

Chapter 13

The Evolution of Sonic Ecosystems

Jon McCormack

This chapter describes a novel type of artistic artificial life software environment. Agents that have the ability to make and listen to sound populate a synthetic world. An evolvable, rule-based classifier system drives agent behavior. Agents compete for limited resources in a virtual environment that is influenced by the presence and movement of people observing the system. Electronic sensors create a link between the real and virtual spaces, virtual agents evolve implicitly to try to maintain the interest of the human audience, whose presence provides them with life-sustaining food.

13.1 Introduction

One thing that foreigners, computers and poets have in common is that they make unexpected linguistic associations.

Jasia Reichardt [26]

Music and art are undoubtedly fundamental qualities that help define the human condition. While many different discourses contribute to our understanding of art making and art interpretation, two implicit themes connect all artworks. The first is the act of creation. Even the most abstract or conceptual artworks cannot escape the fact that, as ideas, objects, or configurations, they must be made. Second, the importance of novelty, either perceived or real, is a fundamental driving force behind any creative impetus or gesture. Artists do not seek to create works that are identical to their previous creations or the previous work of others.

Artificial life (AL) methodologies can play an important role in developing new modes of artistic enquiry and musical composition. For artists, AL can offer new methodologies for the creative arts. For the first time in the history of art, AL suggests that, in theory at least, it may be possible to create artificial organisms that develop their own autonomous creative practices –

to paraphrase the terminology of Langton [16], *life-as-it-could-be* creating *art-as-it-could-be*.

In addition, AL has important contributions to make in our understanding of genuine novelty,¹ often referred to under the generalized term *emergence* [6, 11, 21].

13.1.1 Artificial Life Art

Techniques from cybernetics and artificial life have found numerous applications in the creative arts. General contemporary overviews can be found in [2, 3, 32, 39], for example.

Cybernetics has a rich and often overlooked history in terms of computing and the arts. The seminal ICA exhibition *Cybernetic Serendipity*, held in London in the summer of 1968, was one of the first major exhibitions to look at connections between creativity and technology [25]. Even the title suggests notions of novelty and discovery, a key theme for many works and critics in the decades that have followed the exhibition. Interestingly, the curators shunned distinctions between art and science and instead focused on ideas and methodologies that connected the two.

One particularly relevant concept from cybernetics is that of *open-ended behavior*, what Ashby referred to as *Descartes dictum*: How can a designer build a device that outperforms the designer’s specifications [1]? Cyberneticist Gordon Pask built an “ear” that developed, not through direct design, but by establishing certain electrochemical processes whereby the ear formed and developed in response to external stimuli [7].

The goal of the work described here is to create an open-ended artistic system that is *reactive* to its environment. In order to address this goal, two important problems were explored during the design and development of the work: first, how we can create a virtual AL world that evolves toward some subjective criteria of the audience experiencing it, without the audience needing to explicitly perform fitness selection and, second, how the relationship between real and virtual spaces can be realised in a way that integrates those spaces phenomenologically. The resultant artwork developed by the author is titled *Eden*.

¹ The concept of novelty is a vexed one with many different interpretations in the literature and could easily occupy an entire chapter in itself. Some authors argue that novelty and emergence have no relation [23], whereas others see them as fundamentally the same [6]. In the sense that the term is used in this chapter, novelty suggests that which has never existed before, hence the issues surrounding novelty are connected with determinism [11]. For art, almost every new artwork is in some sense novel; however, we may at least be able to apply criteria that suggest a degree of novelty, such as descriptive causality and explainable causality. Moreover, in an AL sense, we require not only the artwork to be novel, but the behavior of the virtual agents to be novel as well.

In terms of software, *Eden* is an AL environment that operates over a cellular lattice, inhabited by agents who have, among other capabilities, the ability to make and “listen” to sound. Agents use an internal, evolvable, rule-based system to control their behavior in the world. The virtual environment that the agents inhabit develops in response to the presence and movement of people experiencing the system as an artwork. The work is conceptualised and designed as an *artificial ecosystem* where virtual and real species interact with their biotic and a-biotic environment.

This software system will be described more fully in Sections 13.2 and 13.3 of the chapter. Interaction with the work is detailed in Section 13.4, with a summary of results and brief conclusion in Sections 13.5 and 13.6.

13.1.2 Related Work

The software system described in this chapter draws its technical inspiration from John Holland’s *Echo* [15], particularly in the use of classifier systems for the internal decision-making system of agents. Many others have used evolutionary systems as a basis for musical composition, but in the main for compositional *simulation* [33, 38], rather than as a new form of creative tool for the artist and audience.

The *Living Melodies* system [9] uses a genetic programming framework to evolve an ecosystem of musical creatures that communicate using sound. *Living Melodies* assumes that all agents have an innate “listening pleasure” that encourages them to make noise to increase their survival chances. The system described in this chapter, *Eden*, contains no such inducement, beyond the fact that some sonic communication strategies that creatures discover should offer a survival or mating advantage. This results in the observation that only some instances of evolution in *Eden* result in the use of sonic communication, whereas in *Living Melodies*, every instance evolves sonic communication. *Living Melodies* restricts its focus to music composition, whereas *Eden* is both a sonic and visual experience.

