Chapter 14
Enriching Aesthetics with Artificial Life

Alan Dorin

Perhaps one day our machines will be able to take us beyond environments that look more and more familiar as we traverse them, into spaces that continue to surprise us as we explore. Can our engineered artifacts ever rival the intricacy of nature?

This chapter explores the application of artificial life software for the purposes of generating dynamic works of audio-visual art that act to some degree independently of the artist. Autonomous artworks slot neatly into artificial life's fascination with the construction of digitally instantiated living systems. This interest itself reflects the long-standing human desire to construct "living" images, clay statues and idols, hydraulically powered figures, mechanical toys, and electrical robots. The chapter focuses on the elements of dynamic software art that relate to the concept of the *sublime* in nature and computation. It is suggested that artificial life and virtual ecosystems provide an avenue for the production of rich, dynamic works of software art suited to this exploration.

14.1 Introduction

Artificial life is studied largely as a means of furthering our understanding of biology and of complex adaptive systems in general. While it has demonstrated potential in a number of fields, in particular as a means of solving engineering problems, artificial life techniques have also been applied by some in the art community. The field promises to continue to enrich artistic practice and our approach to contemporary aesthetics, even as its initial flash of popularity wanes. It is this continued application to aesthetics which the present chapter begins to address.

The techniques of art based in artificial life form a subset of *generative art*. This is an artistic practice that adopts an aesthetic of *process*. This means that although the final outcome of a work may depend for its appeal on the aesthetics of an image, sound, sculpture, or other form, the process that generates it is also significant. The generative artist is responsible for setting up initial conditions and a process to act upon them. The work that unfolds is the result of this series of changes. This is analogous to a biological phenotype (usually an organism and its behaviour) being the result of the physical and chemical interactions that govern its development from a genotype (its genetic material as stored in DNA). Hence, there are conceptual connections between generative art and artificial life as well as practical ones.

Generative/process-based art is no longer treated as a fresh field for aesthetic exploration. Fascination with its possibilities seemed to fade sometime after its heyday in the late 1960s. This of course was the time of the *Cybernetic Serendipity* exhibition in London, Jack Burnham's text *Beyond Modern Sculpture*, and much other activity linking art and computer technology in innovative ways. Recently, artists involved in artificial life have retrodden some of the ground cleared by their predecessors in cybernetic art and have cleared some new space for themselves. Now that artificial life is also unfashionable, perhaps a serious assessment of its past and future contributions can be made while side-stepping the hype that initially accompanied the field.

The following section begins by discussing the sense of wonder people often feel while contemplating nature or simple physical systems. This is then related to the tradition of the *sublime* in aesthetics. The chapter then discusses means of employing artificial life techniques to explore the *computational sublime* such that a computational system emulating the physical world's capability to generate complexity and novelty might be devised. In particular, the chapter describes the application of metaphors adopted from Ecology to create virtual ecosystems. This avenue has been explored by a few artists to whom it appears to offer the potential to generate and control complexity in creative software systems.

14.2 Wonder and the Sublime in Art and Nature

People stare contemplating the ocean as it swells and crashes on a rocky shore. They gaze fixedly into a fire until the sting of smoke raises their awareness. People may lie on their backs and follow passing clouds or marvel at the glittering of stars. There are many things that fascinate us, that mesmerise us, that cause us to forget ourselves and our situation as we become lost in a timeless appreciation of nature.

There are also circumstances under which we may have an experience that reminds or forces us to consider our insignificance: the feeling that accompanies the vastness of spaces such as the ocean, a desert plain, or an endless mountain range, or the feeling of insignificance when contemplating the age of the universe or the years over which a trickle of water has eroded a canyon, for example. These expanses can nevertheless be made subject to reason. We are mentally equipped to discuss the concept of infinity, even if we are unable to quite fathom it intuitively. This paradoxical experience Kant has labeled the *mathematical sublime* [15].

Likewise, Kant introduced the *dynamically sublime* as relating to encounters with the ferocity of nature and the sense of vulnerability this entails, coupled with the triumph of reason over fear. In his own writings on the sublime, Edmund Burke took the view that the sensation of incomprehensibility, the fear of hopelessness or of danger, coupled with the knowledge that one *is* able to reason about something beyond one's senses or one is not inadequate or in danger, causes a kind of delight through internal conflict – a sublime experience [19]. That is to say, there is an element of the sublime, perhaps not an artistic element, in resisting the urge to flee as a tiger roars behind bars. One's body responds with fear to the situation, but reason easily overcomes this and forces the body to stand its ground – generating a sense of delight tinged with terror. In the case of a painting, the viewer's separation from a wild scene of stormy seas or a vast desert is not enforced by iron bars, but by the picture plane. This is coupled with the knowledge that "it is only a picture."

