

Artificial Ecosystems for Creative Discovery

Jon McCormack
Centre for Electronic Media Art
Faculty of Information Technology
Monash University Clayton 3800, Australia
Jon.McCormack@infotech.monash.edu.au

ABSTRACT

This paper discusses the concept of an *artificial ecosystem* for use in machine-assisted creative discovery. Properties and processes from natural ecosystems are abstracted and applied to the design of creative systems, in a similar way that evolutionary computing methods use the metaphor of Darwinian evolution to solve problems in search and optimisation. The paper examines some appropriate mechanisms and metaphors when applying artificial ecosystems to problems in creative design. General properties and processes of evolutionary artificial ecosystems are presented as a basis for developing individual systems that automate the discovery of novelty without explicit teleological goals. The adaptation of species to fit their environment drives the creative solutions, so the role of the designer shifts to the design of environments. This allows a variety of creative solutions to emerge in simulation without the need for explicit or human-evaluated fitness measures, such as those used in interactive evolution. Two example creative ecosystems are described to highlight the effectiveness of the method presented.

Categories and Subject Descriptors

J.5 [Arts and Humanities]: Fine Arts—*creativity*; I.2.11 [Distributed Artificial Intelligence]: Multiagent systems—*artificial ecosystems, evolutionary algorithms*

General Terms

Algorithms, Design, Experimentation

Keywords

Ecosystem theory; creativity

1. INTRODUCTION

Evolutionary Computing (EC) has successfully exploited the metaphor of Darwinian evolution to solve problems in search, optimisation, and machine learning. Different EC

techniques selectively adopt particular processes and metaphors from biological evolution. For example, many complex multi-modal problems are intractable to algorithms such as the Simple Genetic Algorithm (SGA), hence the use of more sophisticated techniques such as competitive and co-operative evolution, spatially distributed, diffusion model and island model Evolutionary Algorithms (EAs), see e.g. [7].

Search and optimisation problems suited to EC approaches are generally characterised by the requirement of explicit ranking. That is, a machine representable fitness function is required to rank and select potential solutions, which seek to minimise or maximise the fitness function. However, there exist classes of problems where explicit specification of an appropriate fitness function is impractical to compute or impossible (practically or theoretically) to express as a machine-representable function. Selection methods such as tournament selection may help, but such methods still require an explicit *evaluation* of an individual's relative fitness [3].

For creative problems, methods such as *Interactive Evolution* have found wide application and popularity [24]. 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 for this problem, it is not without significant problems [6, 18]. 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 [7, p. 29]).

Traditionally in EC, the emphasis has been placed on selection of individuals by fitness ranking, and the modification of genotypes using mutation and crossover. As EC operates using a *metaphor* based on evolution¹, it is not surprising that it focuses on some features of Darwinian selection, while ignoring or minimising others, such as environment. Genes exist in cells, cells may form multi-cellular groupings, organs, and so on. Organisms develop and interact in a physical environment. Selection pressures come from other individuals, species, and from the environment itself. The evolutionary process typically results in *adaptation* of populations to fit their environment, so fit organisms will devise features suited to the specifics of their evolutionary milieu. In addition, populations of species may (explicitly or implicitly) *modify* their environment, creating feedback

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'07, July 7–11, 2007, London, England, United Kingdom.
Copyright 2007 ACM 978-1-59593-697-4/07/0007 ...\$5.00.

¹Actually, the metaphor is closer to selective breeding than evolution by natural selection.

loops between adapting populations and their environment.

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. This paper introduces the concept of an *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 [23]. 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.

2. ARTIFICIAL ECOSYSTEMS

In biology, an ecosystem is defined as “the interactions of species with each other and their physical environment” [23]. Ecosystem models typically place emphasis on abiotic factors and heterogeneous environments. Natural ecosystems have been well studied in the biological literature from mathematical modelling perspectives [15, 11, 14]. Models are, naturally, diverse and approached from particular research directions (e.g. chemical and energy cycles, predator-prey systems, population genetics, trophic levels, co-evolution, migration, disease propagation, territorial behaviour). Much “traditional” ecological research models a species or population as an individual, normalised unit, rather than as an emergent collection of individual agents. Hence, models are expressed in terms of differential equations. An alternative approach models the interaction and behaviour at an individual level (referred to in the Ecology community as an “individual-based model” or IBM [9]). The approach described here is more closely based on this IBM approach to modelling ecosystems (see Section 2.1.1).