13.2 *Eden*: An Artificial Life Artwork

Eden is a “reactive” AL artwork developed by the author. The artwork is typically experienced in an art gallery setting, but in contrast to more traditional artworks, it is designed as an *experiential* environment, whereby viewers participation and activity within the physical space have important consequences over the development of the virtual environment.

The artwork is exhibited as an installation and can be experienced by any number of people simultaneously. It consists of multiple screens, video projec-




Fig. 13.1 Plan and side views of *Eden* showing the layout screens, speakers, projectors, and camera sensor area. The active sensing area extends approximately one metre past the edge of the screens.

tors, audio speakers, an infrared camera and lighting system, computers, and custom electronics. Figure 13.1 shows layout plans and Figure 13.2 shows a simulated visualisation of the work. As shown in these figures, physically the work consists of two semitransparent screens suspended from the ceiling of the exhibition space. The screens are positioned at 90° to each other, forming an X shape when viewed in plan. The ambient lighting is minimal – making the screens and the light they reflect and transmit the predominant source of visual interest in the space. The screens' transparency enables them to be viewed from either side and creates a layered visual effect that merges the real and virtual boundaries. Multichannel audio is provided by a number of speakers placed on the periphery of the main screen area.

In addition to this audio-visual infrastructure, an infrared digital video camera is placed above the screens, looking down at the space surrounding the projection screens. This area is illuminated by an infrared lighting system, which is invisible to the naked eye and so does not affect the perceptual properties of the work. The purpose of this camera sensor system is to measure the position and movement of people experiencing the work in its immediate vicinity. It is not necessary that the audience has any direct knowledge of this sensing. The purpose of sensing the real space of the artwork is as an environmental stimulus for the virtual agents' world and ultimately to con-




Fig. 13.2 Simulation of *Eden* running in a typical gallery environment, illustrating the effect of using transparent screens to visualise the work.

tribute to selective pressures that aim to encourage a symbiotic relationship between people experiencing the work and the agents populating the virtual world. The role of sensing and its effect on the development of the virtual environment portrayed in the work are detailed in Section 13.4.

13.3 Agents and Environments

This section gives technical details on the major software components of the system, with particular emphasis on the mechanisms that facilitate development of sonic agents within the system. Further details, particularly the *payoff* and *bidding* processes for rule selection, may be found in [18].

13.3.1 The Eden World

The environment projected onto the screens is known as the *Eden world*. In implementation terms, the world consists of a two-dimensional toroidal cellular lattice that develops using a global, discrete, time-step model – a popular AL model based on the theory of cellular automata [8, 35]. Each cell in the lattice may be populated by one of the following entities:

- *Rock*: inert matter that is impervious to other entities and opaque to sound and light. Rock is placed in cells at initialisation time using a variation of the *diffusion limited aggregation* (DLA) model [40]. Rocks provide refuge and contribute to more interesting spatial environmental behavior of the agents.
- *Biomass*: a food source for evolving entities in the world. Biomass grows in yearly² cycles based on a simple feedback model, similar to that of *Daisyworld* [36]. Radiant energy (in “infinite” supply) drives the growth of biomass. The amount of radiant energy falling on a particular cell is dependent on a number of factors, including the local absorption rate of the biomass and global seasonal variation. Probabilistic parameters can be specified at initialisation time to control these rates and variations. The efficiency at which the biomass converts radiant energy into more biomass is also dependent on the presence of people in the real space of the artwork. This dependency is detailed in Section 13.4.
- *Sonic Agents*: mobile agents with an internal, evolvable *performance system*. Agents get energy by eating biomass or by killing and eating other agents. More than one agent may occupy a single cell. Since these agents are the most complex and interesting entity in the world, they are described in detail in Section 13.3.2.

² An *Eden* year lasts 600 *Eden* days, but passes by in about 10 minutes of real time.

A real-time visualisation of the world produces images that are projected onto the screens, as illustrated in Fig. 13.1 (in this case, there are two worlds, each running on a separate computer, but connected as a single logical world running over two computers). The visualisation process is described more fully in Section 13.3.3. The sound the agents make as they move about the world is played with approximate spatial correspondence by a series of loud-speakers.

13.3.2 Agent Implementation

Sonic agents are the main evolving entity in the world. Essentially, the agent system uses a learning classifier system (LCS) similar to that of Holland's *Echo* system [15]. An agent consists of a set of *sensors*, a rule-based *performance system*, and a set of *actuators*. This configuration is illustrated in Figure 13.3. Sensors provide measurement of the environment and internal introspection of an individual agent's status. The performance system relates input messages from the sensors to desired actions. The actuators are used to show intent to carry out actions in the world. The success or failure of an intended action will be dependent on the physical constraints in operation at the time and place the intent is instigated. Actuators and actions are detailed later in this section.