How can a work of art, specifically a work of computer-based, generative software, approach the sublime? While there is a vast amount of literature on the sublime in art dating back as far as Longinus [23], the focus of this chapter is narrower. An earlier publication [25] outlined the *computational sublime*. This arises from viewing the computer in terms of its capability to perform logical operations at a rate and on a scale vastly outside our own abilities in this area. Due to its speed, the computer is able to mediate between our human perceptual apparatus and (practically) infinite computationally defined spaces. Yet, since we are the makers and programmers of these machines, our power of reason is not only able to overcome but also to *define* these very spaces that our senses are unable to grasp fully.

Before discussing computer-based art in more detail, some much simpler artifacts, possibly just for a moment, rival natural wonders in the fascination they hold. How does this relate to the sublime? It seems that there are a number of methods by which the sensation may be encouraged and that some of these have been recreated in artefacts specifically for this purpose. They are categorised below for convenience.

Marking Time

Maybe the simplest examples of our mesmerizing creations include a stream of liquid running from a water-clock or fountain, the shifting sands of the hourglass, and the oscillating pendulum. Each of these marks time in its own manner: one a continuous stream; the next a finely particulate flow;

and the last in clear, discrete stages. The eternal flow, the innumerable sand particles, the never-ceasing oscillations confront us with the infinite through this visible marking of time. Although these natural processes were utilised in desk ornaments to amuse the bored office worker of the 1980s, even in "serious" art, simple processes like this may give a sense of the sublime.

Exposing Spaces

The wire-suspended *mobiles* of Calder successfully employ mechanical processes in a manner accepted within the art world [29]. Calder's playful pieces are captivating and elegant for all their simplicity. Their workings are laid plainly before the viewer, all that they are is apparent at a glance – and yet this is not so, for their movement brings a vitality and opens a space the static sculpture does not possess. The universe a mobile sweeps out is contained within its wires, rods, and solid forms, so in one sense they may be held in (and created from) the palm of a hand. Yet, as they are touched by invisible air currents, their inner complexity is exposed.

Intricacy

Many a visitor has been fascinated by the button-driven clockwork and gearing of exhibits in traditional science museums. Here, a sense of wonder at the machine in its entirety arises, but also a fascination with the intricacy of the mechanism. Each gear meshes with another, each component is configured "just so," and together the pistons and wheels turn in harmony to produce a composite that might drive a clock, crush a quartz boulder, pump water, or power a vehicle. This is not the beauty of a crashing ocean or a sunset, but the charm of peering into an ant's nest or through a microscope at a drop of pond water. There is something in these systems that causes one to marvel at a complexity that is just beyond grasp.

Defying the Natural Order

There is still another wonder to be described here, that of somehow defying the natural order of things. It appears wonderful that a huge boulder might come to be balanced on a slender rock column, that a gyroscopic top remains on its tip despite interference, that a bird or a massive aeroplane can remain suspended in the air or a colossal steel ship can float, and that a magnet can push another away without touching it. These things, of course, are dealt with to some extent by simple science. Sadly, only children may wonder at these things. Yet, implicitly, these interactions remain in need of continual re-explanation since each instils apprehension through not being quite "right." That boulder or spinning top might topple at any moment. That bird should fall from the sky and the ship ought to take on water. The magnets are behaving unnaturally. All of these systems cause one to ponder "How can this be?" even if one can reason about the answer through science.

Curiosity

Related to all of the above phenomena, in particular the previous one, attention can also be held by riddles and intellectual pursuits. Included in this are mind games such as chess, paradoxes, and mathematical puzzles, but also scientific enquiry. These all captivate us through our drive to learn why something behaves as it does, especially if it seems to defy the natural order of things. Of the algorithmic art exhibition at Xerox PARC in 1994, Bern writes: "The appearance of algorithmic visual art should raise the temporal questions of 'How is this made?' and 'Can this be extended?' " [3]. While the use of the term *should* is of questionable accuracy, nevertheless there is something of relevance in the idea that algorithmic art may invite a question and arouse curiosity as much as it satisfies visual aesthetic criteria associated with traditional media.

Any of the preceding devices (and no doubt many others) may be used to convey a sense of the sublime. All of them expose the limitations in our comprehension, experience, or endurance. However, simultaneously, we have the power to reason about these phenomena, to encapsulate them in a word or other representation, or we may just have the opportunity to walk away unharmed. Of course, the context in which the preceding occurrences and objects are encountered will influence the extent to which the sensation derived relates to art. However, the potential to employ any of these basic devices for artistic purposes in search of the sublime experience is present.

14.3 Sublime Software

Now we return to the issue at hand: writing artificial life software for artistic purposes. There are, of course, infinitely many reasons why an artist might wish to do this. Let us suppose in this instance that the goal was software that surprised not only a viewer but also the artist who fashioned the work. This is in keeping with a proposal made elsewhere by the author [10]. The surprise required might not be a trivial one-off shock, but a genuine and continuing fascination with the novelty of the outcome generated, combined with a sense of being lost in a complex world beyond one's grasp. The work requires an aesthetic quality that exhibits the character listed above as *intricacy* and also at the level listed as *curiosity*.