Biologists also experiment with *artificial ecosystems*, i.e. artificially constructed ecosystems made with biological materials [23]. Examples include soil or water ecosystems and bacterial cultures. However, for the remainder of this paper I will use the term “artificial ecosystem” to refer to a type of computer simulation based on dynamic processes that occur in biological ecosystems (either natural or artificial).

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 seen as a dynamic, complex system, one that is 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 and apply them to creative processes of interest to humans.

2.1 Related Work

The idea of abstracting a process from biology and ap-

plying it to other domains, or exploiting it as a more general problem-solving technique is not new, EC being a fine example. The central question in relation to the suitability of translating biological novelty to creative novelty lies in both *metaphors* and *mechanisms*. Metaphors—in terms of how we replace the biological literal with the creative symbolic—determine the scope and framework of the systems to be explored. The mechanisms form the machine representable abstractions of natural ecosystem processes. Real ecosystems are highly complex and diverse systems, so we need to be careful about the processes we wish to mimic and how well suited (or otherwise) they are to the goal of creative discovery.

2.1.1 Artificial Life and Ecological Models

A number of artificial life models employ the concept of an abstract or simplified ecosystem. This concept of the artificial ecosystem was introduced in [4]. 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 [2], ecology [19] and social science [8].

The majority of such systems focus on single-niche, homogeneous environments, and operate at evolutionary time-scales, simulating the “evolution” of a single-species population. This focus, and the use of minimal, broad assumptions are primarily for the purposes of verification and validation of evolution in artificial life models [1]. 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.

2.1.2 Creative Ecosystems

The concept of an ecosystem has been employed in creative works, in addition to the work described in this paper. Most of these works are characterised by some form of user interaction or influence over the environment or individuals. *Avolve*, an interactive artwork by Sommerer and Mignonneau allows users to manually design agents, then release them into an artificial ecosystem in which they interact with other agents created by users [22]. The ecosystem is projected into a pool of real water, which also forms its user interface. Moving one’s hand through the water influences the swimming movement, hunting behaviour, and reproduction of the virtual organisms. Similar in its use of artificial

ecosystems with user-specified agents, the artwork *Technosphere* by Prophet and Selley uses agents built by the user from a fixed catalogue of components, such as wheels, heads and eyes [21]. Agents are then released into a shared virtual world to compete for resources and mate with others. The *E-volver* work of Dutch artists Erwin Driessens and Maria Verstappen uses a population of virtual creatures to construct a continuously changing image [12]. Each pixel-sized creature modifies its local environment, hence creating a dynamic interaction between populations of agents and their environment. The system uses interactive evolution to select and breed creatures that create the images. The pixels of the image and the agents that interact with them form the ecosystem, which is the artwork itself.

Each of these systems are highly individual and specific to the creative concerns of the individuals who created them. The goal of this paper is to begin to formalise some of the metaphors and mechanisms appropriate when using artificial ecosystems for creative discovery. In doing this it is hoped that artificial ecosystem algorithms may be more readily applied as general problem solvers for creative applications, rather than highly specific, ad-hoc mechanisms used by individual artists.

2.2 Processes for Artificial Ecosystems

In many artificial ecosystem models, the designers 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 [9]. 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 (this is further explained in Section 2.3);
- physical 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 a scalar measure of an individual’s success in surviving within its environment;
- 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 gives some species a higher reproduction rate over others, creating an implicit measure of fitness [20, 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 and novel changes.

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, but different strategies may lead to similar health levels. 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 form equally good solutions (e.g. notes or phrases forming a musical composition). When using standard EA methods for search or optimisation, the challenge faced is in choosing appropriate 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.

2.3 Species and Speciation

Biological ecosystems consist of multiple species interacting with each other and their environment. This may lead to inter-species relationships such as symbiosis or parasitism.

“Extended phenotype” effects, such as co-operative building from resources in the environment, may also be observed [5]. In artificial ecosystem models, species may be one of (i) explicitly specified, or (ii) emergent within the simulation. A species is defined as a individual type that may reproduce with other individuals of a matching type, but not with any others. The measure of matching type may be binary or over some equivalence interval.

With explicit specification, species are fixed from the onset of the simulation. Their number and type do not change during simulation. For example, in predator-prey models, predator and prey species, each with their fixed roles, are established *a priori*. One species may not change into another, nor can new species emerge in simulation.

Implicit specification gives individuals some form of “type” information (similar to what Holland refers to as *tags* [10]). This type information is normally encoded in the genome of the individual. Individuals may reproduce with other individuals with *type equivalence* (matching tags). The measure of equivalence may be over some interval or a stochastic measure (e.g. $\geq 95\%$ equivalent). This permits new species to dynamically emerge within the simulation.