Fig. 13.3 A section of the *Eden* cellular lattice in visual form (left). To emphasize the lattice structure, grid lines have been layered over the image. The image shows rocks (solid), biomass (outline), and an agent (thick circle). The diagram (right) shows the agent's internal schematic structure, consisting of a number of sensors, a performance system that evolves, and a set of actuators.

At initialisation of the world, a number of agents are seeded into the population. Each agent maintains a collection of internal data, which includes the following:

- *Current age*, an integer measured in time steps since birth. Agents live up to 100 years and cannot mate in their first year of life.
- *Health index*: an integer value indicating the overall health of the agent. A value of 100 indicates perfect health; if the health index falls to 0, the agent dies. An agent can lose health via a sustained negative *energy level* differential (explained next) by bumping into solid objects, such as rocks, or being hit by other agents. In addition, the loss in health from being hit by another agent depends on both its mass and health index.
- *Energy level*: a measure of the amount of energy the agent currently has. Agents gain energy by eating biomass or other agents. Energy is expended attempting to perform actions (regardless of their success); a small quantity of energy is expended even if no action is performed at a given time step. If an agent's energy level falls to zero, the agent dies and its body is converted to new biomass in the cell in which it died.
- *Mass*: an agent's mass is linearly proportional to its energy level, plus an initial "birth mass" that is normally distributed over the population.

13.3.2.1 Sensors

Sensors provide a way for an agent to measure itself and its environment [24]. Sensor data are presented as bit strings constructed from local environmental conditions and from the internal data structures held by the agent. Sensor data are updated once every time step. An agent can use a range of sensor types, but the sensors themselves do not undergo any evolution and are fixed in function, sensitivity, and morphology. It is up to an individual agent's performance system to make use of a particular sensor, so data from a particular sensor will only be used in the long term if it provides useful information that assists the agent's survival or mating prospects. Sensor use does not incur any cost to the agent. Sensor information available to an agent consists of:

- A simple local vision system that detects the "colour" of objects on facing and neighboring cells (the range is limited to a single cell). Rocks, biomass, and agents all have different "colours," which enables an agent to distinguish between them.
- A sensor to detect the local cell nutritional value. Cells that contain biomass or dead agents have a high nutritional value; rocks and empty cells do not.
- A sound sensor that detects sound pressure levels over a range of frequency bands. Sound can be detected over a much larger spatial range than vision and also with greater fidelity.

- An introspection of *pain*. Pain corresponds to a negative health index differential and would usually indicate attack by another agent or that the agent is bumping into rocks.
- An introspection of the current energy level.

13.3.2.2 Actuators

Actuators are used to signal an agent's intent to carry out an action in the world. The physical laws of the world will determine whether the intended action can be carried out or not. For example, the agent may intend to "walk forward one cell," but if that cell contains a rock, the action will not be possible. Furthermore, all actions cost energy, the amount dependent on the type of action and its context (e.g., attempting to walk into a rock will cost more energy than walking into an empty cell).

As with the sensors, the number and function of actuators are fixed and do not change as the performance system evolves. Actions will only be used in the long term if they benefit the agent. Analysis of actions used by agents who are successful in surviving shows that not all agents make use of the full set of actuators.

Agents may perform any of the following actions:

- *Move* forward in the current direction.
- *Turn* left or right.
- *Hit* whatever else is in the cell occupied by the agent. Hitting another agent reduces that agent's health level using a non-linear combination of the mass, health, and energy level of the agent performing the hit. Hitting other objects or being hit will cause pain and a loss of health.
- *Mate* with whatever is currently occupying the current cell. Obviously, this is only useful if another agent is in the same cell. In addition, mating is only possible if the age of both agents is greater than 1 year.
- *Eat* whatever is currently occupying the current cell. Agent's can only eat biomass or dead agents (which turn into biomass shortly after death).
- *Sing*: make a sound that can be heard by other agents. Sound is detailed more fully in Section 13.3.4.

Performing an action costs energy, so agents quickly learn not to perform certain actions without benefit. For example, attempting to eat when your nutritional sensor is not activated has a cost but no benefit. Attempting to move into a rock has a cost greater than moving into an empty cell.

Agents may also choose not to perform any action at a given time step (a "do nothing" action), but even this costs energy (although less than any other action).

13.3.2.3 Performance System

The performance system connects an agent's sensors to its actuators (Fig. 13.3). It is based on the classification system of [15]. Sensory data arrive from the sensors in the form of a *message*, a binary string of fixed length.³ Messages are placed in an *active message table*, a first-in, first-out (FIFO) list of messages currently undergoing processing. Each agent maintains a collection of *rules*, stored in a database or *rule table*. Rules consist of three components: a *condition string*, an *output message*, and a *credit*. Condition strings are composed from an alphabet of three possible symbols: {1,0,#}. At each time step, the message at the head of the active message table is processed by checking for a match with the condition string of each rule in the rule table. A 1 or 0 in the condition string matches the corresponding value in the message at the same index. A # matches either symbol (0 or 1). So, for example, the message 10010111 is matched by any of the condition strings 10010111, 10010##1, and #####. The condition string #####0, however, would not match.

