The *Creatures* Global Digital Ecosystem

Dave Cliff
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139
davec@ai.mit.edu

Stephen Grand
CyberLife Technology Ltd
Quayside, Bridge Street
Cambridge, CB5 8AB
U.K.
steve.grand@cyberlife.co.uk

Keywords
digital organisms, artificial ecosystems, entertainment software

Abstract An artificial life entertainment-software product called *Creatures* was released in Europe in late 1996 and in the United States and Japan in mid-1997. When installed on a domestic computer (PC or Macintosh), each *Creatures* CD-ROM creates a virtual world in which autonomous software agents exist. The agents, known as "norns," interact with the human user, with each other, and with objects in their virtual world. Each norn coordinates perception and action via its own modular recurrent neural network: Each network has Hebbian learning, plus diffuse modulation of activity via a "hormonal" system that is part of that norn's "biochemistry." Details of each norn's neural network and biochemistry are genetically specified, and norns can breed via sexual reproduction. In the reproduction process, genetic material may be mutated and may also be subjected to "gene duplications" that enable potentially unlimited increases in complexity of the norns' design. Over 500,000 *Creatures* CD-ROMS have now been sold. As each installed copy of *Creatures* can support 5 to 10 simultaneously existing individual norns, it seems reasonable to estimate that there are up to 5 million norns existing in the "cyberspace" provided by the global *Creatures* user community. Continued growth of the global norn population, to figures measured in tens of millions, is quite likely. Although a commercial product, the *Creatures* digital ecosystem should be of interest to artificial life scientists: There are obvious parallels with Yaeger's PolyWorld and Ray's NetTierra systems. This article provides a detailed discussion of the links between the artificial life literature and the technology used in *Creatures* and includes anecdotal discussion of the "digital naturalism" witnessed on the many independent websites maintained by *Creatures* enthusiasts.

1 Introduction

In November 1996 an artificial life entertainment-software product (i.e., a computer game) called *Creatures* was released in Europe, to immediate success. Over 100,000 units were sold in the first month, and the product attracted significant media attention. Release in the United States and Japan followed in mid-1997, and global sales totaled over 500,000 units by the second quarter of 1998. A revised and extended version, *Creatures 2*, was released globally in September 1998 with an immediate shipment of 200,000 units.[1]

The merits of *Creatures* as entertainment are not discussed here. Rather, this article is concerned with issues surrounding the underlying technology on which *Creatures* is

1 For details of the differences between *Creatures* and *Creatures 2*, see [44].

based, and with the "digital ecosystem" created by the global *Creatures* user community.

The underlying technology used in *Creatures* draws on and combines ideas developed in artificial life (alife) research. The technology is referred to as *CyberLife*.[2] Although *Creatures* was the first practical demonstration of CyberLife technology, there is nothing within the technology that limits it to entertainment applications. CyberLife has been developed by CTL (the developers and owners of the technology) as a route to the generation of artificial intelligence: It could reasonably be employed in a wide variety of simulation domains, and to add extra intelligence or autonomy to existing products.

CyberLife allows for the creation of autonomous agents based on three core techniques: artificial neural networks, artificial biochemistries, and artificial evolution. None of these techniques are new or proprietary: All have been studied and developed in alife research for over a decade. However, the way in which these three techniques are integrated in CyberLife is novel and represents a significant investment in engineering development. The method of integration stems from a particular set of methodological principles employed by CTL (see [12]). Detailed technical descriptions of how CyberLife is used in *Creatures* have been published elsewhere [19, 20]. The information in those papers is sufficient to motivate an extended discussion, drawing explicit links and comparisons with research published in the alife literature and discussing the implications of the success of the product, both for the science and for the sociology of alife. And that is what we present here.

We start with an overview of *Creatures*, followed by brief summaries of other entertainment software products that are comparable to *Creatures*. Next, we discuss the relationship between *Creatures* and work in alife: first in terms of the component technologies, then with reference to other alife projects in which complete "digital ecosystems" have been constructed. In particular, we focus on Ray's [40, 41] NetTierra, an ongoing project to develop a digital biodiversity "game reserve" by using spare computing capacity on the Internet. One of Ray's arguments for developing NetTierra is that, once large-scale self-sustaining high-diversity digital ecosystems are established and stabilized, it should be possible for people to engage in "digital naturalism," where individuals observe and experiment with the digital life forms that evolve within the digital ecosystem. We close this article with anecdotal discussion demonstrating that such digital naturalism is already being practiced by members of the global *Creatures* user community.

2 *Creatures*

Creatures involves a virtual environment, populated by two types of simulated creatures: "norns" and "grendels." The user is responsible for the maintenance and breeding of norns. The *Creatures* packaging actually contains two disks: the CD-ROM containing all the program files, and an "egg disk" containing six "eggs." Each "egg" contains genetic material for an individual norn: Eggs may be taken from the egg disk and placed in an incubator to hatch into a juvenile norn. This is how each user creates the first generation of norns. In principle each norn hatched from the egg disk has a unique genetic makeup.