A “species” might also be defined implicitly without the need for explicit type information by the percentage of alleles shared with other individuals (e.g. individuals with $\geq n\%$ of their alleles identical or within some normally-distributed limits are permitted to mate). Such a method may form the basis of speciation mechanism that restricts intra-species reproduction and permits the evolution of new species during simulation. Combining this mechanism with the concept of a spatial distribution of phenotypes over a heterogeneous landscape encourages allopatric speciation, for example.

3. EXAMPLES

I now highlight the way the artificial ecosystem model, described in this paper, has been successfully used as an alternative to other EC methods for creative discovery.

3.1 Colourfield

Colourfield is a simple one-species ecosystem of colour patterns and sounds. It consists of a one-dimensional discrete world of fixed width, n , populated by $0 \leq m \leq n$ individuals (Figure 1). 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 and sonically as a sine tone: frequency maps logarithmically to hue; phase to saturation; and amplitude to area. A large, red colour produces a loud, deep tone; a small blue colour produces a quiet, high tone. A population of individuals thus produces a field of colour and sound simultaneously.

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

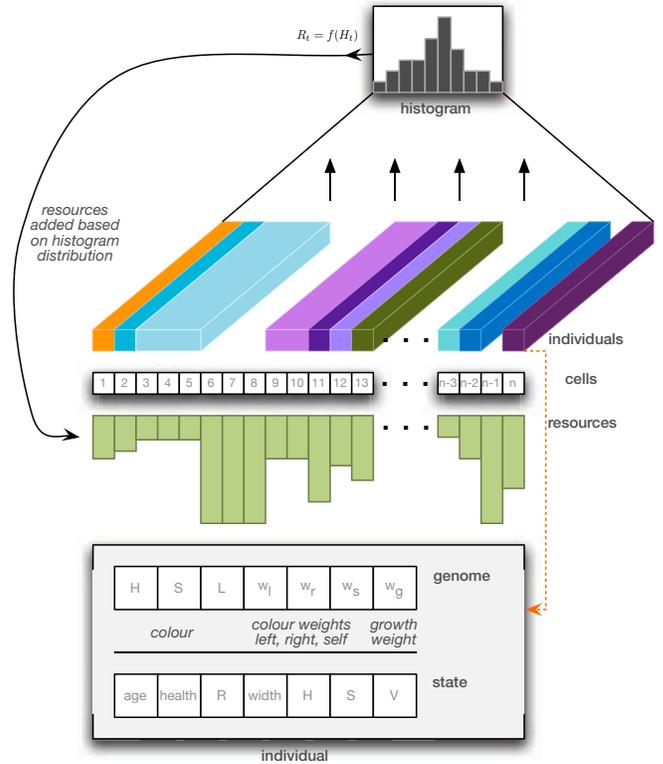


Figure 1: Schematic overview of *Colourfield*

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 an 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 this 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, t , 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 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.

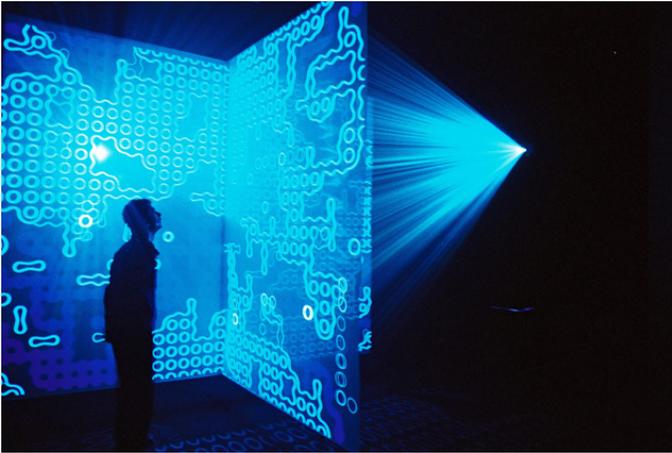


Figure 2: Installation view of *Eden*

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 some 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 and sound). Let us now move to a more complex ecosystem example.

3.2 *Eden*: an Evolutionary Sonic Ecosystem

Eden is an installation artwork 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 [16, 17].

The work consists of a complex artificial ecosystem run-

ning in real-time on a two-dimensional lattice of cells. This world is projected into a three-dimensional environment, approximately 6m x 6m (see Figure 2). The ecosystem 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 [13], with the growth rate β_i for an individual biomass element, i , a Gaussian function of local temperature at the location (x, y) of the element, $T_{x,y}$:

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

Local temperature is determined by the amount of biomass present at (x, y) . 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 give rise to seasonal variations of temperature, 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 [25]). 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 single-distance 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 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 that the more interesting the work, the longer the audience will stay. If people find the work uninteresting (either visually or sonically), 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.

