

Creative Ecosystems

Jon McCormack

Centre for Electronic Media Art
Clayton School of IT, Monash University
Clayton 3800, Australia

Jon.McCormack@infotech.monash.edu.au

Abstract

This paper addresses problems in computational creative discovery, either autonomous or in synergetic tandem with humans. A computer program generates output as a combination of base primitives whose interpretation must lie outside the program itself. Concepts of combinatoric and creative emergence are analysed in relation to creative outputs being novel and appropriate combinations of base primitives, with the conclusion that the choice of the generative process that builds and combines the primitives is of high importance. The generalised concept of an artificial ecosystem, which adapts concepts and processes from a biological ecosystem at a metaphoric level, is an appropriate generative system for creative discovery. The fundamental properties of artificial ecosystems are discussed and examples given in two different creative problem domains. Systems are implemented as pure simulation, and where the ecosystem concept is expanded to include real environments and people as ecosystem components, offer an alternative to the ‘software tool’ approach of conventional creative software.

Keywords: Artificial ecosystems, Combinationalism, Emergence.

“Theories are important and indispensable because without them we could not orientate ourselves in the world — we could not live. Even our observations are interpreted with their help.”

— Karl Popper, *The Myth of the Framework*

1 Introduction

We are interested in problems of *computational creative discovery* where computer processes assist in enhancing human creativity or may autonomously exhibit creative behaviour independently. The intention is to develop ways

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.

©2007 Goldsmiths, University of London

of working with technology that achieve creative possibilities unattainable from any existing software tools or methods. These goals will be addressed here in the context of artistic creation, however the results may be applicable to many forms of creative discovery.

Darwinian evolution has been described as the only theory with the “explanatory power for the design and function of living systems...accounting for the amazing diversity and astonishing complexity of life” (Nowak, 2006). Evolutionary synthesis is a process capable of generating unprecedented novelty, i.e. *it is creative*. It has been able to create things like prokaryotes, eukaryotes, higher multicellularity and language through a non-teleological process of replication and selection. We would like to adapt, on a metaphoric level, the mechanisms of biological evolution in order to develop new approaches to computational creativity. In Biology, the physical processes of replication and selection take place in an environment, populated by species that interact with and modify this environment, i.e. an ecosystem. Processes from biological ecosystems serve as inspiration for computational artificial ecosystems. The aim is to structure these artificial ecosystems in such a way that they exhibit novel discovery in a creative context rather than a biological one.

We consider creativity in terms that it involves the generation of something *novel and appropriate* (i.e. unexpected, valuable) to the particular aesthetic domain. Van Langen et al. conclude the necessary conditions for any artificial creative system must be the ability to interact with its environment, learn, and self-organise (van Langen et al., 2004). In this paper, the aim is for creative discovery by machines, or humans and machines working synergistically, rather than a computational model of human creativity or knowledge-based models for a particular domain.

Before looking at how artificial ecosystem concepts can be used as processes for creative discovery, the next section examines how such processes fit into computational creative discovery in general.

2 Combinationalism

A major controversy regarding computational creativity relates to the concept of ‘combinationalism’: the understanding that “creativity is the creative combination or re-

combination of previously existing elements” (Dartnall, 2002). This understanding is based on the intuition that one cannot create something new from nothing, hence we require a “combination or recombination of what we already had” — the opposing view being that creativity begins with knowledge, skill and abilities, and emerges from these faculties through interaction with the environment. The challenge is to account for how these cognitive properties give rise to creative output (McCormack, 2005b).

Clearly, many creative outputs are indeed a combination of basic primitives organised in a new way. Let us consider an arbitrary system that generates some creative output from a fixed set of indivisible, distinct primitives (basic building blocks, fundamental units). We will call this set of n distinct primitives V , i.e.: $V = \{p_1, p_2, \dots, p_n\}$. A generative process, G selects elements from V to make $S \in V^r$, an output composed by some permutation of primitives from V . We will assume that:

- the ordering of primitives in S is important;
- repetitions of primitives are permitted;
- The size of S is fixed¹ and $|S| = r$, where $r > 0$.

The process of generating S from V by G is denoted:

$$V \xrightarrow{G} S$$

Denote each specific possibility, S_i , $i = 1, 2, \dots, n^r$ (since there are n^r possibilities for S) and $Q \in S^* = \{S_1, S_2, \dots, S_{n^r}\}$ the set of all possible outputs. Further, let us define $Q_G \subseteq Q$ the set of all outputs generated by G . The *conceptual space*, C , is defined as the base primitives, V and the rules for combining them, i.e.: $C = \langle V, G \rangle$.