It so happens, of course, that the techniques of artificial life are well suited to this application, Conway's *Game of Life* cellular automata (CA) being a case in point [17]. CAs are fascinating for their ability to (practically, if not theoretically) produce ever-fresh patterns. Like clockwork automata or Calder's mobiles, the patterns produced may fall well within a limited domain, often even a cycle. However, especially in a large-scale configuration, there is always some aspect of the system that remains to be discovered in the vast universe these interactions define.

The innovative design pair, the Eames, understood this aspect of a complex visual field. They used their understanding effectively in the 1964 New York World Fair exhibit for IBM to build a multiscreen cinema. The screens displayed related material simultaneously in such a way as to make it impossible to view all of the footage in detail. Viewers could watch the spectacle as a whole and see it as a multiscreen montage or they could pick and choose elements to observe in detail – much as one examines a complex clockwork machine, a microscopic world, or a beehive – piecemeal.

On a related topic, Tufte indicated the benefits of presenting information in parts grouped carefully together, "Panorama, vista and prospect deliver to viewers the freedom of choice that derives from an overview, a capacity to compare and sort through detail" [33]. Such considerations also assist painters in creating worlds that reward careful study of their detail. For example, Hieronymous Bosch's *Garden of Delights* (ca. 1500) portrays a world over which the eye may wander at will as it takes in relationships between couples, groups of people, sections of garden, and so on. Of course, where the object under study is static, a viewer may take the time to examine and re-examine the various parts of a display and avoid missing anything of importance (as Tufte would prefer). When the object is constantly changing, as with multiscreen cinema or a *Game of Life* run, an investigation of one element results in irrevocably missed information about another. Under these circumstances, no matter where a viewers' attention is focused, they are left with the impression that an abundance of chances were missed. Their faculties are fundamentally unable to absorb all that is before them.

While one of the attractions of the *Game of Life* occurs at the visual level, the software is also an attractive subject for philosophical discussions concerning emergence and complex systems [5]. It contains a vast number of possibilities within the bounds of its simple virtual physics. While its transition rules may be considered in terms of biological analogy (overcrowding and mutual support of living cells, for instance), the result is still an intrinsically digital system. Cells are in one of exactly two states, and their behaviour is completely deterministic. Yet, from this emerges a milieu of flickering forms that somehow *suggest* life, without mimicking it.

The question "How does it work?" is implied by any given run of the *Game of Life* in which it is recognised that the output is not at all random, but highly organised and structured according to its own internal rules. Even once the viewer know these rules, that they are capable of creating such a bewildering outcome remains a source of fascination to the artificial life researcher and to the newly informed student alike.

14.4 The Betrayal of Points and Lines

Having briefly discussed the concepts of wonder and the sublime, this section now addresses the representational schemes that may be used to bring the digital realm to the realm of sublime experience. Of particular interest is the way in which points of light may be used to give the bounded computer screen the appearance of *intricacy* and an ability to expose *unbounded* space (see the earlier categorization of methods for approaching the sublime). This has already been touched upon while discussing the *Game of Life* in the previous section.

Taking the lead of Kandinsky, the discussion begins with the geometric point, an invisible thing too perfect to exist [21]. An artist may approximate it by instigating a collision between a sharp tool and a flat plane. The artist's point has character. It is not one of Plato's Ideal Forms, but it is beautiful for all its imperfection.

The point may also be displayed on a CRT or LCD monitor. The size and shape of the smallest displayable shape are dictated by the display technology. Beyond a single pixel, the character of the point may also be altered, albeit with less potential for variation than in the analog world.

Anyone familiar with Islamic or Roman mosaics, medieval tapestry, Pointillist painting, the television screen, or a computer monitor is aware that, in vast numbers, the point may vanish in an ocean of its kind. Here it becomes simply one atom in a larger texture – mobs of points are not as honest as individuals.

If the draftsman drags a pen across the page, another form results: a line. Like the point, the line seems unpretentious, especially if it is straight and completely surrounded by white space. However, place a few straight lines together and important transformations may take place. For example, the *triangle* is born of three lines and is eventful for its ability to enclose space on the plane and thereby define a boundary and, with it, an object. The second transformation of relevance here is the apparent shift into three dimensions brought about by the *Cartesian axes*. This is also born of three lines and is eventful for its ability to *suggest* space that extends beyond the plane in which the lines are drawn.

Under some circumstances the three lines that represent Cartesian space may appear simple and flat; under others, however, the third dimension "pops" out of the plane at the viewer. It is outside the scope of this chapter to discuss how this second transformation occurs; the interested reader is referred to [20]. For now, it is sufficient if it is acknowledged that at least in some cases, three Cartesian axes suggest a volume where there is none, this being (of course) the principle that underlies Renaissance painting and a host of styles up to and, of relevance to this discussion, including modern computer graphics. Using only lines, one may build a two-dimensional representation of a solid object in unbounded three-dimensional space.

To return at this stage to the *Game of Life*, the system is typically displayed as, strangely enough, a grid of points. The choice of the point as the basic element of the simulation excludes the misleading suggestions of line drawings. Still, even within this limited visual space, the CA is successful at its evocation of life and with it *intricacy*. How can similar systems expose or suggest spaces that are infinitely large *and* intricate?