A user's norn interacts constantly with its environment and intermittently with its human user. The environment is a "$2\frac{1}{2}$-dimensional" world: essentially a number of horizontal platforms at different heights (i.e., a 2D space), with a depth ordering established when more than one object occupies the same position in the two-dimensional

2 CyberLife is a trademark of CyberLife Technology Ltd. Throughout this article, CyberLife Technology Ltd is referred to as CTL.

space. In addition to the fixed spatial layout, the environment includes multiple creatures, plant life (some of which is edible), and a large number of both fixed and movable objects including toys and transportation devices. There is one environment in each installation of *Creatures*, and all copies carry the same initial environment conditions. Thus, if a user has more than one norn active at a time, they must inhabit the same environment, but the spatial extent of the environment is sufficient to avoid forcing the norns into continual interaction with each other.

The environment also supports a small number of grendels. These are norn-like creatures that the user has little influence over and that are hostile toward norns. Grendels, although not rapacious, are likely to attack a norn if the norn comes too close. They are also a potent source of "bacterial diseases."

Norns can learn on the basis of experience: both from interacting with the environment, and from direct "rewards" and "punishments" administered by the user. Norns have a primitive linguistic capability: The user can teach a norn nouns and verbs, and norns can communicate with and learn from each other. Norns have a finite lifetime and move through distinct developmental stages of "childhood," "adolescence," "maturity," and "senesence" (old age).

The norn's behavior is generated by its "brain," an artificial neural network that coordinates the norn's perceptions and actions, according to a set of behavioral "drives and needs." The network is heterogeneous (the neurons are not all identical in their response) and has intrinsic dynamics at multiple time scales (i.e., the responses of the individual neurons may change over time, and connections between neurons may also change over time: Both types of changes can affect the overall behavior of the norn). Technically, the network is a continuous-time recurrent neural network (CTRNN) similar to that studied by [6].

For practical reasons, the norn neural network is not *directly* involved in either perceiving the environment or generating actions at the level of altering body joint angles. It would require too much computation to simulate accurately the physics of sound, vision, touch, smell, and taste, along with the neural processing that would be required to recognize particular sensory patterns as objects in the environment (e.g., to perceive a particular pattern of pixels on the norn's "retina" as a particular object). Instead, objects in the world are perceived "directly": Individual objects within a certain range of the norn directly affect the activities of individual neurons in the norn's network. The range can be dependent on the norn's position and orientation: They can only see in the direction they are looking, and sounds attenuate with distance and are muffled when they pass through solid objects such as walls. Similarly, it would require too much processing if a norn's goal of moving in a particular direction was carried out by the neural network with individual neurons controlling individual "muscles" in the norn's "body." So there are a fixed number of predefined action scripts (e.g., "move left," "push object"), written in a higher-level language. In the output stage of the network, there is one neuron for each script, and the output neuron with the highest activity is allowed to initiate the execution of its action script, which then generates the required pattern of synchronized body movements. This approach is entirely justifiable, given the processing capabilities of the target platforms, and there are many precedents for this approach in the alife research literature.

It is important to note that the *Creatures* environment acts as a virtual reality *for the norns themselves*: From the perspective of an individual norn, it is interacting with a "real" world, but a real world rather different from the real world we humans are accustomed to. In this sense, the norns are *situated autonomous agents*: They are capable of coordinating perceptions and actions over extended periods of time without human intervention. Interactions with the human user may alter the norns' behavior, but the human can only *influence* a norn, not control it.

The operation of a norn's neural network is affected by that norn's "biochemistry." Events or actions such as eating can release reactive "chemicals" into the norn's "bloodstream," where chain reactions may occur (releasing new chemicals); some chemicals in the "bloodstream" may bind to particular neurons in the network, altering their performance. This gives the possibility of the norns eating "toxic" foodstuffs, which would require the administration of care (and possibly feeding the norn "medicinal" food) to revive the norn. More significantly, it allows for modeling changes in motivational state such as those that occur when a human releases hormones or ingests artificial stimulants such as caffeine or amphetamines.

The environment includes a very simple model of infectious agents or "bacteria" that may cause potentially harmful changes to the norn's metabolism. It may be necessary to administer chemicals to an infected norn, and some of the "plant life" in the environment may, when eaten, counteract the effects of the bacteria. The details of how a bacterium affects a norn, and what chemicals or plants affect it, are genetically encoded for each strain of bacterium, so there is an opportunity for coevolutionary interactions between the bacteria and their hosts.

The structure of a norn's neural network, the details of its biochemistry, and also aspects of its body morphology (i.e., its "physical" build and appearance) are all specified by that norn's "genes," encoded as strings of characters. Sequences of characters on the genome of a norn are read and interpreted to determine parameters such as the number of neurons in a module in the norn's network, the manner of their interconnection, the number and nature of the possible chemical interactions in the biochemistry, and the types of interaction between the biochemistry and the brain. The timing and effects of the different stages of development are also encoded on the genes, including "senesence" genes that affect the age at which the norn dies. At different stages of development, new processes may become active: A norn is not sexually mature until it has passed through adolescence, when new hormones are released and new behaviors exhibited.