Rules whose condition strings match the current message bid for their output message (also a bit string of the same length as sensor messages) to be placed in the active message table. This bid is achieved by calculating the rule's *strength*. Strength is the product of the rule's credit (detailed shortly) and its *specificity*. Specificity is a unit normalized value, equal to 1 – the normalised number of # symbols in the condition string. So, for example, a condition string consisting entirely of # symbols has a specificity of 0; a string with 75% # symbols has a specificity of 0.25; and so on.

For each rule that matches the current message under consideration, its strength is calculated. The rule with the highest strength is selected and then places its output message into the active message table. If more than one rule has the highest strength, then a uniform random selection is made from the winning rules. The selected rule places its output message into the active message table. Most output messages are *action messages*⁴, that is, they trigger an actuator. Action messages are removed from the table once they have been translated into actuator instructions.

The process outlined in this section is illustrated in Figure 13.4.

13.3.2.4 Credits and Payoffs

Each rule maintains a credit, essentially a measure of how useful this particular rule has been in the past. Rules begin with a default credit value and

³ A message length of 32 bits is used, but the actual length does not concern the processes described. Larger message lengths allow more bandwidth in sensor messages, but require more storage.

⁴ Action messages are distinguished from other messages by a marker bit in the string being set – all other message types are guaranteed not to set this bit.




Fig. 13.4 The rule matching and bidding process. The top message from the active message table is selected and becomes the current message. Rules whose condition string matches the current message have a strength calculated as the product of their credit and specificity. The rule with the highest strength then becomes the selected rule and its output message is added to the active message table. The current message is then discarded and the process repeats. Some messages are action messages and trigger actions.

earn or lose credit based on how useful the rule is in assisting the agent to live and mate in the world. As described earlier in this section, agents maintain an energy level and health index. The differentials of these quantities are monitored, and when they reach a certain threshold, a *credit payoff* is performed. The credit payoff rewards or punishes rules that have been used since the last payoff (held in a separate list), by altering their credit according to frequency of use and the magnitude of the change in energy since the last payoff. Further details regarding this process may be found in [18].

The credit payoff system enables rules that, over time, assist in increasing health and energy to be rewarded; those that decrease health and energy will decrease in credit. The rationale being that the next time rules have a chance to bid, if they have been useful in the past, they will probably be useful in the current situation.

The number of time steps between successive payoffs is dependent on the rate of change in the agent’s health (i.e., the magnitude of the differential). For example, if a creature is being attacked and losing health quickly, payoffs will be more frequent. The rules involved in letting the agent get hit will also decrease in credit quickly (hopefully soon being outbid by other rules that may prove more successful if the agent is to survive).

Maintaining a list of rules that have been used since the previous payoff allows rules that indirectly increase health to receive appropriate credit. For example, while the rule to “eat when you find food” is a good one, you may need to walk around and look for food first to find it. The rules for walking and turning, although they decrease health in the short term, may result in finding food. This increases health in the longer term. Overall, if such rules are helpful in increasing health, their credit will increase. A rule whose strength falls to zero will be automatically removed from the agent’s rule table, since it is no longer able to bid anything for use.

As specified in Section 13.3.2.3, a rule’s strength is the product of its credit and specificity. This is necessary, since rules that are more specific will be used less often, as they match fewer messages. Rules that are more specific will

have less chance to receive credit payoffs but still may be useful. When two or more rules with the same credit match a message, the more specific rule will have greater strength and thus will be selected over the more general one.

13.3.2.5 Agent Evolution

The credit payoff system allows rules that have contributed to the agent's survival to be used more often. However, this will only select the best rules from the currently available set. The problem remains as to how the agent can discover better rules than those it currently uses.

Genetic algorithms follow a Darwinian metaphor in that they operate as a search method over the phase space of possible phenotypes in a given system, searching over the *fitness landscape* for individuals of higher *fitness* [12, 22]. In the *Eden* system, a rule functions as the genetic unit of selection and new rules are brought into an agent's genome via the standard operations of *crossover* and *mutation* (see the references for explanations of these terms).

Recall from Section 13.3.2 that mating is a possible action an agent can perform. If two agents successfully mate, they produce a new agent whose rule table is a combination of the parents' tables. A proportion of rules from each parent selected based on their strength – the rules of highest strength from each parent being favoured for selection. These selected rules undergo crossover and mutation operations, as per the schema system of Holland [14], resulting in the creation of new rules. Mutation rates vary according to the behavior of people experiencing the artwork in the exhibition space.

Since rules of highest strength are selected from each parent and those rules may have been discovered during the parents' lifetime, the evolutionary process is Lamarckian [4]. This design decision was used to allow more rapid adaptation to changing environmental conditions: a necessary feature if the agents' in the artificial ecosystem are to adapt to the behavior of people experiencing the work in real time. Another way to consider this approach is that parents teach their offspring all the good things they have learnt so far in their lifetime: a kind of social learning.