As a simple example, let us suppose V is a set of musical notes, i.e. $V = \{A, B, C, D, E, F, G\}$ and $r = 12$, so each S is a 12 note melody composed from the notes in V . In this case $n^r = 13,841,287,201$. Clearly, for non-trivial problems the number of possibilities for S is very large, in many cases beyond astronomical proportions such as the estimated number of particles in the universe.

This vast space of potential combinatorial possibilities for S illustrates why such systems are said to display *combinatoric emergence*, that is, configurations generated by G appear to express new properties or structures not found in the individual primitive components p . Note that such new properties or structures are generally *observed*, not defined quantitatively (Baas, 1994; Dorin and McCormack, 2002).

While the potential output generated by G may be vast, any individual output S can *only* be composed of elements from V . In the case of our musical example, we could generate a large number of melodies from V , but none of those melodies could contain the note $C\sharp$, for example, because it is not a member of V .

2.1 Creative Emergence

In the case of what is termed *creative emergence*, it is proposed that fundamentally new primitives enter the system,


¹Arbitrary size outputs are possible by incorporating an empty primitive into V , i.e. $V \cup \{\emptyset\}$.

opening up a new set of possibilities that were not previously possible (Cariani, 1991, 1997). In more formal terms, this process modifies the conceptual space:

$$C \rightsquigarrow C^\Lambda$$

Where C^Λ is the new conceptual space. According to Bird (2004), in an analogy with letters generating words, creative emergence “involves expanding the alphabet of letters by transforming the underlying generative system as well as combining the letters into new words” (Fig. 1). In the terminology used in this paper, creative emergence can introduce new members into V , i.e.: $V \Rightarrow V^\Lambda$. The introduction of new primitives in V would by necessity involve some transformation of G , since by definition G only knows how to generate things from the original V .

COMBINATORIC EMERGENCE



CREATIVE EMERGENCE




Figure 1: Combinatoric and creative emergence (redrawn from (Bird, 2004))

A computational system that combines primitives must provide a semantic interpretation for the members of V . For example, the symbol ‘A’ must be interpreted as a musical note before it can represent music. It is easy to generate additional primitive *symbols* that can be added to V , but seemingly impossible to computationally discover new interpretations for those symbols, because the interpretation of those symbols is done outside of the software itself (by a listener in the case of a musical example).

There are two conclusions to be drawn from this discussion. The first, rather obviously, is that in any combinatorial system, you will only get combinations of the base primitives for which you provide an *a priori* interpretation. The knowledge of how to interpret symbols is provided by the programmer, not the program (the computer can only differentiate one symbol from another).

The second point is that a combinatorial approach is still a useful one if we get our base primitives right. As we have seen, the scope of possibilities is very large in any practicable system. Composers, for example, seem in the main content with composing from a fixed set of base primitives. Architects can design great architecture from a fixed set of building materials. Any digital image can be made by combining pixels in the right order.

3 Generative processes for Creative Discovery

At this point, we have said nothing about the quality or utility of G and the output it generates. Having a vast range of possibilities in a combinatorial system represents only a *potential* for actually finding good combinations. It is trivial to construct a generative process, G that can generate all the possible members of Q (i.e. $Q_G = Q$). Each combination S_i generated by G will be new, so finding novelty is not the problem, it is finding *appropriate* novelty. The exponential expansion of possibilities, dependent on n and r , means that for any non-trivial system, brute-force methods such as an iterative or random search will not be practical.

Ideally, we would like a generative process that finds the creatively interesting combinations and avoids the uninteresting ones. For many domains the proportion of what we might be inclined to call “interesting” is likely to be extremely small. Randomly sampling individual members from Q is not, in general, a useful strategy for finding the appropriate members of that set.

If our approach is to relegate creative discovery to being a search or optimisation problem, then a number of general algorithms already exist for this task, e.g. (Michalewicz and Fogel, 1999). One popular choice has been the use of Evolutionary Computing (EC) methods, such as genetic algorithms, evolution strategies or genetic programming.

Standard EC methods require an explicit evaluation of *fitness*, that is a comparative ranking between possible solutions in order to determine the composition of the population for the next generation. For creative discovery, this is a difficult problem for two reasons:

- evaluation of the quality of creative output is highly subjective and context dependent, relying on much domain specific knowledge that is difficult to quantify;
- the type of knowledge and evaluation necessary depends specifically on the creative task or activity being simulated, i.e. it is difficult to generalise or abstract.

