new theory of how we think*. New York: Free Press.
Baluja, S., Pomerleau, D., & Jochem, T. (1994). Simulating user's preferences: towards automated artificial evolution for computer generated images. *Connection Science*, *6*, 325–354.
Basalla, G. (1998). *The evolution of technology*. Cambridge: Cambridge University Press.
Begon, M., Townsend, C., & Harper, J. (2006). *Ecology: from individuals to ecosystems*. New York: Wiley-Blackwell.
Bell, S. (1999). *Landscape: pattern, perception and process*. London: E & F N Spon.
Bentley, P. J., & Corne, D. W. (Eds.) (2002). *Creative evolutionary systems*. London: Academic Press.
Bird, J., Husbands, P., Perris, M., Bigge, B., & Brown, P. (2008). Implicit fitness functions for evolving a drawing robot. In M. Giacobini et al. (Eds.), *Lecture notes in computer science: Vol. 4974. Applications of evolutionary computing, EvoWorkshops 2008: EvoCOMNET, EvoFIN, EvoHOT, EvoIASP, EvoMUSART, EvoNUM, EvoSTOC, and EvoTransLog, Proceedings*, Naples, Italy, March 26–28, 2008 (pp. 473–478). Berlin: Springer.
Birkhoff, G. D. (1933). *Aesthetic measure*. Cambridge: Harvard University Press.
Boden, M. A. (2010). *Creativity and art: three roads to surprise*. London: Oxford University Press.
Bown, O., & McCormack, J. (2010). Taming nature: tapping the creative potential of ecosystem models in the arts. *Digital Creativity*, *21*(4), 215–231. http://www.csse.monash.edu.au/~jonmc/resources/DC2010/.
Brown, D. E. (1991). *Human universals*. New York: McGraw-Hill.
Dahlstedt, P. (2006). A mutasynth in parameter space: interactive composition through evolution. *Organised Sound*, *6*(2), 121–124.
Dawkins, R. (1999). *The extended phenotype: the long reach of the gene* (rev. ed.). Oxford: Oxford University Press.
De Landa, M. (2000). *A thousand years of nonlinear history*. Cambridge: MIT Press.
Di Scipio, A. (2003). 'Sound is the interface': from interactive to ecosystemic signal processing. *Organised Sound*, *8*(3), 269–277.
Dissanayake, E. (1995). *Homo aestheticus: where art comes from and why*. Seattle: University of Washington Press.
Dorin, A. (2001). Aesthetic fitness and artificial evolution for the selection of imagery from the mythical infinite library. In J. Kelemen & P. Sosík (Eds.), *LNAI: Vol. 2159. Advances in artificial life* (pp. 659–668). Prague: Springer. http://www.csse.monash.edu.au/~aland/PAPERS/aestheticFitness_ECAL2001.pdf.
Driessens, E., & Verstappen, M. (2008). Natural processes and artificial procedures. In P. F. Hingston, L. C. Barone & Z. Michalewicz (Eds.), *Natural computing series. Design by evolution: advances in evolutionary design* (pp. 101–120). Berlin: Springer.
Dutton, D. (2002). Aesthetic universals. In B. Gaut & D. M. Lopes (Eds.), *The Routledge companion to aesthetics*. London: Routledge. http://www.denisdutton.com/universals.htm.

Eiben, A. E., & Smith, J. E. (2003). *Introduction to evolutionary computing*. *Natural computing series*. Berlin: Springer.

Eldridge, A. C., & Dorin, A. (2009). Filterscape: energy recycling in a creative ecosystem. In M. Giacobini et al. (Eds.), *Lecture notes in computer science: Vol. 5484*. *Applications of evolutionary computing, EvoWorkshops 2009: EvoCOMNET, EvoENVIRONMENT, EvoFIN, EvoGAMES, EvoHOT, EvoIASP, EvoINTERACTION, EvoMUSART, EvoNUM, EvoSTOC, EvoTRANSLOG, Proceedings*, Tübingen, Germany, April 15–17, 2009 (pp. 508–517). Berlin: Springer.

Eldridge, A. C., Dorin, A., & McCormack, J. (2008). Manipulating artificial ecosystems. In M. Giacobini et al. (Eds.), *Lecture notes in computer science: Vol. 4974*. *Applications of evolutionary computing,EvoWorkshops 2008: EvoCOMNET, EvoFIN, EvoHOT, EvoIASP, EvoMUSART, EvoNUM, EvoSTOC, and EvoTransLog, Proceedings*, Naples, Italy, March 26–28, 2008 (pp. 392–401). Berlin: Springer.

Fuller, M. (2005). *Media ecologies: materialist energies in art and technoculture*. Cambridge: MIT Press.

Gamma, E., Helm, R., Johnson, R., & Vlissides, J. M. (1995). *Design patterns: elements of reusable object-oriented software*. *Addison-Wesley professional computing series*. Reading: Addison-Wesley.

Harvey, I. (2004). Homeostasis and rein control: from daisyworld to active perception. In J. B. Pollack, M. A. Bedau, P. Husbands, T. Ikegami & R. A. Watson (Eds.), *Ninth international conference on artificial life* (pp. 309–314). Cambridge: MIT Press.

Kaplinsky, J. (2006). Biomimicry versus humanism. *Architectural Design*, *76*(1), 66–71.

Keane, A. J., & Brown, S. M. (1996). The design of a satellite boom with enhanced vibration performance using genetic algorithm techniques. In I. C. Parmee (Ed.), *Conference on adaptive computing in engineering design and control 96, P.E.D.C.* (pp. 107–113).