13.3.3 Image

Representation of the entities of *Eden* is achieved using tiling patterns, loosely based on Islamic ornamental patterns [13]. Only the representation of biomass will be considered here. The visual representation of the biomass is based on a 16-tile set. A tile is chosen for a particular cell based on the neighbor relationships of adjacent cells. For the purposes of tile selection, tiles are selected based on the binary occupancy condition of the cell's neighbors, considering

only cells of the same type. For the 16-tile set, only immediate orthogonal cells are considered – thus there are 16 possible configurations of neighboring cells. Figure 13.5 shows the individual tiles and the neighbor relation necessary for the tile to be used. The resultant images formed by a grid of cells (illustrated in Figure 13.6) form a continuous mass of substance, as opposed to squares containing individual entities. These minimalist geometric textures suggest abstract landscapes rather than the iconic or literal individual representations that are common in many artificial life simulations. This design decision forms an integral aesthetic component of the work.




Fig. 13.5 Cellular tiling set for *Eden's* biomass. Each cell considers the four immediate neighboring cells (north, south, east, and west). The neighboring relations determine the image used for each cell. A function returns the bit pattern representing the neighborhood state for the cell and the tile is selected based on the supplied index. Four bits are required, each representing the four directions. The bits are encoded NESW (from most to least significant bit). The symbols above each cell pattern shown here illustrate the bit pattern and corresponding neighborhood relationships.




Fig. 13.6 Visualisation of the *Eden* world, showing rocks (solid shapes), biomass (outline shapes), and agents (circular elements).

13.3.4 Sound

One of the key elements of *Eden* is the ability of agents to create and listen to sound. A large proportion of sensor bandwidth is devoted to sound, allowing orthogonal sensing of both frequency and sound pressure (volume). Some basic physical modeling is performed on sound pressure levels, however many physical sound propagation aspects are simplified in the interests of efficiency.

13.3.4.1 Sound Generation

Actuator messages requesting sound generation need to be converted into a generated sound. As described in Section 13.3.2, actuator messages are bit strings. A portion of the string encodes the sound generation command (“sing”); the remainder encodes the sound generation data (spectral levels over a range of frequency bands). The current implementation has three distinct frequency bands, each occupying one-third of the total of the sound generation data for the “sing” actuator message (see Fig. 13.7).

When an agent “sings,” the spectral signature determined by the sing data in the actuator message is registered for the current time step. In addition, the same signature is used to drive a sonification process, so that people in the exhibition space can hear sounds that correspond to the “singing” activities of the agents. To drive this sonification process, the three frequency bands are assigned labels L, M, and H corresponding to low-, medium-, and high-pitched sounds (e.g., the majority of spectral energy in the 100, 1000,




Fig. 13.7 The “sing” actuator message contains two parts. The first is the command requesting the agent to perform a sing operation; the remainder contains the sing data: volume levels for three distinct frequency bands. Using three bits per frequency band results in 2^9 , or 512, distinct sounds.

and 10,000 Hz regions respectively). When an agent makes a sound, the corresponding selection from a precomputed library of sounds is triggered and sent to the audio subsystem. The audio subsystem does basic sound spatialisation using the four-channel audio system that is part of the artwork. Sounds are spatialised according to the position of the agent making the sound on the screen. Thus, as an agent making sound moves across the screen, that sound will appear to move with the agent to human observers. The audio subsystem allows many agents to be making sound simultaneously.

13.3.4.2 Sound Reception

Agents have a significant amount of sensor bandwidth devoted to sound reception. An agent's sound reception area is a forward-facing conical pattern that, like the sound generation, is sensitive across three separate frequency bands (see Figure 13.8). Each band has the same propagation and reception pattern, that is, there are no frequency-dependent differences in the modelling.

At each time step, the conical reception area for each agent is checked for any other agent that is singing within that area. A simple physical model [28] controls the propagation of sound through the environment.⁵ Sounds arriving at the agent's cell are summed on a per-frequency basis and the resultant sensor message is instantiated.

13.4 Interaction

The *Eden* system has a unique relationship between the physical and virtual components of the system. As shown in Figure 13.1, an infrared video camera monitors the immediate space of the installation environment, recognising the presence and movement of people in the space. A video digitisation sub-system performs basic image processing and analysis, such as background removal and feature detection. This data is converted into a stream of vectors indicating the location and movement of individuals in the exhibition area, and used to drive environmental parameters in the virtual simulation. Before discussing the details of the mappings between position and movement vectors and the simulated environment, we will present a background discussion on the rationale for the mappings used.

⁵ When sound propagates in a medium such as air at standard temperature and pressure, the perceptual mechanism for loudness behaves in an exponential way, as it does for humans. The relationship between distance and perceived levels is $L = 20 \log_{10}(P/P_o)$, where L is the sound pressure level in decibels (dB), and P_o is a reference pressure corresponding roughly to the threshold of hearing in humans [27].