The norns are allowed to sexually reproduce (i.e., using crossover between two parent genomes, with random mutations): This may result in offspring norns with altered networks, metabolisms, morphologies, and development patterns, any or all of which may affect the norn's observable behavior. Furthermore, the *number* of genes can vary: Two "parents" can give rise to a "child" that has more (or fewer) genes than its parents. Extra genes may encode for bigger brains, more complex biochemistries, or both. The net effect is that new behavior patterns can evolve over a number of generations. Because the norns can "learn" from one another, there is also the possibility of cultural transmission and accumulation of information: If an old norn teaches something to a young norn, then when the old norn dies its experience "lives on" in the young norn.

To summarize, the key features of CyberLife as used in *Creatures* are as follows: one or more simulated autonomous agents within a rich virtual environment, where the agents have sensory (sound, vision, etc.) and motor (movement, speaking, etc.) capabilities, coordinated by an artificial neural network that can learn from experience and be affected by the agent's biochemistry, with details of both the network and the biochemistry being specified by the agent's genetic composition. The use of CyberLife in *Creatures* demonstrates that the technology can be encoded in a computationally efficient manner: *Creatures* runs on domestic computing platforms (e.g., Windows-PC compatible and Macintosh computers), allowing for up to around 10 norns to be simulated in real time. This may involve processing around 20,000 neurons with 100,000 interconnections every second, in addition to the processor load imposed by the rest of the system, such as the simulated physics of perception and action, real-time graphic displays, and so forth [20].

3 Other Alife Entertainment Products

Maes [28] presented an overview of a number of academic projects in entertainment software developed using ideas from alife research. Publications in the scientific literature describing commercial interactive entertainment software products are rare, so many of the citations in this section are to promotional material available on the World Wide Web.[3]

The first commercial entertainment software that bears any comparison with *Creatures* predates the first alife conference. In 1985, Activision marketed a product called *Little Computer People* (*House on a Disk*), written by David Crane for the Commodore64 computer and subsequently ported to the Commodore Amiga, Apple II, and Atari ST platforms. The screen display for this program resembled a doll's house with the front wall removed: The interior rooms of a suburban dwelling were shown, with items of furniture including a bed, TV, a music center, home computer, and seats. Inside the house, an animated man and his dog went about routine daily activities such as watching the TV, sleeping, listening to music, or playing on the computer. An elementary natural-language front end allowed the user to communicate with the man, and the man could type letters back to the user. The man (and the dog) were also reliant on the user sending food and drink, new music, and such. If left unfed, the inhabitants of the house would eventually die. If left unattended for too long, the man could become bored, sad, angry, or ill. Activision took the innovative step of duplicating all copies of the program in-house, encoding in each copy a unique machine-readable serial number. When the program was run, it would initialize by using the serial number to determine aspects of the man's behavioral characteristics (i.e., his "personality"). Thus, while there was no adaptation (learning or evolution) in this product, there was a degree of diversity that was novel for its time. Although a commercial success, and now regarded as something of a cult classic, *Little Computer People* failed to establish a "genre" of games; more than 10 years passed before a comparable product was released.

One of the first pieces of entertainment software explicitly promoted as drawing on alife research was *SimLife*, released by Maxis in 1993 [30]. In essence, *SimLife* allows a user to observe and interact with a "simulated ecosystem" with a variable terrain and climate, and a variety of species of plant life, herbivorous animals, and carnivorous animals. The product package includes a "laboratory notebook" for recording experiments. *SimLife* was designed and programmed by Ken Karakotsios, who subsequently coauthored a strategy guide that gives detailed descriptions of the system [23].

SimLife provided a simulated "landscape" with geographic features and resources such as water or nutrients. Cellular automata techniques were used to generate the terrain and to distribute the resources for each landscape. The landscape was then "populated" with simple artificial life forms, referred to as "orgots" in the *SimLife* literature. The orgots were independent mobile objects that could move around the landscape and interact with one another; their behavior depended both on external environmental stimuli and also on their internal state; their internal state included a simple model of energy metabolism: For instance, orgots could be plants or animals, and animal-type orgots could be herbivores, carnivores, or omnivores. Aspects of each orgot's life cycle and behavior were controlled by its "genotype": genotypes specified phenotypic parameters such as speed of movement, permissible food sources, reproductive characteristics, and so on. When orgots met the criteria for reproduction (which, in sexually reproducing species, included being within a maximum spatial distance of another orgot), they produced an offspring that inherited genetic material from the par-

3 The text in this section is extended and updated from the brief survey we presented in [20].

ent(s) and was subject to random variation introduced via mutation—the mutation rate could be spatially dependent: In addition to a general "background" mutation rate, the landscapes could include mutagen sources, the intensity of effect of which followed a reciprocal function of the distance from the source to the reproducing orgots. Thus, in *SimLife* there was no explicit a priori fitness function used to evaluate each individual before reproduction to give a differential selection pressure, as is common in standard genetic algorithm applications (an approach Packard [36] described as "extrinsic adaptation"). Rather, *SimLife*, like *Creatures*, employed what Packard [36] described as "intrinsic adaptation."