3.3 Feedback Loops

The success of the systems described is in part due to the feedback loops of resources within the environments. Figure 3 shows the causal relations between entities and resources within each of the artificial ecosystem models discussed in this section. As can be seen, *Eden* has a more complex network of relations which mirrors, in part, the more complex range of creative behaviours it is capable of producing.

3.4 Discussion

Table 1 summarises the properties and processes (as described in Section 2.2) for the example creative ecosystems. Adopting metaphors and mechanisms appropriate to the problem domain is a major consideration in developing a successful ecosystem. Note that both systems operate using both individual (I) and evolutionary (E) time scales, that is, each individual has a simulated life-cycle and the simulation is run long enough for populations to discover adaptations to their environment via evolution. In the case of *Eden*,

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!

Table 1: Summary of Ecosystem Characteristics

	<i>Colourfield</i>	<i>Eden</i>
Genotype	scalar vector	XCS rules
Phenotype	colours/sounds	sonic agents
Species model	single	multiple
Spatial model	1D discrete	2D discrete
Environment	colours	agents, biomass, etc.
Health	from environment	from food
Energy source	colour histogram	sunlight, users
Metabolism	colour histogram	food processing
Timescales	I, E	I, E
Evolution	Darwinian	Lamarckian
Discovery	Colour composition	Musical composition

an agent’s classifiers evolve during the lifetime of the agent, and successful rules are passed to offspring, so adaptations may occur faster than in a Darwinian model. This essentially circumvents the need for agents to re-learn how to live in their world from birth, akin to parents teaching their children what they have learned in their lifetime.

Ecosystem algorithms circumvent the problems of interactive evolution by replacing explicit fitness evaluation with environmental adaptations. The problem thus becomes one of designing environments from which creative solutions are likely to emerge. As illustrated in the examples, the specific nature of the discovery domain determines the composition of the environment and the interactions between its elements.

4. CONCLUSIONS

The concept of an artificial ecosystem algorithm for creative discovery is still formative. However, as demonstrated by the examples in this paper, it holds potential as a method to permit creative discovery where other methods, such as interactive evolution or traditional EC methods are inappropriate. Artificial ecosystems exploit the novel discovery of real ecosystem processes. The big question remains as to *how* to exploit these properties. This paper proposed a focus on metaphor and mechanism as a suitable approach. The examples show how suitable metaphors (e.g. individual = colour, species = colour spectrum in *Colourfield*) and mechanisms (e.g. symbiosis between audience and artificial agents in *Eden*) can be applied to create successful creative works.

The potential of artificial ecosystems is only beginning to be discovered. It is hoped that future work may further formalise the concepts and process and lead to more general modes of creative discovery.

5. REFERENCES

- [1] C. Adami. Ab initio modeling of ecosystems with artificial life. *Natural Resource Modeling*, 15:133–146, 2002.
- [2] W. B. Arthur, S. Durlauf, and D. A. Lane, editors. *The economy as an evolving complex system II*. Addison-Wesley, Reading, MA, 1997.
- [3] T. Bickle and L. Thiele. A comparison of selection schemes used in genetic algorithms. Technical Report 11, Swiss Federal Institute of Technology, December 1995.