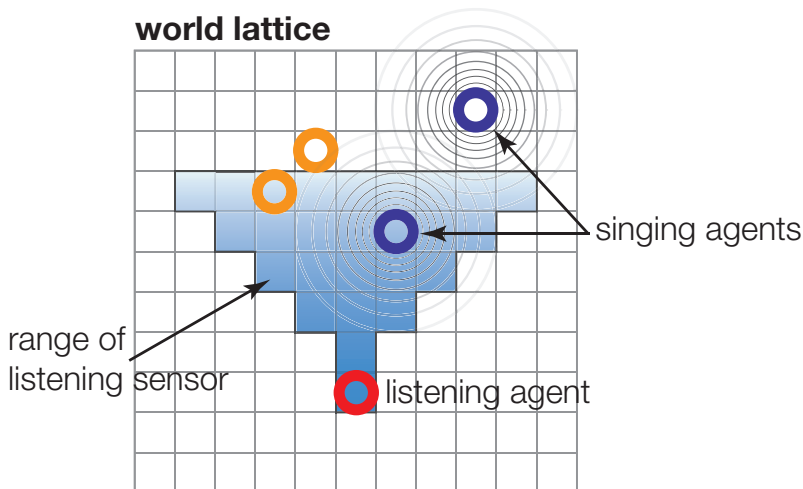


Fig. 13.8 The reception area of an agent. The listening agent will hear agents who are making sound within the blue area only. A simple physical model controls the perceptual volume levels for the agent.

13.4.1 The Problem of Aesthetic Evolution

Typically, genetic algorithms evolve toward finding maxima in *fitness*, where fitness is some criterion that can be evaluated for each phenotype of the population. Many systems define an explicit *fitness function* that can be machine evaluated for every phenotype at each generation [22].

The *Interactive Genetic Algorithm* (IGA, also known as *aesthetic evolution* or *aesthetic selection*) is a popular technique that replaces the machine-evaluated fitness function with the subjective criteria of the human operator. Aesthetic evolution was first used by Dawkins [10] in his “Blind Watchmaker” software to evolve two-dimensional, insect-like shapes. Aesthetic selection has been used to successfully evolve images [29, 30], dynamic systems [31], morphogenic forms [34, 17], even musical patterns and structures [5]. Regardless of the system or form being evolved, aesthetic selection relies on the user to explicitly select the highest-fitness phenotypes at each generation. Users typically evolve to some subjective criteria – often described as “beautiful,” “strange” or “interesting” – criteria that prove difficult to quantify or express in a machine-representable form (hence the use of the technique in the first place).

However, aesthetic evolution has two significant problems:

- The number of phenotypes that can be evaluated at each generation is limited by both screen area (in the case of visual representation) and the ability of people to perform subjective comparisons on large numbers of

objects (simultaneously comparing 16 different phenotypes is relatively easy; comparing 10,000 would be significantly more difficult).

- The subjective comparison process, even for a small number of phenotypes, is slow and forms a bottleneck in the evolutionary process. Human users may take hours to evaluate many successive generations that in an automated system could be performed in a matter of seconds.

What we would like is a system that combines the ability to subjectively evolve toward phenotypes that people find “interesting” without the bottleneck and selection problems inherent in traditional aesthetic evolution.

13.4.2 *Eden as a Reactive System*

The solution to the problem described in the previous section is to map the presence and motion data of people experiencing the artwork to the environmental parameters of the virtual environment. Thus, the virtual world in which sonic agents live and evolve is dependent not only on the simulated qualities discussed so far, but also on the presence (or absence) of people experiencing the work and their behavior within the exhibition space. This couples the real and virtual worlds into a kind of ecosystem where behaviours in one world influence behaviours in the other, creating *feedback loops* of behaviour.

Eden has no explicit fitness function. Agents continue to be part of the system based on how well they can survive and mate in the current environment. If certain selection pressures are applied, such as food becoming scarce, only those agents who can adapt and find food will prosper. By driving environmental conditions from the presence and movement of people in the exhibition space, agents must implicitly adapt to an environment that includes aspects of the world outside the simulation.

In the current system, the following mappings are used:

- *Presence in the real environment maps to biomass growth rates.* The presence of people around the screen area affects the rate of biomass growth in that local area of the *Eden* world. Areas with no people correspond to a barren environment: little biomass will grow without the presence of people in the real environment.
- *Movement in the real environment maps to genotype mutation rates.* The greater the movement of people in the space, the higher the mutation rate for rule evolution (see Section 13.3.2.5).

These mappings are based on certain assumptions. First, people will generally spend time experiencing something only if it interests them. In the context of experiencing an artwork, people generally may spend a short time evaluating their interest in an artwork, but, after a short time, if it no longer

interests them, they will leave. There may be other reasons for leaving, but in general, the duration of stay will have some relation to how “interesting” the experience is.

Agents require food to survive. If people are in the real environment, then food will grow at a more rapid rate. An agent who is making “interesting” noises, for instance, would have a better chance of keeping a person’s attention than one who is not. Moreover, an agent making a progression of sounds, rather than a just a single, repeating sound, is likely to hold a person’s attention even longer. Agents who encourage and hold a person’s attention in the space implicitly give the environment a more plentiful food supply.