In our opinion, the most significant distinction between *SimLife* and *Creatures* is that *SimLife* operated primarily at the population level whereas *Creatures* is individual-centered: In *SimLife*, the user observed and interacted with sizeable populations of comparatively very simple digital organisms, whereas *Creatures* users observe and interact with small numbers of comparatively very complex digital organisms. Although much of the development work on *Creatures* was done before the phrase entered common usage, the *Creatures* norns can reasonably be categorized as "virtual pets,"[4] unlike *SimLife*'s orgots.

Over the past 5 years, decreases in the real costs of processor power, RAM, and disk storage have been little short of astounding. For this reason, it is important to remember that although *SimLife* was released only 6 years ago, it was targeted at what would today be considered as impoverished domestic computing platforms. Thus, while it served as a valuable introduction to alife for many members of the general public, we believe that very few observers of the autonomous agents available in today's marketplace would recognize *SimLife*'s orgots as virtual life forms. Nevertheless, *SimLife* deserves to be remembered as an important step forward on the path initially broken by Crane's *Little Computer People*.

More recent products have had stronger links to autonomous agent research directed at creating virtual pets. Another Maxis product, *El-Fish*, was presented as an "electronic aquarium" where users could design and breed virtual tropical fish that could then be observed swimming in a virtual fish tank. The similarities between this product and the work of Terzopoulos, Tu, and Grzeszczuk [48] are manifest.

It should also be noted that Maxis pioneered the concept of *software toys* as opposed to *computer games*. The metaphor of *toy* rather than *game* is intended to highlight a different style of interaction: A game is usually played in one (extended) session, until an "end condition" or "goal state" is reached, and a score or high score is awarded; in contrast, use of a toy does not imply a score or an aim to achieve some end condition, and interaction with a toy is a more creative, ongoing, open-ended experience.

Subsequently, PFMagic Inc. [37] released two products, *Dogz* and *Catz*, which gave users on-screen animations of virtual pets based on dogs and cats. Users can interact with their virtual pets and train them to perform simple tricks. There is some superficial similarity between these products and Blumberg's work [8, 9], but there is little indication that the underlying technology employs ideas developed in alife.

Two products released since *Creatures* are *Fin-Fin* by Fujitsu Interactive Inc. [18] and *Galapagos* by Anark [1]. Both are marketed as alife software. Of the two, *Galapagos* has stronger genuine links to alife, involving a three-dimensional kinematically realistic model of a six-legged agent in a three-dimensional maze-like environment with an adaptive neural-network-like controller based on Anark's proprietary "NERM" (nonstationary entropic reduction mapping) technology. *Fin-Fin* involves three-dimensional rendering of a hybrid dolphin/bird creature with which the user can engage in sim-

4 A comprehensive list of web links to sites describing "virtual pet" products is available from http://www.virtualpet.com/vp/.

ple interactions, via a specialized input device combining a proximity detector with a microphone that detects amplitude and pitch of nearby sounds (e.g., voices).[5]

Finally, there has recently been a consumer craze for *Tamagotchi*, cheap battery-powered hand-held dedicated processors that run a "virtual pet" program, requiring periodic attention in some form. The user interfaces of Tamagotchi devices typically consist of a custom LCD display and a small number of buttons. Given the simplicity of the environment in which Tamagotchi "pets" exist, the restrictions of the user interface, and the simplicity of the behavioral repertoire of these "pets," we find it difficult to draw useful comparisons between these toys and any of the products or technologies mentioned elsewhere in this article. Nevertheless, the sociology, psychology, or population memetics of the spread of the craze may be of some academic interest.

Dogz, *Catz*, *Fin-Fin* and *Galapagos* are all presented as involving alife technologies, but none of them (yet) employ genetically encoded neural network architectures or artificial biochemistries as used in *Creatures*. Nor do they allow for the development of "culture" in communities of artificial agents, or "digital naturalism" in communities of users, both of which are possibilities with *Creatures* that we discuss later in this article. Before that, we draw explicit links between the alife research literature and the CyberLife technology used in *Creatures*.

4 Related Work in Artificial Life

Having now introduced the core components of CyberLife technology as used in *Creatures*, it is possible to make detailed comparisons between *Creatures* and other work in the alife research literature. We first examine details of the components of CyberLife in relation to work in artificial life, prior to discussion of the relationship between *Creatures* and complete integrated "ecosystems" such as PolyWorld and NetTierra.

4.1 Component Technologies

The most novel development in CyberLife is the interaction between the neural network and the biochemistry. In comparison to other issues with alife, there is comparatively little work in the literature dealing with computational models of biochemical systems, and most of that literature deals with issues in the origins of complex organic compounds and self-sustaining cyclic chain reactions or autocatalytic sets (see, e.g., [2–4, 17, 24, 32]). None of this work addresses the interaction of biochemical activity with the operation of a neural network.