14.5 Moving Beyond Two Dimensions

The *Game of Life* and CA-based generative art such as *IMA Traveler* [12] work in the domain of points and suggestion, as outlined. With a computer monitor it is easy to represent lines and surfaces using a multitude of similarly lit pixels.

IMA Traveler, while working with points in a plane, recursively sub-divides these, giving the viewer the sensation of a bottomless plummet. This is an interesting perceptual phenomenon when it is considered that by continuously zooming in on a two-dimensional image (the points are not modeled in three-dimensional space), one can instil the sensation of motion through a continuous and infinitely detailed universe. The fractal zooms of the 1980s are perhaps the most familiar and overused form of this same trick. In their work *Powers of Ten* (1968), the Eames also utilised rapid scale shifts to produce a film of great effect.

The CA and its artistic derivatives draw attention for evoking natural phenomena and for prompting the viewer to determine the underlying principles by which they operate. Not all artificial life software is of this type, even if it is important to the research field for other reasons. One may propose a full range of works of which the CA-based software falls near the center. This spectrum runs from abstract systems such as Ray's *Tierra* [28] to clearly representational systems such as Sims' *Virtual Creatures* [31].

Ray's work emphasised the principles of selection and evolution through rapid reproduction and has itself proved to be a useful model and a starting point for other researchers such as Pargellis [26] (discussed later) to further the study of evolutionary systems. However, Ray's work does not have the visceral appeal that visual cellular automata may have. *Tierra*'s space of machine code instructions is abstract and not typically represented in such a way as to be comprehended by the eye. A fascination with *Tierra* arises through careful study and is predicated upon an understanding of the core workings of the system and the way in which the instructions interact and occupy memory. This contrasts with the CA, in which the rules need not be known to understand the system at one level. It is this first level of visceral comprehension that prompts theorizing about what underlies the behaviour of the CA cells.

At the opposite end of the spectrum, Sims' *Virtual Creatures* fascinate researcher and layperson alike. His jumping, swimming, running and limping creations are so full of character that one easily takes a leap of faith by referring to them as "creatures," even though their bodies are clearly rendered cuboids. This leap erases any chance of a viewer posing the question "How does it work?" because it is implied by the visualization that these *are* creatures! While a graphics researcher or software engineer may see Sims' video and wonder at the means of producing such marvelous results, this question is not implied by the creatures' visualization as it is by the CA. The virtual creatures do not operate according to unknown rules; we are tricked into believing that they operate according to the rules all creatures obey. There is seemingly nothing here to discover.

This aspect of Sims' creatures is largely due to the way in which they are visualised – the representations of three-dimensional space, of solids and their surfaces, of friction and other forces are all customary. In this case, the visualization is prescriptive, rather than evocative. McCormack's *Eden*, which was constructed as an artwork (see Chapter 13 in this volume and [24]), similarly represents organisms and their real-world behaviours. However, it displays them using iconic forms on a two-dimensional grid (mapped to an "X" in three-dimensional space) and in a world nevertheless governed by rules of survival, energy, mating, and space, based on those of the real world. In contrast with Sims' work, this underlying link to the physical world is not prescribed and takes some experience to decipher – if a viewer is able to decipher it at all. Although the system is still considerably more representational than the *Game of Life*, what remains is an aesthetic visual experience akin to viewing an intricate CA, coupled with the implicit question "How does it work?"

This discussion of Ray's, Sims', and McCormack's work is in no way intended as an evaluation of their worth. These examples only serve to illustrate the various ways in which a viewer may contemplate a generative process and its outcome and, therefore, the various ways in which the devices may be employed for artistic purposes. The next section takes current generative artworks and artificial life software as a starting point. It explores areas into which artists might move to further expand the sense of wonder their works inspire.

14.6 Spaces That Build Themselves

An artist aiming to produce a system capable of sustaining a continuous increase in complexity shares a goal with many artificial life researchers [2]. Why would an artist wish to do such a thing? To return to the ideas discussed earlier, the artist searching for the computational sublime would find it, perhaps in its ultimate form, by generating an experience of open-ended complexity governed by processes instantiated on a computer. The loss of control would be complete – the system would build its own structures and define its own universe. It would do this according to human-engineered rules, yet in a way that defies humans to anticipate its outcome.

This system would be following the code a programmer laid down and run on a machine an engineer designed. Would this device behave in a way about which we could easily reason? In this conflict between control and riot lies the sublime. Mary Shelley knew this well – her friends and contemporaries were much interested in the sublime – when she vividly penned *Frankenstein* and his monster run amok [30]. For a historical overview of theories of the sublime in this period, see [19].

How *does* one write code that will produce the hierarchically organised composite structures associated with life and ever-increasing complexity? If we consider the cells of a CA grid as analogous to molecules and, consider higher-level emergent structures such as gliders as analogous to organelles (a far stretch when one considers the complexity of interactions a physical molecule or an organelle may undergo, and the feeble interactions between neighbouring cellular automata), can we code the system so that still larger-scale groupings of structure occur? Can it produce structures at the level of single cells, a multicellular organism, or an ecosystem?