The mapping of people’s movement in the space to mutation rates is based on the assumption that people will move over an area looking for something that interests them and, when they find it, will stay relatively still and observe it. Hence, the movement of people within the real space serves to inject “noise” into the genome of agents who are close to the source of movement. Higher mutation rates result in more variation of rules.⁶ If an agent or group of agents are holding the viewer’s attention, then less rule discovery is needed in the current environment, whereas if people are continually moving, looking for something “interesting,” this will aid in the generation of new rules.

Further details on the dynamics of this component of the system can be found in [19].

13.5 Results

At the time of this writing, a number of exhibitions of the work have been completed. An image from a recent exhibition of the work is shown in Fig. 13.9. A typical exhibition may last several weeks, giving plenty of opportunity for the agent evolutionary system to take into account the behavior of people experiencing the work. Certain factors have a marked effect on this behavior and require specific compensations in the software. For example, when the gallery is closed, there will be no people in the space, which diminishes the food supply for the agents. Thus without compensation for gallery opening hours, the entire population dies out each night!

Analysis of the rules agents use shows that sound is often used to assist in mating, as would be expected [37], and with the influence of people, sound is used in other ways as well. Once the environmental pressures from audience behavior are incorporated into the system, the generation of sound shows a marked increase and analysis of the rules discovered shows that making sound is not only used for mating purposes.

⁶ Most child rules that mutate will not be “better” than the parent rule, but, in general, the use of mutation does provide the possibility for the system to discover rules that would not be possible by crossover alone.




Fig. 13.9 An image of *Eden* in operation.

The resulting sounds produced by the system *do* appear to interest and engage the audience, and combined with the visual experience, *Eden* does incite, from visitors, an interest and curiosity in what it is doing. In many cases, people are not aware of the learning system, camera sensing, even the fact that what they are experiencing is a complex artificial life system.

Since the system is reactive to people (rather than interactive), there is no correct or incorrect way to behave except to appreciate the experience. Anecdotal accounts from people who have experienced the work describe it as “having a sense that it is somehow alive,” or “like being in a strange forest at night.” In a number of exhibitions, people returned to the work over a period of several days, to see how the qualitative behaviour of the virtual environment had changed. In one recent exhibition, a local businessman visited the work during his lunch-hour every day for 3 weeks, describing the experience as “fascinating. . . one that made me more sensitive to my own environment.” While these are, of course, subjective evaluations, it does appear that *Eden* is able to adapt and evolve to create an ongoing interest for its audience.

13.6 Conclusion

This chapter has described a novel evolutionary ecosystem, where agents make use of sound to assist in survival. While the main impetus and methodologies are based around the development of an artistic system, it is hoped that some of the ideas presented here may be of interest to those working in artificial life from other perspectives or with different agendas and applications.

In summary, a system has been produced that attempts to integrate the open-ended nature of synthetic evolutionary systems into a reactive virtual space. The approach used has been to measure components of the real environment, incorporating them into that of the virtual one, thus enabling a symbiotic relationship between virtual agents and the artwork's audience, without need for explicit selection of phenotypes that engage in "interesting" behavior.

Further information, including sample sound recordings and video documentation of the work, is available on-line [20].

References

1. Ashby WR (1956) *An Introduction to Cybernetics*. Chapman & Hall, London.
2. Bentley PJ (1999) *Evolutionary Design by Computers*. Morgan Kaufmann Publishers, San Francisco, CA.
3. Bentley PJ and Corne DW (eds.) (2002) *Creative Evolutionary Systems*. Academic Press, London.
4. Bowler PJ (1992) Lamarckism, In: Keller EF and Lloyd EA (eds.). *Keywords in Evolutionary Biology*, Harvard University Press, Cambridge, MA, pp. 188–193.
5. Bulhak A (1999) *Evolving Automata for Dynamic Rearrangement of Sampled Rhythm Loops*. In: Dorin A and J McCormack (eds.) *First Iteration: A Conference on Generative Systems in the Electronic Arts*, CEMA, Melbourne, Australia, pp. 46–54.
6. Cariani P (1991) *Emergence and Artificial Life*. In: Langton CG et al. (eds.) *Artificial Life II*, SFI Studies in the Sciences of Complexity, Vol 10, Addison-Wesley, Reading, MA, pp. 775–797.
7. Cariani P (1993) *To Evolve an Ear: Epistemological Implications of Gordon Pask's Electrochemical Devices*. *Systems Research* 10:19–33.
8. Conrad M and Pattee HH (1970) *Evolution Experiments with an Artificial Ecosystem*. *Journal of Theoretical Biology* 28:393.
9. Dahlstedt P and Nordahl MG (1999) *Living Melodies: Coevolution of Sonic Communication*. In: Dorin A and McCormack J. (eds.) *First Iteration: A Conference on Generative Systems in the Electronic Arts*, Centre for Electronic Media Art, Melbourne, Australia, pp. 56–66.
10. Dawkins R (1986) *The Blind Watchmaker*. Longman Scientific & Technical, Essex, UK.
11. Emmeche C, Køppe S and Stjernfelt F (1997) *Explaining Emergence: Towards an Ontology of Levels*. *Journal for General Philosophy of Science* 28:83–119.
12. Goldberg DE (1989) *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA.