It is for these reasons (and many others) that machine representable fitness functions for “creativity”, or “aesthetics” have largely unsuccessful (though not for want of trying, e.g. (Birkhoff, 1933)).

Evaluation of subjective criteria is relatively easy for humans, so a natural approach incorporate human evaluation of fitness into the algorithm. *Interactive Evolution* (also know as *aesthetic selection* or *aesthetic evolution*) have found wide application and popularity for a variety of problems in creative discovery (Takagi, 2001). In this approach, the problem of finding machine-representable fitness functions for aesthetic or subjective properties is circumvented in favour of human fitness evaluation and ranking. While this is a popular method, it is not without significant problems (Dorin, 2001; McCormack, 2005b). These problems include: difficulty in fine-grained evaluation; limited population sizes; slow evaluation times; poor balancing between exploration and exploitation (one

of the GA’s main benefits as a search method (Eiben and Smith, 2003, p. 29)).

The central question addressed by this paper, then, is this: in a combinatorial system, how can we search and optimise using EC techniques without an explicit fitness evaluation, either by human or machine? That is, what kinds of processes, G are best suited to creative discovery from a combinatorial system? The answer proposed here is through the use of an *artificial ecosystem* approach. This approach is detailed in the following sections.

4 Artificial Ecosystems

The design of environments from which creative behaviour is expected to emerge is at least as important as the design of the individuals who are expected to evolve this behaviour. The *Artificial Ecosystem* as a generalised evolutionary approach for creative discovery. Natural ecosystems exhibit a vast array of complex phenomena, including homeostasis, food-webs, wide causal dependencies and feedback loops, even (controversially) evolution at the ecosystem level (Swenson et al., 2000). Species within the ecosystem compete for resources in order to survive and reproduce. Typical co-operative and competitive evolutionary strategies are observed, such as mutualism, symbiosis, predation and parasitism. To be glib, it could be said that the ecosystem has a lot of interesting features going for it. We would like to harness some of these features for the purposes of *creative discovery* — the discovery of novelty in a system without explicit teleology.

The concept of an artificial ecosystem used here is formative and based on abstractions of selected processes found in biology. We are interested in developing general algorithms for creative discovery. These algorithms are based on dynamic evolutionary processes observed in biological ecosystems. Just as genetic algorithms are not a simulation of natural selection, the artificial ecosystem algorithms presented here are not intended to simulate real biological ecosystems. The ecosystem is viewed as a dynamic, complex system, essential for selection and a driving force behind biological novelty when established with the appropriate conditions. We would like to harness the novel potential of ecosystem processes at a metaphoric level and apply them to creative processes of interest to humans.

4.1 Simulated Ecosystem Studies

Simulated artificial ecosystems have been well studied in the sciences. A number of artificial life models employ the concept of an abstract or simplified ecosystem. This concept of the artificial ecosystem was introduced in (Conrad and Pattee, 1970). A population of independent software agents interact within a programmer-specified artificial physics and chemistry. Agent interaction is simplistically analogous to that which occurs in a real ecosystem. Agents must gain sufficient resources from their environment in order to survive and reproduce. Typically, a number of successful survival strategies will emerge (niches) often with inter-dependencies between individual species (e.g. symbiosis and parasitism). Similar artificial

ecosystem methods have been useful in modelling problems in economics (Arthur et al., 1997), ecology (Mitchell and Taylor, 1999) and social science (Epstein and Axtell, 1996).

The majority of such systems focus on single-niche, homogeneous environments, and operate at evolutionary time-scales, simulating the evolution a single species over time. This focus, and the use of minimal, broad assumptions is primarily for the purposes of verification and validation (Adami, 2002). Artificial life agents adapt their behaviour through an evolutionary process to best fit their (typically homogeneous) environment.

Ecological models, on the other hand, tend to operate on far smaller time scales, simulating periods typically ranging from hours to several decades, with a focus on *fitness seeking* through organisational changes or behavioural adaptation of an individual species. This level of simulation reflects the practical questions asked by ecologists in relation to real ecosystems, whereas artificial life research tends to focus on abstract evolutionary dynamics. Important to both styles of investigation is the emergence of macro phenomena or properties from micro interactions. The micro interactions (typically interacting agents) being formally specified in the model; the macro properties an emergent outcome of the simulation.