In theory, even gliders, spinners, and other structures of the *Game of Life* may be carefully arranged into larger-scale units (such as a self-reproducing machine incorporating a universal Turing machine, if one has the patience to arrange its 10^{13} cells [5]). The question remains though "Is it possible that such a higher-level structure will appear of its own accord?" If software could be arranged to facilitate this, the structures that arose would do so on their own terms and *might* therefore behave in ways not envisaged by the creator. This *might* provide an effective source of complexity to assist an artist in his search for the computational sublime.

If theories about the process of natural increase in complexity (such as those extensively discussed by Kauffman, for example [22]) hold true in virtual systems, then the elements in the virtual space might self-assemble into simple stable structures. Perhaps, if the virtual physics and chemistry of the world allowed it, simple reactions might occur, possibly some in autocatalytic and cross-catalytic sets such as those described by Dorin [6]. These might form the basis of a recognizable topology with the bare bones of a metabolism. What next?

As natural evolution demonstrates, one way to achieve an increase in complexity beyond this is to have the structures engage in a reproductive battle against one another for resources (see [32]). Pargellis' system *Amoeba* manages to initiate reproduction randomly. It establishes conditions in which practical CPU and memory resources are sufficient for a replicator to appear spontaneously from the prebiotic soup [26]. In a *Tierra* run, such an event is much less likely than in *Amoeba*. Hence, Ray initially seeded *Tierra*'s population with a replicator to allow evolution to commence. The problem of coding a simulation that can make the leap from self-organization to spontaneous evolution of structure seems (as far as this author is aware) unmade by any researcher.

Setting aside the leap from self-organization to evolution, even though systems solely employing artificial evolution are readily implemented, getting these to mimic natural evolution's progression from molecules to organelles to cells and on to multicellular creatures and ecosystems has proved a stumbling block. Although worlds such as *SOCA* give rise to auto-catalytic and cross-catalytic sets, *Amoeba* may randomly give rise to replicators and *Polyworld* [35] gives rise to simple communities, there has been limited success (arguably no success) in writing software that encompasses more than one of these important level shifts without resorting to abstractions so high that the simulations they are contained within become trivial.

The reasons for this difficulty are not yet clear. In part, current computational resources may be to blame. However, this is quite likely only a part of the story, and maybe only a small part at that. Rasmussen et al., for instance, has suggested that the bottom level of our simulations are not complex enough to give rise to the kind of multilayered outcomes we desire. They proposed that only by adding complexity to the bottom-level elements of a simulation can we expect to gain extra levels of organised structure on top of any earlier ones [27]. This claim seems to contradict the "complex systems dogma" that explicitly treats complex phenomena as emergent from simple interactions. Since in Rasmussen et al.'s work, notions of "complexity" and "adding complexity" are only loosely defined, it is not clear exactly how and to what extent this might be the case. This is discussed in detail elsewhere [8].

Rasmussen et al. supported their view with a claim about a model they had constructed. Gross and McMullin [18] argued that this model does *not*, in fact, demonstrate the emergence of multiple levels and that a similar outcome can be obtained without adding complexity at the base layer.

It is outside the scope of this chapter to become too deeply embroiled in this battle. Rasmussen et al.'s suggestion may in some sense prove true. Either way, further questions need to be addressed simultaneously. Might there be a limit to all this "complexity adding"? How much added complexity is necessary to move from one level to the next? Questions like these remain open and continue to be debated.

Besides issues raised above, there may be fundamental problems with current approaches to solving the problem. It is possible that simulations on current computer architectures and employing computer programs as they are currently understood will turn out to be practically limited in their ability to produce the kind of truly open-ended complexity increase required. Issues of available resource consumption are an obvious reason why infinite increase is impossible; however, are there reasons why *any* interesting string of increases in complexity may be impossible with current programming and computer technology?

One potential means for overcoming these limitations has been identified – several artists, the author included, have employed *virtual ecosystems* in order to harness the complexity of interactions between organisms and their environments over evolutionary time periods. The goal of one day creating a virtual space as multifaceted as nature may perhaps be realised by adopting ideas from Ecology. This idea is discussed next.

14.7 An Ecosystemic Approach to Writing Artistic Software

A virtual ecosystem is a software simulation of organism and environmental interactions within a real ecosystem. There is a tradition in artificial life, ecology, and even other fields, of constructing such models. For instance, some early systems include *Polyworld* [35], *Tierra* [28], *Echo* [16], *Sugarscape* [14] and *Avida* [1]. These all model relationships amongst organisms within virtual environments inspired by biological counterparts. Purely creative and artistic applications of these systems have also been devised. The author has been directly involved with a few as follows:

Listening Sky is an interactive sonic virtual reality environment in which evolving inhabitants, represented visually as patches of particles, move across the surface of a globe singing to one another and passing their songs onto their offspring (Fig. 14.1). Songs are employed to attract mates and therefore govern the fitness of individuals [4].