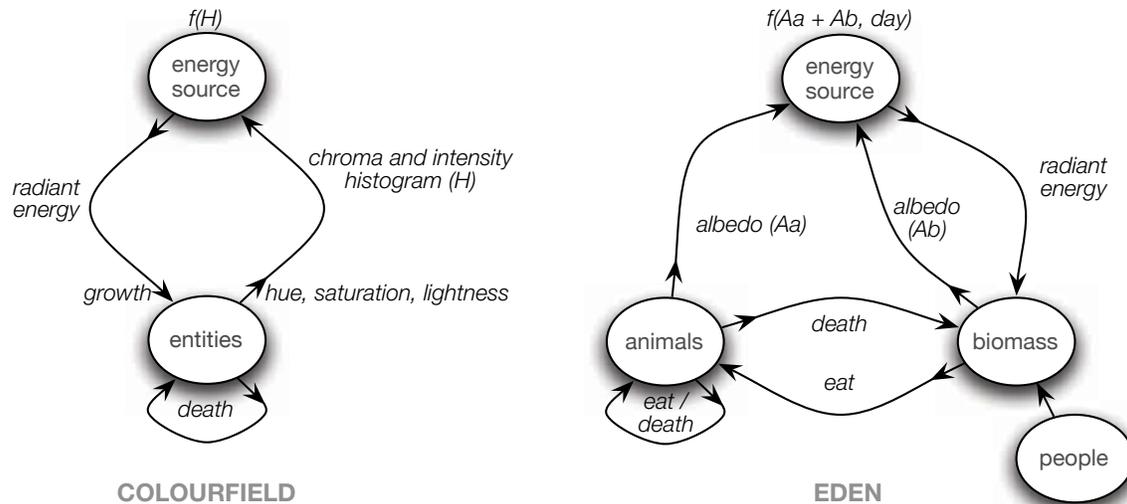


Figure 3: Causal resource flows for *Colourfield* and *Eden*

- [4] M. Conrad and H. H. Pattee. Evolution experiments with an artificial ecosystem. *Journal of Theoretical Biology*, 28:393, 1970.
- [5] R. Dawkins. *The extended phenotype: the gene as the unit of selection*. Freeman, Oxford; San Francisco, 1982.
- [6] A. Dorin. Aesthetic fitness and artificial evolution for the selection of imagery from the mythical infinite library. In J. Kelemen and P. Sosik (eds), *Advances in Artificial Life, Proceedings of the Sixth European Conference, ECAL*, LNCS 2159:659–668, 2001. Springer-Verlag.
- [7] A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Natural Computing Series. Springer, Berlin, 2003.
- [8] J. M. Epstein and R. Axtell. *Growing Artificial Societies*. MIT Press, Cambridge, MA, 1996.
- [9] V. Grimm and S. F. Railsback. *Individual-based Modeling and Ecology*. Princeton Series in Theoretical and Computational Biology. Princeton University Press, 2005.
- [10] J. H. Holland. *Hidden order: how adaptation builds complexity*. Helix books. Addison-Wesley, Reading, MA, 1995.
- [11] S. E. Jørgensen. *Integration of Ecosystem Theories: A Pattern*. Kluwer Academic Publishers, Dordrecht, second revised edition, 1997.
- [12] K. Kuitenbrouwer and W. Lentz. *E-volver*. SKOR (Stichting Kunst en Openbare Ruimte/Foundation Art and Public Space), Amsterdam (The Netherlands), 2006.
- [13] T. M. Lenton and J. E. Lovelock. Daisyworld revisited: quantifying biological effects on planetary self-regulation. *Tellus*, 53B(3):288–305, 2001.
- [14] R. M. May. *Stability and Complexity in Model Ecosystems*. Princeton University Press, Princeton, NJ, second edition, 2001.
- [15] J. Maynard Smith. *Models in Ecology*. Cambridge University Press, London, 1974.
- [16] J. McCormack. Eden: An evolutionary sonic ecosystem. In J. Kelemen and P. Sosik (eds), *Advances in Artificial Life, Proceedings of the Sixth European Conference, ECAL*, LNCS 2159:133–142, 2001.
- [17] J. McCormack. On the Evolution of Sonic Ecosystems. In A. Adamatzky and M. Komosinski (eds), *Artificial Life Models in Software*, pages 211–230. Springer-Verlag, London, 2005.
- [18] J. McCormack. Open problems in evolutionary music and art. In F. Rothlauf, et. al. (eds), *EvoWorkshops*, LNCS 3449, pages 428–436. Springer, 2005.
- [19] M. Mitchell and C. E. Taylor. Evolutionary computation: An overview. *Annual Review of Ecology and Systematics*, 30:593–616, 1999.
- [20] M. A. Nowak. *Evolutionary Dynamics: exploring the equations of life*. The Bekknap Press of Harvard University Press, Cambridge, Massachusetts, and London, England, 2006.
- [21] J. Prophet and G. Selley. Technosphere: “real” time “artificial” life. *Leonardo*, 34:309–312, 2001.
- [22] C. Sommerer and L. Mignonneau. Art as a Living System. In C. Sommerer, and L. Mignonneau (eds), *Art@Science*, pages 148–161. Springer, Wein, 1998.
- [23] W. Swenson, D. S. Wilson, and R. Elias. Artificial ecosystem selection. *PNAS*, 97(16):9110–9114, August 1 2000.
- [24] H. Takagi. Interactive evolutionary computation: Fusion of the capabilities of EC optimization and human evaluation. *Proceedings of the IEEE*, 89:1275–1296, Sep 2001.
- [25] S. W. Wilson. State of XCS classifier system research. Technical report, Concord, MA, March 1999.