4.2 Processes for Artificial Ecosystems

In many artificial ecosystem models, the designers of the model are driven by specific applications or outcomes, so the mechanisms, abstractions and terminology differ between systems. This section attempts to define both properties and concepts for general artificial ecosystems. They are positioned at a “middle level” of abstraction: for example an individual is an indivisible unit, it is not represented as a combination of self-organising sub-units, even though this might be possible. In any agent or individual-based model there is always a conflicting tension between model complexity, model validation and simulation outcomes. In contrast to ecological models, the focus of creative discovery is on the suitability and sophistication of creative outcomes, not the verification of models with empirical data or their validation in terms of answering questions not explicit in the original model (Grimm and Railsback, 2005). This allows us some creative licence in our interpretation, but we would still hope for some (at least) semi-formal validation of any general ecosystem models for creative discovery.

While not an essential characteristic of ecosystem models, the use of evolution and the operation on evolutionary time scales is an assumption of the ecosystem models proposed here. This does not preclude the possibility of the model operating at other time scales.

The basic concepts and processes for artificial ecosystems are:

- the concepts of *genotype* and *phenotype* as used in standard EC algorithms. A genotype undergoes a process of *translation* to the phenotype. The genotype and phenotype form the basis of an *individual* in the model;
- a collection of individuals represent a *species* and the

system may potentially accommodate multiple, interacting species;

- spatial distribution and (optionally) movement of individuals;
- the ability of individuals to modify and change their environment (either directly or indirectly as a result of their development within, and interaction with, the environment);
- the concept of individual *health* as an abstract scalar measure of an individual’s success in surviving within its environment over its lifetime;
- the concept of an individual *life-cycle*, in that an individual undergoes stages of development that may affect its properties, physical interaction and behaviour;
- the concept of an *environment* as a physical model with consistent physical rules on interaction and causality between the elements of the environment;
- an *energy-metabolism resource model*, which describes the process for converting energy into resources that may be utilised by species in the environment to perform *actions* (including the production of resources).

For populations to evolve, there must be some kind of *selection pressure* that implicitly gives some species a higher reproduction rate than others, creating an implicit measure of fitness (Nowak, 2006, Chapter 2). Let us assume any given environment has finite resources and a total population carrying capacity, κ . Species compete for finite resources. These resources are used by individuals to better their reproductive success, until the total population reaches κ . Hence, those able to discover successful strategies for efficiently exploiting those resources are able to reproduce at a higher rate, dominating the population. In contrast to EAs with explicit fitness functions, selection is implicit: successful strategies (individuals) emerge in response to the challenges set by the environment. Moreover, in locating and processing resources, species may alter the environment itself. In this case, adaptation is a dynamic process involving feedback loops and possibly delicate balances.

Individuals maintain a scalar measure of “health” which indicates the success of the individual during its lifetime. This is roughly akin to a fitness measure in traditional EC algorithms. If the health level of an individual falls to zero, the individual dies and is removed from the population (normally returning its resources to the environment). Health is normally affected by the individual’s ability to acquire resources from the environment (which may include other individuals). Other internal factors, such as age, may also change an individual’s health measure.

In the context of problem solving, individual species may represent competing or co-operating parts of a global solution. This is highly suitable when many different combinations of components may form equally good solutions (e.g. notes or phrases forming a musical composition). When using standard EC methods for search or optimisation, the challenge faced is in choosing appropriate




Figure 2: Schematic overview of *Colourfield*

genotype representations, selection methods, and fitness functions. The challenge for artificial ecosystems is in the design of environments and the interaction of species within them.

An example of a simple artificial ecosystem model follows.

4.3 Colourfield

Colourfield is a simple one-species ecosystem of colour patterns. It consists of a one-dimensional discrete world of fixed width, populated by individuals (Fig. 2). Each space in the world is called a *cell* and may be occupied by at most one individual. Individuals occupy one or more cells and are represented visually as lines of colour. A population of individuals produces a field of one or more colours.

An individual’s genome is a fixed-length array of real numbers representing: the natural colour (hue, saturation, lightness: HSL); propensity to change to the natural colour, and to the colour of the individual to the left and right of this individual (a normalised weight); propensity to grow into empty neighbouring cells. Each individual in the population maintains a separate *state*, which consists of: the age of the agent, health, current resources held, number of cells currently occupied, and current colour.