Koren, L. (2010). *Which "Aesthetics" do you mean?: ten definitions*. Imperfect Publishing.

Lenton, T. M., & Lovelock, J. E. (2001). Daisyworld revisited: quantifying biological effects on planetary self-regulation. *Tellus*, *53B*(3), 288–305.

Luke, S. (2009). *Essentials of metaheuristics*. Lulu Publishing, Department of Computer Science, George Mason University

Lumsden, C. J. (1999). Evolving creative minds: stories and mechanisms. In R. J. Sternberg (Ed.), *Handbook of creativity* (pp. 153–169). Cambridge: Cambridge University Press. Chap. 8.

Machado, P., & Cardoso, A. (2002). All the truth about NEvAr. *Applied Intelligence*, *16*(2), 101–118.

Machado, P., Romero, J., & Manaris, B. (2008). Experiments in computational aesthetics. In J. Romero & P. Machado (Eds.), *The art of artificial evolution: a handbook on evolutionary art and music* (pp. 381–415). Berlin: Springer.

Martindale, C. (1999). Biological bases of creativity. In R. J. Sternberg (Ed.), *Handbook of creativity* (pp. 137–152). Cambridge: Cambridge University Press. Chap. 7.

McCormack, J. (2001). Eden: an evolutionary sonic ecosystem. In *Lecture notes in computer science: Vol. 2159*. *Advances in artificial life, proceedings of the sixth European conference, ECAL* (pp. 133–142).

McCormack, J. (2005). On the evolution of sonic ecosystems. In *Artificial life models in software* (pp. 211–230). London: Springer. http://www.springeronline.com/sgw/cda/frontpage/0,11855,5-40007-22-39144451-0,00.html.

McCormack, J. (2007a). Artificial ecosystems for creative discovery. In *Proceedings of the 9th annual conference on genetic and evolutionary computation (GECCO 2007)* (pp. 301–307). New York: ACM.

McCormack, J. (2007b). Creative ecosystems. In A. Cardoso & G. Wiggins (Eds.), *Proceedings of the 4th international joint workshop on computational creativity* (pp. 129–136).

McCormack, J. (2008a). Evolutionary L-systems. In P. F. Hingston, L. C. Barone & Z. Michalewicz (Eds.), *Natural computing series*. *Design by evolution: advances in evolutionary design* (pp. 168–196). Berlin: Springer.

McCormack, J. (2008b). Facing the future: evolutionary possibilities for human-machine creativity. In P. Machado & J. Romero (Eds.), *The art of artificial evolution: a handbook on evolutionary art and music* (pp. 417–451). Berlin: Springer.

McCormack, J. (2010). Enhancing creativity with niche construction. In H. Fellerman et al. (Eds.), *Artificial life XII* (pp. 525–532). Cambridge: MIT Press.

McCormack, J., Dorin, A., & Innocent, T. (2004). Generative design: a paradigm for design research. In J. Redmond, D. Durling, & A. de Bono (Eds.), *Futureground, vol. 1: abstracts, 2: proceedings* (p. 156). Melbourne: Design Research Society.

Murray, S. (2011). Design ecologies: editorial. *Design Ecologies*, *1*(1), 7–9.

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

Perkins, D. N. (1996). Creativity: beyond the Darwinian paradigm. In M. Boden (Ed.), *Dimensions of creativity* (pp. 119–142). Cambridge: MIT Press. Chap. 5.

Ramachandran, V. S. (2003). The emerging mind. In *Reith lectures; 2003*. London: BBC in association with profile Books.

Ramachandran, V. S., & Hirstein, W. (1999). The science of art: a neurological theory of aesthetic experience. *Journal of Consciousness Studies*, *6*, 15–51.

Romero, J., & Machado, P. (Eds.) (2008). *The art of artificial evolution: a handbook on evolutionary art and music*. *Natural computing series*. Berlin: Springer.

Scheiner, S. M., & Willig, M. R. (2008). A general theory of ecology, *Theoretical Ecology 1*, 21–28.

Shapshak, T. (2011). Why Nokia got into bed with Microsoft. http://www.bizcommunity.com/Article/410/78/57030.html.

Staudek, T. (2002). *Exact aesthetics. Object and scene to message*. PhD thesis, Faculty of Informatics, Masaryk University of Brno.

Svangåard, N., & Nordin, P. (2004). Automated aesthetic selection of evolutionary art by distance based classification of genomes and phenomes using the universal similarity metric. In G. R. Raidl, S. Cagnoni, J. Branke, D. Corne, R. Drechsler, Y. Jin, C. G. Johnson, P. Machado, E. Marchiori, F. Rothlauf, G. D. Smith & G. Squillero (Eds.), *Lecture notes in computer science: Vol. 3005*. *EvoWorkshops 2004* (pp. 447–456). Berlin: Springer.

Takagi, H. (2001). Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation. *Proceedings of the IEEE*, *89*, 1275–1296.

Tansley, A. G. (1939). British ecology during the past quarter-century: the plant community and the ecosystem. *Journal of Ecology*, *27*(2), 513–530.

Waters, S. (2007). Performance ecosystems: ecological approaches to musical interaction. In *EMS07—the 'languages' of electroacoustic music*, Leicester.

Willis, A. J. (1997). The ecosystem: an evolving concept viewed historically. *Functional Ecology*, *11*(2), 268–271.

Wilson, S. W. (1999). *State of XCS classifier system research* (Technical report). Concord, MA.