Although the effects of psychoactive hormones and drugs have long been studied in neurobiology, the assumption was long held that all signaling from one neuron to another was via *direct* connections at the synapses (the point where an output terminal or *axon* of a signaling neuron connects with the membrane of a receiving neuron). Most such connections are chemical: Electrical activity on the membrane of the signaling neuron initiates the release of small packets of *neurotransmitter* chemical from the axon onto the membrane of the receiving neuron, a process that alters the level of electrical activity on the receiving neuron's membrane. Less common are electrical connections, where electrical activity on the signaling neuron's membrane directly affects that of the receiving neuron, without intervening neurotransmitters being released. However, recent developments in neuroscience have identified the presence

5 In June 1997 Fujitsu Interactive released another virtual pet product, *K-9 Cyber Companion*, to lukewarm reviews (e.g., http://www.gamesdomain.co.uk/gdreview/zones/reviews/pc/jul97/k9.html); and currently only very vague details of this product appear on the Fujitsu Interactive website at http://www.fujitsu-interactive.com. Fujitsu Interactive has also announced (at http://fujitsu-interactive.com/press/BioPunxPress1.html) a forthcoming product, *BioPunx*, which will feature "autonomous A-life characters . . . [that] . . . go way beyond AI [artificial intelligence] and bring you true Artificial Life." But, at the time of writing this article, no technical details were available and the product's web page at http://www.biopunx.com was empty.

of *gaseous* neurotransmitters: nitric oxide in particular. This discovery indicates that neurons may be capable of signaling in a *diffuse* manner, by release of gases to nearby neurons. For further details, see, for example, [14, 15].

We know of no papers in either the alife literature or the artificial neural network literature that discuss this issue, and indeed the discovery is so recent that the full implications are yet to be explored by neuroscientists. Yet the biochemistry in CyberLife is implemented in such a way that it offers the opportunity to model diffuse gaseous signaling: In addition to providing direct stimulation to one or more neurons, a CyberLife neuron can emit chemicals that diffusely affect the operation of other neurons. The chemicals a neuron emits, and those to which it is receptive, can both be genetically specified.

Furthermore, there is nothing in the biochemistry that prevents the formation of autocatalytic sets, so in principle there is the possibility for complex self-sustaining biochemical pathways to interact in subtle and hopefully useful ways with the neural network. It could be possible for CyberLife agents to exhibit “moods” and “emotions” as a result of interactions between the nervous system and the biochemical system. The need for such phenomena in artificial agents has long been argued for by Sloman (see, e.g., [46]).

Independent of the biochemical system, CyberLife neurons are units with rich intrinsic dynamics, which can be connected to form CTRNNs exhibiting complex dynamical behavior. The CyberLife networks are divided into a number of distinct “lobes”: There are connections both within and between lobes, in a manner reminiscent of the computational neuroscience work of Edelman [13]; this contrasts with the majority of research in artificial neural networks, where no such global structure is imposed. For a recent collection on modular approaches to artificial neural networks, see [43]. Empirically, the dynamics of the networks in *Creatures* are stable while avoiding pathological point or limit-cycle attractors. Analytic explanations of the dynamics of such CTRNNs have only recently been developed, and then only for small networks (fewer than 10 units, say): See, e.g., [6].

The norn’s neural network and the biochemistry are both encoded in such a way that genotype length can vary, and deletion and duplication operators are used in breeding to introduce length variations. The genetic encoding scheme allows for variations in length by employing “marker” sequences on the genotypes to act as punctuation, dividing the genotype into distinct zones. Marker-sequence punctuation is an idea borrowed from natural genetic systems that has also been explored by other alife researchers (e.g., [35, 47]). Thus, CyberLife is similar to other approaches in artificial evolution where genotypes can increase in length allowing for increases in complexity of the genetically encoded objects, such as Harvey’s SAGA [21], Koza’s genetic programming [25] or Sims’ evolved autonomous agents [45], and thus CyberLife has the same potential for incremental evolution of ever more complex designs.

A number of factors make it likely that rapid and productive evolution will occur in CyberLife systems. The provision of a Hebbian reinforcement learning mechanism, with genetically specified parameters, allows for the Baldwin effect (see, e.g., [22, 31]) to operate, speeding the process of searching for better genotypes. CyberLife systems could readily be used in coevolutionary scenarios, thereby harnessing the Red Queen effect [10, 42, 51]. Additionally, the provision of sex-linked genetically specified characteristics makes it possible that sexual selection can be employed to encourage diversity in the search process, in the manner demonstrated by Todd and Miller [49]. The use of distinct development stages (or *ontogeny*) in determining (possibly via a biochemistry) what behaviors artificial neural networks exhibit is a topic that has received little attention in the literature (one example is the work of Nolfi and Parisi [34]).