Meniscus is an interactive work in which virtual invertebrates breed and swim. The invertebrates are visualised as a series of connected discs with tufts of cilia-like hair (Fig. 14.2). Humans can adjust the depth and agitation of the water to differentially favour the creatures that evolve. The creatures have preferred depths and certain levels of agitation they find favourable for breeding; hence, humans indirectly control the virtual population and therefore the appearance of the work at any time [9].

Eden is an immersive sonic and visual composition written by Jon McCormack in which an evolving population of creatures roam a space of virtual plants and minerals that is projected into a gallery onto central screens arranged in a large X. The creatures communicate their whereabouts and the location of resources to one another audibly. As they produce offspring adapted to their environment, musical patterns emerge for mating calls, food

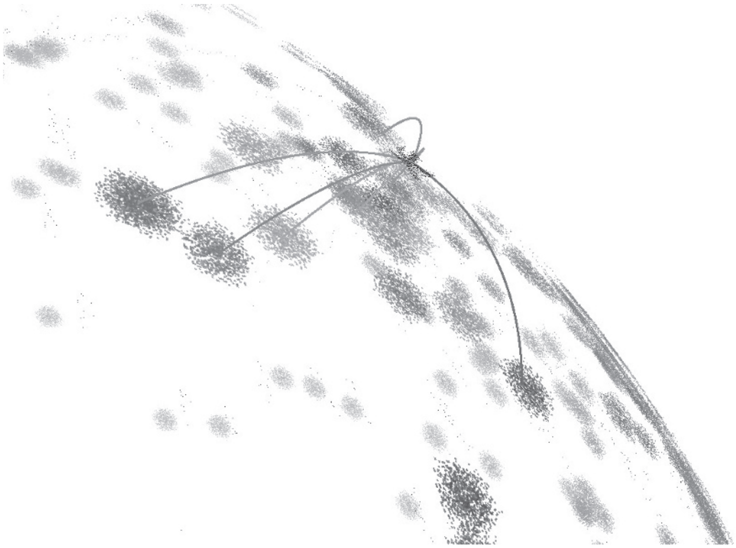

Fig. 14.1 *Listening Sky*, generative, interactive virtual-reality installation (detail) [4].

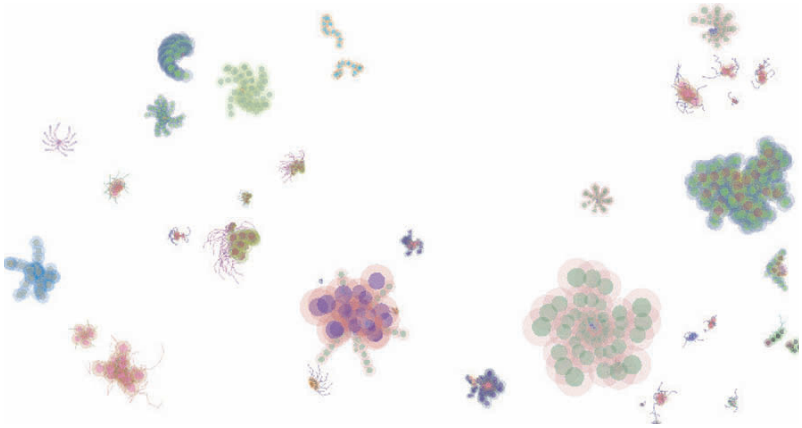

Fig. 14.2 *Meniscus*, generative, interactive installation (detail) [9].

indicators, and so forth. Humans indirectly alter the conditions in the virtual environment by moving in physical space around the screens monitored by infrared sensors [24].

Autumn Squares is a textural, tapestry-like generative animation in which populations of coloured rectangles roam a two-dimensional grid. The grid is representative of the paths through any human construction (a city, an office building); the rectangles (people who populate the construction) wander down its paths meeting and avoiding one another depending on the kind of "boxes" they are (or "fit into"). Rectangle communities form fanning, intermingling clusters of colour in an evolving visual field [10].

Plague (Diseased Squares) is an epidemiological visualisation and sonification exploring the impact of disease on the population dynamics of simple, evolving coloured agents. A replicating disease is spread by inter-agent contact. It infects agents of particular colours more or less virulently depending on the disease's co-evolved preferences for hosts. Agents attempt to evolve immunity to the co-evolving diseases whilst being constrained by their need to attract mates [11].

The interactive aesthetic selection and artificial evolutionary technique introduced by Richard Dawkins and popularised for computer graphics by Karl Sims has been widely used in artificial life art. One merit of the technique is its avoidance of prespecified, explicit fitness functions [7]. Whereas the need for these functions may not present many difficulties in engineering applications where the final goal is well understood (e.g., a bridge must span a river without collapsing under load, a sorting algorithm must order numbers, etc.), in artistic and creative endeavours these functions are as troublesome to specify as the question "What is beauty?" is to answer. A drawback of the interactive technique for artistic exploration is the bottleneck introduced by the user's manual selection of fit individuals for reproduction. Only a single evolutionary chain may be explored in this manner at any one time. Selections of one or a few fit individuals are usually made from a tiny population of only around sixteen to twenty-five individuals. Many potentially interesting avenues will always remain unexplored and even unimagined.