13. Grünbaum B and Shephard GC (1993) Interlace Patterns in Islamic and Moorish Art. In: Emmer M (ed.) *The Visual Mind: Art and Mathematics*, MIT Press, Cambridge, MA.
14. Holland, JH (1992) *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. MIT Press, Cambridge, MA.
15. Holland JH (1995) *Hidden Order: How Adaptation Builds Complexity*, Addison-Wesley, Reading, MA.
16. Langton CG (1989) Artificial Life. In: Langton CG (ed.) *Artificial Life*, SFI Studies in the Sciences of Complexity, Vol. 6, Addison-Wesley, Reading, MA, pp. 1–47.
17. McCormack J (1993) Interactive Evolution of L-System Grammars for Computer Graphics Modelling. In: Green D and Bossomaier T. (eds.) *Complex Systems: From Biology to Computation*, ISO Press, Amsterdam, pp. 118–130.
18. McCormack J (2001) Eden: An Evolutionary Sonic Ecosystem. In: Kelemen J and Sosik P. (eds.) *Advances in Artificial Life*, 6th European Conference, ECAL 2001, Springer, Berlin, pp. 133–142.
19. McCormack J (2002) Evolving for the Audience. *International Journal of Design Computing* 4 (Special Issue on Designing Virtual Worlds). Available online: <http://www.arch.usyd.edu.au/kcdc/journal/vol4/index.html>.
20. McCormack J (2006) Eden: An Evolutionary Sonic Ecosystem, web site, <http://www.csse.monash.edu.au/~jonmc/projects/eden/eden.html>.
21. McCormack J and Dorin A (2001) Art, Emergence and the Computational Sublime. In: Dorin A (ed.) *Second Iteration: A Conference on Generative Systems in the Electronic Arts*, CEMA, Melbourne, Australia, pp. 67–81.
22. Mitchell M (1996) *Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA.
23. Nagel E (1961) *The Structure of Science: Problems in the Logic of Scientific Explanation*. Routledge, London.
24. Pattee HH (1988) Simulations, Realizations, and Theories of Life. In: Langton CG (ed.) *Artificial Life*, Vol. VI, Addison-Wesley, Reading, MA, pp. 63–77.
25. Reichardt J (1971) *Cybernetics, Art and Ideas*. New York Graphic Society, Greenwich, CT.
26. Reichardt J (1971) *Cybernetics, Art and Ideas*. In: Reichardt J (ed.) *Cybernetics, Art and Ideas*, Studio Vista, London, pp. 11–17.
27. Roads C (1996) *The Computer Music Tutorial*. MIT Press, Cambridge, MA.
28. Roederer J (1975) *Introduction to the Physics and Psychophysics of Music*. Springer-Verlag, New York.
29. Rooke S (2002) Eons of Genetically Evolved Algorithmic Images. In: Bentley PJ and Corne DW (eds.) *Creative Evolutionary Systems*, Academic Press, London, pp. 339–365.
30. Sims K (1991) Artificial Evolution for Computer Graphics. *Computer Graphics* 25:319–328.
31. Sims K (1991) Interactive Evolution of Dynamical Systems. In: *First European Conference on Artificial Life*. MIT Press, Cambridge, MA, pp. 171–178.
32. Sommerer C and Mignonneau L (eds.) (1998) *Art@Science*. Springer-Verlag, Wien, Austria.
33. Todd PM and Werner GM (1998) Frankensteinian Methods for Evolutionary Music Composition. In: Griffith N and Todd PM (eds.) *Musical Networks: Parallel Distributed Perception and Performance*. MIT Press/Bradford Books, Cambridge, MA.
34. Todd S and Latham W (1992) *Evolutionary Art and Computers*. Academic Press, London.
35. Ulam S (1952) Random Processes and Transformations. *Proceedings of the International Congress on Mathematics*, Vol. 2, pp. 264–275.
36. Watson AJ and JE Lovelock (1983) Biological Homeostasis of the Global Environment: The Parable of Daisyworld. *Tellus* 35B:284–289.

37. Werner GM and Dyer MG (1991) Evolution of Communication in Artificial Systems. In: Langton CG (ed.) *Artificial Life II*, Addison-Wesley, Redwood City, CA, pp. 659–682.
38. Wiggins G et al (1999) Evolutionary Methods for Musical Composition. Proceedings of the CASYS98 Workshop on Anticipation, Music & Cognition.
39. Wilson S (2002) *Information Arts: A Survey of Art and Research at the Intersection of Art, Science, and Technology*. MIT Press, Cambridge, MA.
40. Witten TA and Sander LM (1981) Diffusion-Limited Aggregation, a Kinetic Critical Phenomenon. *Physical Review Letters* 47:1400–1403.