All individuals begin with no colour (black) and attempt to acquire resources to reach their target colour (a weighted sum, as determined by the genome, of the natural colour and the current colours of neighbours). Resources are required to change and maintain a particular colour, proportionate to the rate of change. If a neighbouring cell is empty, the individual may “grow” into that cell, the propensity to grow determined by the genome. The more cells occupied, the more resources are required to change colour, but the greater the contribution to the overall *colour histogram* of the world (detailed shortly).

Let the current colour of individual i in RGB colour space be the vector $C_i = (r_i, g_i, b_i)$ and the width w_i . The resources required by the individual are:

$$r_i = w_i^2 \left(k_0 + k_1 \log \left(\frac{d\|C_i\|}{dt} \right) \right) + k_2 \frac{dw_i}{dt},$$

where k_0 , k_1 and k_2 are constants.

Individuals receive resources from the environment via a feedback process based on the composition of the world. At each timestep, a histogram of chroma and intensity values for the world is built. This histogram, H_t is used as a basis for delivering resources to the world. A total resource R_t for the timestep t , is calculated via a function $f: \mathbb{R}^n \rightarrow \mathbb{R}$:

$$R_t = f(H_t)$$

and then distributed equally to all the cells in the world, e.g.:

$$r_{k,t+1} = r_{k,t} + \frac{R_t}{n}, \quad k = 1, 2, \dots, n$$

where n is the size of the world. Individuals that occupy more cells therefore receive a greater amount of resources, as they make a greater contribution to the histogram.

A number of different versions of the function f have been tested. These include: favouring chroma values with peaks at equal division, maximising chroma or intensity variation; matching a normal distribution; matching histograms based on paintings recognised for their skilful use of colour.

Given sufficient resources, and following a period of “growth” an individual may reach its desired colour and width (which may be dependent on the individual’s neighbour states). At this time, it may choose to reproduce, either by crossover with an immediate neighbour, or — if there are no neighbours — by mutation. In the case of two immediate neighbours, the mating partner is selected with probability weighted to the normalised Euclidean distance between the colour of the individual and its neighbours, so individuals are more likely to mate with others who produce colours similar to themselves. Offspring are placed in the nearest empty cell, or if none exists, they replace parent cells. If there are insufficient resources, the agent is unable to maintain its target colour, causing it to fade and eventually die.

Over time, the system evolves to maximise the production of resources according to the composition of the histogram, which is determined by the size and colour of all the individuals in the world. The system exhibits novel colour patterns with patterns of stasis followed by large-scale change as new optimal configurations are discovered. Due to the configuration of co-dependencies, *Colourfield* exhibits classic ecosystem phenomena such as parasitism (a rogue colour contributing little to resource production but “feeding off” other resource producing colours) and mutualism (co-operative combinations of colours mutually contributing to high resource production).

Colourfield is a simple experiment in adapting ecosystem concepts to a simple creative system. It demonstrates creative discovery in a limited domain (creative relationships between fields of colour).

4.4 Tools and Ecosystems

The concept of an ecosystem as a mechanism for creative discovery is not limited to the simulation of ecosystems within the computer. In a creative context it is useful to consider human-machine interaction as forming an ecosystem, replacing the concept of machine as creative tool. This discussion is similar to that used by Di Scipio (2003) and the approach used in the design of the *Eden* system, discussed in Section 4.5.

Humans have always worked with tools. Physical tools are useful because: (i) they enable a manipulation of the environment (a chisel sculpts wood); (ii) their constraints focus the user to their proper function (a pencil is used for drawing on surfaces); and (iii) their organisation encompasses knowledge (we cannot imagine in our mind the correct positioning of a slide rule to evaluate the multiplication of two numbers, yet by physically using a real slide rule it is easy).

Today, computer use is widespread in many areas of creative production, but this use is almost exclusively in the role of “computer as a tool”. Moreover, many of the metaphors used by software tools borrow from physical counterparts or historical lineage (e.g. Adobe Photoshop is a “digital darkroom”, Paint programs use a “virtual paint brush”, etc.). Often these metaphors are poorly translated or simply lack the physicality of their real counterparts (playing a “virtual piano” is just not as good as the real thing).