Finally, although a number of researchers have reported on studies of the evolution of communication in animats (e.g., [5, 26, 27, 33, 52]), none of the studies conducted so far have involved agents as complex as the norns in *Creatures*, and so the interaction between genetic evolution, lifetime learning, and cultural transmission of information (i.e., “population memetics”: See, e.g., [7]) remains a topic open for further research.

Furthermore, although it is highly unlikely in the current version of *Creatures*, it is tempting to speculate about the possibility of social structures emerging in the *Creatures* environment. Given that the norns can communicate with one another, and that supplies of some environmental resources (such as food or “medicine”) can sometimes be limited or scarce, it is not inconceivable that simple economic interactions such as bartering, bargaining, and trade occur between norns, allowing for comparison with recent work in simulated societies such as that by Epstein and Axtell [16].

4.2 Integrated Systems

There are very few publications in the research literature that describe systems such as *Creatures*, where complete integrated systems have been constructed. Nevertheless, there are a few projects where worthwhile comparison can be made.

Ray’s Tierra [38] allows for large numbers of simple digital organisms to interact, compete, and evolve in a virtual environment. Ray has argued forcefully that such software systems are not necessarily models or simulations of life on earth; rather, they may be independent instances of life [39]. The Tierra system provided an elegantly minimal virtual environment in which primitive digital organisms could replicate, compete for limited resources, and (potentially) evolve. In a landmark paper, Ray [38] described how, by seeding the Tierra environment with a simple self-replicating organism, a rich evolutionary process with complex dynamics was unleashed, giving rise to diverse groups (or “species”) of agents, including species that were parasitic.

However, the agents in *Creatures* are vastly more complex than those in Tierra. In colloquial terms, if the agents in *Creatures* are similar to animals in their complexity of design and behaviors, then the agents in Tierra are similar to bacteria or viruses. This is not intended as criticism of Ray’s work: The intention in Tierra was to create systems of minimal complexity to allow for greater ease of analysis, to reduce computational load, and to give clearer indication of the necessary and sufficient design criteria for agents exhibiting the desired evolutionary phenomena. It would not have been practicable to use software agents as complex as the norns in *Creatures* to study the issues explored by Ray.

Prior to Ray’s work on Tierra, Apple Computer, the MIT Media Laboratory, and the Los Angeles Open School collaborated on the *Vivarium Project* [11, 50]. The aim of the Vivarium was to develop an interactive virtual environment populated by software agents that could be used to develop computational models for ethology, and as an educational aid. In Coderre’s work [11], the “pets” in the vivarium could not reproduce or learn, but Travers reported on the addition of simple feed-forward neural networks that could be either evolved or edited via a graphical user interface, and on the use of the Vivarium to simulate the behaviors of real animals (three-spined stickleback fish). Subsequently, Yaeger [53], who was also involved with the Vivarium Project, reported on the development of *PolyWorld*,[6] a complex and sophisticated system that is probably the closest to *Creatures*. Yaeger’s motivation for creating PolyWorld is best summarized in his own words [53]:

> PolyWorld attempts to bring together all the principle components of real living systems into a single artificial (manmade) living system. PolyWorld brings

6 Yaeger’s PolyWorld website is at http://pobox.com/larryy/PolyWorld.html.

> together biologically motivated genetics, simple simulated physiologies and metabolisms, Hebbian learning in arbitrary neural network architectures, a visual perceptive mechanism, and a suite of primitive behaviors in artificial organisms grounded in an ecology just complex enough to foster speciation and interspecies competition. Predation, sexual reproduction, and even communication are all supported in a straightforward fashion. The resulting survival strategies, both individual and group, are purely emergent, as are the functionalities embodied in their neural network "brains." Complex behaviors resulting from the simulated neural activity are unpredictable, and change as natural selection acts over multiple generations ([53], p. 264).

Despite the clear similarities between PolyWorld and *Creatures*, there are several significant differences. PolyWorld is simpler in a number of respects: There are only 18 parameters encoded on the genotype ([53], p. 270), and the environment is a flat ground-plane, possibly with some barriers inhibiting movement ([53], p. 280). More importantly, the model metabolism in PolyWorld only affects the intake and expenditure of energy: There is no interaction between the metabolism and the neural network other than the fact that the network can have an input unit sensitive to the animat's energy level ([53], p. 273–274).

There is also a difference between the aims of CTL in developing CyberLife and the aims of Yaeger in developing PolyWorld: PolyWorld is presented primarily as a tool for scientific enquiry, as a more complex and sophisticated system for exploring the issues addressed by Ray's Tierra, in addition to providing a test-bed for developing theories or models in evolutionary biology, behavioral ecology, ethology, or neurobiology. There is no indication in Yaeger's paper that he intends PolyWorld to be used in industrial engineering or entertainment applications, and so it should not be judged by the needs of industry. In this sense then, there is a clear difference in the intended uses of *Creatures* and PolyWorld. Nevertheless, there is no reason why CyberLife systems such as *Creatures* could not be used for scientific purposes.