In contrast, the "ecosystemic" approach permits simultaneous, multidirectional and automatic exploration of a space of virtual agent traits without any need for a prespecified fitness function. Instead, the fitness function is implicit in the design of the agents, their virtual environment, and its physics and chemistry. If the system is designed to mimic real ecosystems in its formation of niches and habitats, it gives rise to parallel and simultaneous evolutionary trees that interact with one another over evolutionary time periods. Although even this system cannot explore vast areas of the potential phenotypic space, it is at least able to cope with large, interacting populations of phenotypes. There are several ways in which the design of such a system may be used as the basis for generative software art. A few include the following:

- Agents' individual appearance or sonic outbursts may be artistic works (*Meniscus* and *Listening Sky* above).
- Agents' individual behaviours or constructions may be artistic works (swarm-built architecture, ant pheromone trails).
- Agents may generate (draw or sound) a collective structure that becomes the observable artistic output (*Eden*).
- A property of the ecosystem as a whole (e.g., its population dynamics) may constitute or generate a dynamic artistic output (*Plague*).

Combinations and alternatives to these ideas are also potentially interesting avenues for artistic expression. The character of the virtual ecosystem that the author has found most appealing is also fourfold:

- It demonstrates complex dynamics over fine and coarse timescales.
- It may explore large search spaces independently of human input.
- It has the potential for user-events to influence its behaviour.
- It has the potential to allow artist-laid constraints on the search spaces.

These traits are discussed in detail elsewhere [10] and the third, in particular, is discussed in [13]. These are by no means intended to represent universally appreciated characteristics of generative art, merely aspects of computer-based works of interest to the author. They are also traits he has observed in the works of others. It should come as no surprise then that an artist in search of computationally generated creativity described by these traits would settle on ecosystems as an avenue worthy of exploration. Ecosystems are, after all, the dominant source of the Earth's (maybe even the Universe's) novelty and complexity. It is this complexity that has given rise to us and also our concept of the sublime! Since the evolutionary process can be simulated in isolation and turned to aesthetic purposes, it seems only a small step to attempt to turn the processes of Ecology to creative endeavour.

14.8 Conclusion

Although the limitations of our abilities to code multiple-level hierarchies are apparent, clearly this does not imply that our art is similarly constrained. The element of the sublime in a Caspar David Friedrich canvas does not arise from the intricacy of its mechanism, but from contemplating nature and our place in it from behind the safety of a picture plane. Even more apparent is the irrelevance of intricacy and nature (taken literally) to postmodern interpretations of the sublime such as those discussed by Jean-Francois Lyotard [34]. Hence, works such as the dark canvasses produced in Mark Rothko's later years may be discussed in terms of their contribution to the *postmodern sublime*. The sublime does not lie *in* a work, rather the work may act to trigger a sublime experience in a viewer. In the case of Lyotard's ideas, this relates to a sense of formlessness and therefore of things that may be better left unpresented.

Since our machines are faster at mathematics than we are, they will always maintain the ability to play the role of mediator between us and the vast computational spaces outside our direct experience. Perhaps one day these same machines will be able to take us beyond spaces that look more and more familiar as we travel through them, into spaces that increase in complexity and continue to surprise us. Here the sublime experience of nature's vastness and ferocity may be rivaled through a sense of the computational sublime. We will be sensing a space rendered maybe with points on a plane and computed on-the-fly by our fastest machines, and it will seem to us as terrible and delightful as standing on an icy summit surveying all the world.