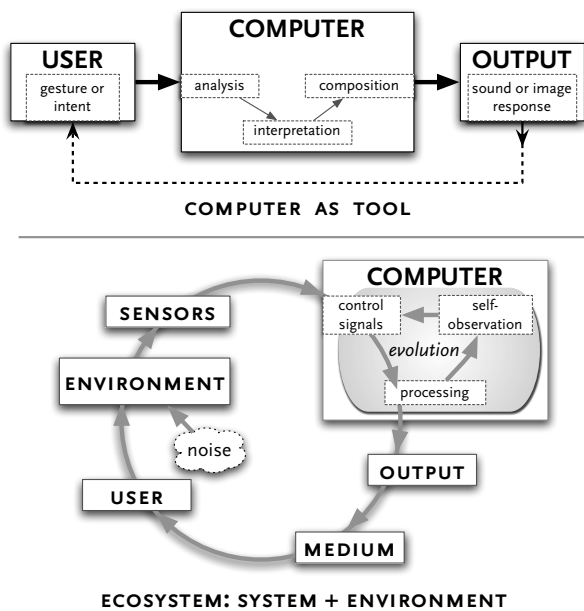


Figure 3: Computer use as a tool (top) and as part of an ecosystem (bottom)

The ecosystem approach does not conceptualise the machine as an object. Rather, the *processes*, both internal and external, are conceived as interdependent, connected *components*, which self-organise into a *system*. Components innately seek to replicate themselves within a noisy environment. With a limited carrying capacity, those com-

ponents best able to fit the environment dominate the population.

When the system interacts with the environment, it forms an ecosystem. The environment is the medium in which the system develops, and in a creative context may be the creative medium itself (e.g. sound, light, 3D form, and so on). In this mode of working, interdependent processes form an evolutionary, dynamical system, with adaptive behaviour to environmental conditions including the ability to interfere with, and modify, the environment. The machine becomes a synergistic partner in a collaborative creative process, as opposed to a passive tool manipulated by a user. As shown in Fig. 3, the computer, the physical environment and the user all form part of a coupled feedback system.

The powerful properties of tools outlined above are still preserved in the ecosystems scenario, along with additional features not normally associated with the human creative use of tools:

1. *Manipulation of the environment*: components are able to manipulate their environment, moreover due to the recursive coupling (Ashby, 1952) between system and environment we gain additional properties such as homeostasis (the ability for self-maintenance of particular dynamic configurations in changing external conditions) and system ‘memory’ through environmental modification.
2. *Constraints are created by the environment*: evolutionary adaptations are fitness seeking, leading to novel solutions imposed by the constraints, not determined by explicit fitness functions as is the case with conventional EC methods.
3. *Organisation encompasses knowledge*: the dynamic configuration of system components represents the knowledge of the system. As this configuration is dynamic and adaptive, the system is able to ‘learn’.

We are interested in new properties and interactions being indirectly implemented: arising as emergent by-products of carefully designed interdependencies between system components.

There are three important considerations in this interactive ecosystem approach to creativity: (i) the design of the individual system components and their interdependencies; (ii) the metaphors used in interpreting the function of components and their dependencies; and (iii) the composition of the environment in which the system interacts. A careful analysis of these considerations remains on-going research.

4.5 Eden: an evolutionary sonic ecosystem

Eden is a artwork installation that makes extensive use of the concepts discussed in this paper. The details presented here focus on the ecosystem aspects of the work. For detailed technical descriptions, see (McCormack, 2001, 2005a).

The work consists of a complex artificial ecosystem running in real-time on a two-dimensional lattice of cells. This world is projected into a three-dimensional environment, approximately 6m x 6m (see Fig. 4). The ecosystem




Figure 4: Installation view of *Eden*

consists of three basic types of matter: rocks, biomass, and evolving agents. If a rock occupies a cell, agents or biomass may not. Agents attempting to move into a cell occupied by a rock will “feel” pain and suffer energy loss.

Biomass provides a food source for the agents. Biomass is modelled on an extended *Daisworld* model (Lenton and Lovelock, 2001), with growth rate, β_i for individual biomass element i , a Gaussian function of local temperature at the location (x, y) of the element, $T_{x, y}$:

$$\beta_i = e^{-0.01(22.5 - T_{x, y})^2}.$$

The *Eden* world exists on an imaginary, Earth-like planet, orbiting a sun with a period of 600 days. The orbit eccentricity and polar orientation result in seasonal variations of temperature, thus affecting biomass growth. As with Lenton and Lovelock’s model, the system exhibits self-regulation and stability under a range of conditions. However, overpopulation by agents may reduce biomass to negligible levels, resulting in a temperature increase. The increased temperature lowers the growth rate of the biomass, leading to agent extinction and a dead planet. The system detects such conditions, at which time the planet is “rebooted” to initial conditions and a fresh batch of agents and biomass seeded into the world.