Since 1994, Ray has been working on the development of an ambitious but profound extension to Tierra, known as NetTierra [40, 41]. Whereas the original Tierra was a self-contained process running on a single processor, NetTierra is intended to run as a massively parallel process on the entire Internet (or, more realistically, the largest subnetwork of the Internet that NetTierra is granted access to). NetTierra programs will run as background processes, consuming spare processor capacity, on as many computers as possible. From time to time, digital organisms may autonomously migrate from one machine to another via Internet connections. In the NetTierra "ecosystem," the primary scarce resource that organisms compete for is spare processor time, and so Ray predicts a flow of digital organisms to areas of the Internet where there are high concentrations of idle machines: typically on the dark side of the planet, where the majority of users are asleep. Ray argues that the global NetTierra system will present an opportunity for creating a "digital ecosystem" capable of supporting self-sustaining evolutionary processes with a high degree of diversity.

Ray [40] hypothesizes that it is possible that, given enough time, innovative new software, with significant potential for industrial application, may evolve. Ray notes that it will be necessary for "digital naturalists" to observe and experiment with the evolving digital organisms, possibly removing promising-looking "wild" organisms for isolation to allow "domestication" and subsequent "farming." Ray supports his argument with analogies to the natural global ecosystem, noting that "useful" life forms such as corn, cattle, and horses may have been inconceivable before they arose from self-sustaining evolutionary processes. Thus NetTierra is not only of scientific interest: Ray argues that it also has the potential for (eventually) spawning useful new software technologies.

This brings us back to the use of CyberLife in applications other than *Creatures*. In systems where predictive simulations of the activity of groups of autonomous agents (from ants to humans) are required, it may be possible to employ CyberLife to create *Creatures*-like integrated systems where the required behaviors emerge or evolve, but there is a significant problem in that such systems (including Tierra, NetTierra, and PolyWorld) exhibit what Packard [36] refers to as *intrinsic adaptation*. That is, in these systems, there is *no fitness function except for survival* [53, p. 269]. This is no problem if the system is intended as a scientific model, because real evolutionary systems also have no explicit fitness function. But for industrial applications, great care (and possibly much trial and error) will be required to ensure that the intrinsic adaptation gives rise to desired behaviors.

For this reason, it may make more sense to employ components of CyberLife in less wholly integrated systems, so that explicit learning error functions or fitness evaluation functions can be employed to steer the search of parameter space(s) and monitor progress of the system toward the desired goal. This may require that large amounts of training data are available, and if the aim is to simulate human activity, generating such training data may be a costly process. Even then, formulating appropriate functions can be a major problem [54], compounded by other issues when using evolutionary techniques to design autonomous agents [29]. And even when these problems are overcome, monitoring evolutionary progress can be a difficult and expensive process [10].

The text in this section has demonstrated the strong links between the existing alife research and the principles and techniques integrated in CyberLife. The way in which these techniques are integrated in CyberLife appears to be unique. Furthermore, the demonstration in *Creatures* of interactive real-time use of CyberLife on domestic computing platforms indicates that the technology is computationally efficient. Nevertheless, the commercial-scale engineering of complex adaptive systems directly inspired by biology is in its infancy, and many currently unseen problems and pitfalls could lie ahead. Even with the development of CyberLife, there is still no such thing as a free lunch.

5 Digital Naturalism with Creatures

Sales of 500,000 units indicate that *Creatures* has attracted a sizeable number of dedicated users. Confirmation of this comes from noting that a Usenet newsgroup, alt.games.creatures, has been established (independent of CTL); and there are now over 200 independent websites dedicated to *Creatures*. It seems fair to assume that, for a large majority of the users, *Creatures* has been their primary introduction to concepts in alife.

Prior to *Creatures*' release, we anticipated that the philosophical question of whether the norns are truly alive would be a major issue discussed among users and observers of the system, but this appears not to have happened. We also foresaw that groups of users would breed and exchange norns. This did happen, and at a rate much faster than we expected. Furthermore, and to our surprise, we saw the rapid appearance of users reporting the results of "hacking" genomes, producing new "genetically engineered" strains of creatures. Several of the *Creatures* websites are devoted entirely to discussions or exchanging data or software for breeding or genome hacking.[7] Without any prompting, these users appear to be engaging in exactly the kind of digital naturalism that Ray foresaw the need for in NetTierra. In the rest of this section we further illustrate this by describing some of the less-expected developments.

7 Notable examples include "Slink's Burrow Online" (http://www.netins.net/showcase/slink/); "The Norn Underground" (http://www.dreamscape.com/lummoxjr/creatures/nuclearchamber.html); "GeNorNics" (http://www.mindspring.com/~elemkay/creatures/); and "The Creatures Exchange" (http://mudhole.spodnet.uk.com/~addicted/creatures/).

5.1 Grenorns

The first genetic manipulation we heard about was the creation of a hybrid from the two original species of creature (norns and grendels). We didn't think this was possible, since grendels had been deliberately made sterile (by removing their reproductive system genes) to prevent them from overrunning the world. Some users had realized that the way the original six norn eggs on the egg disk were made to appear unique for every user was to store two genomes in each. When an egg was hatched, those genomes were crossed to make an essentially unique individual. What the "genetic engineers" did was manually insert a genome from a grendel in place of one parent. When the eggs hatched, the result was a random cross between the two species. Much newsgroup discussion has been devoted to the behavioral characteristics of this new species.