References

1. Adami C, Brown CT (1994) Evolutionary Learning in the 2D Artificial Life System Avida. In: Artificial Life IV. Brooks RA and Maes P (eds.), MIT Press, pp 377–381.
2. Bedau MA, McCaskill JS, Packard NH, Rasmussen S, Adami C, Green DG, Harvey I, Ikegami T, Kaneko K, Ray TS (2000) Open Problems in Artificial Life. Artificial Life 6:363–376.
3. Bern M (1999) Art Shows at PARC. In: Art and Innovation, Harris C (ed). MIT Press, pp 259–277.
4. Berry R, Rungsarityotin W, Dorin A, Dahlstedt P, Haw C (2001) Unfinished Symphonies – Songs of 3.5 worlds. In: Workshop on Artificial Life Models for Musical Applications, Sixth European Conference on Artificial Life, Bilotta E, et al. (eds.). Editoriale Bios, pp 51–64.
5. Dennett DC (1991) Real Patterns. Journal of Philosophy 88:27–51.
6. Dorin A (2000) Creating a Physically-Based, Virtual-Metabolism with Solid Cellular Automata. In: Proceedings Artificial Life 7, Bedau M, et al. (eds.). MIT Press, pp 13–20.
7. Dorin A (2001) Aesthetic Fitness and Artificial Evolution for the Selection of Imagery from the Mythical Infinite Library. In: Advances In Artificial Life, 6th European Conference on Artificial Life, Kelemen J and Sosik P (eds.). Springer-Verlag, pp 10–14.
8. Dorin A, McCormack J (2002) Self-Assembling Dynamical Hierarchies. In: Proceedings of Artificial Life 8, Standish R, et al. (eds.). MIT Press, pp 423–428.
9. Dorin A (2003) Meniscus. In: Experimenta, House of Tomorrow Catalogue, Taylor A (ed.). Experimenta Media Arts, Australia, p 32.
10. Dorin A (2004) The Virtual Ecosystem as Generative Electronic Art. In: Proceedings of 2nd European Workshop on Evolutionary Music and Art, Applications of Evolutionary Computing: EvoWorkshops 2004, Coimbra, Portugal, April 5–7, Günther RR, et al. (eds.). Springer-Verlag, pp 467–470.
11. Dorin A (2006) The Sonic Artificial Ecosystem. In: Proceedings of the Australasian Computer Music Conference (ACMC 2006), Haines C (ed.). ACMA, Adelaide, Australia, pp 32–37.
12. Driessens E, Verstappen M (2001) Keynote presentation in Proceedings of Second Iteration, Second International Conference on Generative Systems in the Electronic Arts, Dorin A (ed.). CEMA, Melbourne, Australia, pp 12–13.
13. Eldridge A, Dorin A, McCormack J (2008) Manipulating Artificial Ecosystems. In: 6th European Workshop on Evolutionary and Biologically Inspired Music, Sound, Art and Design: EvoWorkshops 2008, Napoli, Italy, Giacobini M, et al. (eds.). Springer-Verlag, pp 392–401.
14. Epstein JM, Axtell R (1996) Growing Artificial Societies, Social Science from the Bottom Up. Brookings Institution, MIT Press.
15. Feagin SL (1995) Sublime. In: The Cambridge Dictionary of Philosophy, Audi R (ed.). Cambridge University Press, p 774.
16. Forrest S, Jones T (1994) Modelling Adaptive Systems with Echo In: Complex Systems: Mechanisms of Adaptation. IOS Press, pp 3–21.
17. Gardner M (1970) The Fantastic Combinations of John Conway's New Solitaire Game, "Life". Scientific American 223:120–123.
18. Gross D, McMullin B (2001) Is It the Right Ansatz? Artificial Life 7:355–365.

19. Hipple WJ (1957) The Beautiful, the Sublime, and the Picturesque in Eighteenth-Century British Aesthetic Theory. Southern Illinois University Press.
20. Hoffman DD (2000) Visual Intelligence: How We Create What We See. Norton.
21. Kandinsky W (1979) Point and Line to Plane. Dover Publications, (first published 1926).
22. Kauffman SA (1993) The Origins of Order. Oxford University Press.
23. Longinus (1963) On the Sublime. Translated by Havell. In: Aristotle's Poetics, Demetritus on Style, Longinus on the Sublime. Everyman's Library, Dutton, pp 133–202 (first written ca. 250 AD).
24. McCormack J (2001) Eden: An Evolutionary Sonic Ecosystem. In: Advances in Artificial Life, 6th European Conference, Kelemen J and Sosik P (eds.). Springer, pp 133–142.
25. McCormack J, Dorin A (2001) Art, Emergence, and the Computational Sublime. In: Proceedings of Second Iteration, Second International Conference on Generative Systems in the Electronic Arts, Dorin A (ed.), CEMA, Melbourne, Australia, pp 67–81.
26. Pargellis AN (2001) Digital Life Behaviour in the Amoeba World. Artificial Life 7:63–75.
27. Rasmussen S, Baas NA, Mayer B, Nilsson M, Olesen MW (2001) Ansatz for Dynamical Hierarchies, Artificial Life 7:329–353.
28. Ray TS (1991) An Approach to the Synthesis of Life. In: Artificial Life II, Langton C, et al. (eds.). Addison-Wesley, pp 371–408.
29. Rower ASC (1998) Calder Sculpture, National Gallery of Art, Washington, Universe.
30. Shelley M (1989) Frankenstein, the Modern Prometheus. Joseph MK (ed.). Oxford Univeristy Press (first published 1818).
31. Sims K (1994) Evolving Virtual Creatures. In: Proceedings of SIGRRAPH 1994, ACM Press, pp 15–34.
32. Taylor T (2002) Creativity in Evolution: Individuals, Interactions and Environments. In: Creative Evolutionary Systems, Bentley PJ and Corne DW (eds.). Academic Press, pp 79–108.
33. Tufte ER (1990) Envisioning Information. Graphics Press.
34. van de Vall R (1995) Silent Visions, Lyotard on the Sublime. In: The Contemporary Sublime, Sensibilities of Transcendence and Shock, Hodges (ed.). Art and Design. VCH Publishers, pp 69–75.
35. Yaeger L (1994) Computational Genetics, Physiology, Metabolism, Neural Systems, Learning, Vision and Behavior or Polyworld: Life in a New Context. In: Proceedings, Artificial Life III, SFI Studies in the Sciences of Complexity, Langton C (ed.). Addison-Wesley, pp 263–298.