Agents are oriented, omnivorous, autonomous, mobile entities with a collection of sensors and actuators controlled by a learning system, based on classifier systems (a version of Wilson’s XCS (Wilson, 1999)). Agents are able to metabolise biomass into energy, which is required to perform *actions* via the agent’s actuators. Possible actions include: eating, resting, moving, turning left or right, singing, attacking whatever occupies the cell in front of the agent, mating. The energy cost of these actions varies according to the action (attacking costs more energy than resting, for example), and to physical factors, such as the mass of the agent (mass also increases the power of attacking — a big, heavy agent is more likely to injure or kill a smaller agent). If an agent’s energy (health) level falls to 0, the agent dies. Dead agents may be eaten by other agents for a certain time period following death.

Agent sensors are both internal (enabling introspection) and external (enabling sensation of the environment). They include: sensation of cell contents within the Moore neighbourhood of the agent; sound intensity

and frequency arriving at the agent’s location according to a simple physical model; introspection of pain; introspection of low energy (health). The LCS evolves sets of rules based on past experience and performance of successful rules. At regular periods the agent’s health and resource acquisition differentials are examined and a credit or penalty is provided to those rules used since the previous evaluation. A positive differential pays credit proportional to its magnitude, likewise a negative differential penalises. Successful rules gain credit and so are more likely to be selected in the future. Rules that consistently receive penalty are eventually removed.

Rules evolve during an agent’s lifetime, with a penalty imposed on energy for large rule sets to encourage efficiency. Two agents may mate — the resultant offspring inherit the most successful rules of their parents, hence the system uses *Lamarckian evolution*.

The *Eden* environment is visualised and sonified in the installation space. The two-dimensional world is projected onto two translucent screens, configured in an ‘X’ shape. This enables people experiencing the work to move freely around the screens at close range, examining details of the world as it updates in realtime. The sounds made by the agents are spatially mapped to four speakers located at the two corners of each screen. This rough spatialisation permits the listener to approximately locate the sound source within the *Eden* world. The bandwidth devoted to sound is much higher than any other sensory information used by the agent. Agents are able to differentiate and make sound over a range of frequency bands, giving rich opportunities for the use of sound in an ecosystem context.

In addition to the internal ecosystem model, the *Eden* world is also connected to the physical world of the installation space via an infrared video camera which tracks the presence and motion of people looking and listening to the artwork². The presence of people in the installation space influences the growth of biomass in the virtual space. The longer people spend with the work, the more food is likely to grow in the virtual environment. The rationale for this is driven by the idea that the more interesting people find the work, the longer they will stay. If they find the work uninteresting, they will not spend much time with it. A good way to maintain people’s interest is to produce sounds, moreover, *interesting, changing* sounds.

Over time, the agents evolve to make complex sounds in order to maintain their food supply. The agents have no specific knowledge of people in the environment, however, by making interesting combinations of sounds they attract and maintain the interest of the human audience in the environment³. This interest translates to a more stable supply of food, hence improving chances of survival in the environment. Therefore, *Eden* is a symbiotic ecosystem, which includes the human audience experiencing the work.

²The original version of the work used infrared distance sensors.

³When shown in a gallery environment, it is important to remember to compensate for opening hours, otherwise the population dies out each night when the gallery is closed!

5 Conclusions

In contrast with previous attempts to model creativity, which have applied psychological, cognitive, or knowledge-based models of human creativity, the ecosystem approach sees creativity as an emergent phenomenon of dynamic interaction between interconnected, self-organising components and their environment. These components and their environment may be internal to computer simulation (as in the *Colourfield* system) or part of a system that incorporates humans and the physical environment (as with the *Eden* system).

Combinatorial systems do not practically impose the limitations that might be suggested by the opposing concepts of combinatoric and creative emergence. Necessarily, all base primitives must contain an interpretation that lies outside the software itself. What is important is the process used to derive a creative result from a set of base primitives. The goal is to enable the synergistic exploration of new conceptual spaces in creative partnership with the machine. In the artificial ecosystem approach, this can be achieved by developing a formal understanding of the appropriate design of components, their interconnections, and the environment in which they operate.

References

- Adami, C. (2002). Ab initio modeling of ecosystems with artificial life. *Natural Resource Modeling*, 15:133–146.
- Arthur, W. B., Durlauf, S., and Lane, D. A., editors (1997). *The economy as an evolving complex system II*. Addison-Wesley, Reading, MA.
- Ashby, W. R. (1952). *Design for a Brain*. Chapman & Hall, London.
- Baas, N. A. (1994). Emergence, Hierarchies and Hyperstructures. In Langton, C. G., editor: *Artificial Life III*, 515–537. Addison-Wesley, Reading, MA.
- Bird, J. (2004). Containing Reality: Epistemological Issues in Generative Art and Science. In *Impossible Nature: the art of Jon McCormack*, 40–53. Australian Centre for the Moving Image, Melbourne.
- Birkhoff, G. D. (1933). *Aesthetic Measure*. Harvard University Press, Cambridge, MA.
- Cariani, P. (1991). Emergence and Artificial Life. In Langton, C. G. et. al., editors: *Artificial Life II, SFI Studies in the Sciences of Complexity*, 775–797. Addison-Wesley, Redwood City, CA.
- Cariani, P. (1997). Emergence of new signal-primitives in neural systems. *Intellectica*, 2:95–143.
- Conrad, M. and Pattee, H. H. (1970). Evolution experiments with an artificial ecosystem. *Journal of Theoretical Biology*, 28:393.
- Dartnall, T., editor (2002). *Creativity, Cognition, and Knowledge: An Interaction*. Praeger, Westport, Connecticut.
- Di Scipio, A. (2003). ‘Sound is the interface’: from interactive to ecosystemic signal processing. *Organised Sound*, 8(3):269–277.
- Dorin, A. (2001). Aesthetic fitness and artificial evolution for the selection of imagery from the mythical infinite library. In Kelemen, J. and Sosik, P., editors, *Advances in Artificial Life*, LNAI 2159, 659–668. Springer-Verlag, Berlin.
- Dorin, A. and McCormack, J. (2002). Self-Assembling Dynamical Hierarchies. In Standish, R. K., et. al. editors: *Artificial Life VIII: Proceedings of the Eight International Conference on Artificial Life*, 423–428. MIT Press, Cambridge, MA.
- Eiben, A. E. and Smith, J. E. (2003). *Introduction to Evolutionary Computing*. Natural Computing Series. Springer, Berlin.
- Epstein, J. M. and Axtell, R. (1996). *Growing Artificial Societies*. MIT Press, Cambridge, MA.
- Grimm, V. and Railsback, S. F. (2005). *Individual-based Modeling and Ecology*. Princeton Series in Theoretical and Computational Biology. Princeton University Press.
- Lenton, T. M. and Lovelock, J. E. (2001). Daisyworld revisited: quantifying biological effects on planetary self-regulation. *Tellus*, 53B(3):288–305.
- McCormack, J. (2001). Eden: An evolutionary sonic ecosystem. In Kelemen, J. and Sosik, P., editors, *Advances in Artificial Life*, LNAI 2159, 133–142. Springer-Verlag, Berlin.
- McCormack, J. (2005a). On the Evolution of Sonic Ecosystems. In Adamatzky, A. and Komosinski, M., editors: *Artificial Life Models in Software* 211–230. Springer-Verlag, London.
- McCormack, J. (2005b). Open problems in evolutionary music and art. In Rothlauf, F. et. al. editors, *EvoWorkshops*, LNCS 3449, 428–436. Springer, Berlin.
- Michalewicz, Z. and Fogel, D. B. (1999). *How to solve it: modern heuristics*. Springer, New York.
- Mitchell, M. and Taylor, C. E. (1999). Evolutionary computation: An overview. *Annual Review of Ecology and Systematics*, 30:593–616.
- Nowak, M. A. (2006). *Evolutionary Dynamics: exploring the equations of life*. The Bekknap Press of Harvard University Press, Cambridge, MA, and London, England.
- Swenson, W., Wilson, D. S., and Elias, R. (2000). Artificial ecosystem selection. *PNAS*, 97(16):9110–9114.
- Takagi, H. (2001). Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE*, 89:1275–1296.
- van Langen, P. H. G., Wijngaards, N. J. E., and Brazier, F. M. T. (2004). Towards designing creative artificial systems. *AIEDAM, Special Issue on Learning and Creativity in Design*, 18(4):217–225. A. H. B. Duffy and F. M. T. Brazier (editors).
- Wilson, S. W. (1999). State of XCS classifier system research. Technical report, Concord, MA.