5.2 Female Grendels

Because the grendels were deliberately sterile, only male grendels originally existed. However, one of us (Grand) has seen one female grendel, developed in Portugal. It is possible that she was a throwback descended from grenorns (which may be of either sex), but genetically she was pure grendel, except for the added genes for reproductive chemistry. We still do not know how that was done.

5.3 Natural Mutations

Of the many possible naturally occurring mutations, a few have developed such prominence that they have gained proper names and can be found in norns downloadable from many websites. Their ubiquity is indicative of the power of artificial selection by users.

One such is the "Highlander Gene," which exploits the lack of a conservation principle in the *Creatures* chemistry model. This mutation increases the amount of glucose that can be obtained from stored glycogen, thus making its owner immortal. Whether immortality can be considered an evolutionary development is a moot point. Another mutation is the "Bacchus Gene," which (bizarrely) creates starch in response to sexual stimuli, instead of the normal sex hormone. A third rather morbidly popular phenotype is created due to the "Saturn Gene" mutation, which generates a shivering response in norns regardless of their body temperature. This is caused by a mislocated chemo-emitter.

5.4 Engineered Mutations

By using hex editors (now superceded by the gene editor discussed further below), many "digital genetic engineers" have created deliberate modifications to individual genes. Popular mutant strains have spread via the web. Some of the more notable are

- *American Cardinal norns:* so called because they have been given sex-linked coloration (as in the bird of the same name).
- *New instincts:* "instinct genes" have the ability to prewire certain relationships in the brain. These are used to instill initial behavioral tendencies in newborn norns, such as "when bored, don't just stand there," or to change innate behavior at important life stages, such as the instinct to find the opposite sex suddenly attractive at puberty. Several new instincts have been added by users, such as "eat (potentially medicinal) plants when in pain," or "push things when hungry."
- *Improved immunity:* the efficacy of the antibody reactions and the amounts of initial antibodies in the newborn have all been manipulated to improve disease resistance.

- *G-defense gene:* turns the fear response into an anger response, making norns more aggressive.
- *Cough medicine gene:* modifies the response to a natural cough-suppressing substance to make it suppress sneezing, too (coughs and sneezes spread diseases).

5.5 Diseases

The bacteria that coevolve in the virtual world are relatively trivial and limited in their evolutionary potential. However, this did not stop a pair of German medical students sending one of us (Grand) a complete treatise on a disease that they had identified. The paper described the symptoms, prognosis, and even treatment. The students had taken the trouble to name the disease after themselves, so Schrey-Leonard Syndrome is possibly the first researched and named virtual disease afflicting digital organisms.

5.6 Tools and Objects

Shortly after the initial release of *Creatures*, it became clear that some users had worked out the details of the script language that controls the virtual objects within the *Creatures* environment and discovered how to inject new objects into the world. This has led to a plethora of new objects, including medicines, toys, antigrendel devices, and so on. Thus, the users are not only monitoring or altering the agents but also the environment that the agents are embedded within.

5.7 Gene Editor

The level of interest in norn genetics has become so great that people have tried to write their own manipulation and observation tools. Some website owners have entirely devoted themselves to the task of understanding the norn genome and the implications of mutations and deliberate changes: A "Norn Genome Project." The limited facilities appropriate for this task within the original *Creatures* product led CTL to release an improved (but not in the least simplified) version of the original gene editor tool that was used to create the creatures during product development. This complicated piece of software is only available on the Internet and has not been advertised, but many hundreds of copies have been downloaded, and new genetic variants created by this method have been placed on the web.

5.8 Sociological Issues

In addition to digital naturalism as envisioned by Ray, the *Creatures* user community has shown the formation of social groups, often centered on somewhat sentimentalized or anthropomorphic concerns for the creatures. Presumably, such group dynamics are to be expected among the human observers of (or participants in) other large-scale digital ecosystems. Here we briefly recount three such events.

A "save the grendels" campaign attracted significant support. The grendels were put into the game principally to create a little stress in the norns' lives, by being aggressive and spreading diseases. However, the setting up of the "Grendel Liberation Front" on the newsgroup soon challenged this, tapping into an apparent fondness for the underdog. Grendels are nowadays being nurtured and "domesticated," rather than despised and hounded.

A similar degree of concern was shown by a number of European norn owners who were worried that their creatures would have language problems when passed around the Internet. After some discussion, the general opinion was that people should teach their norns English, as a suitable international language. The users were clearly prepared to put in extra effort themselves to ease the lives of their creatures.

An Australian family e-mailed one of us (Grand) about a norn that they were worried about and asked for help. The norn simply stood still after birth and did not move. Due to lack of food she was clearly close to death. Grand examined her and determined that her problem was that she was deaf, blind, and insensitive to touch: In the absence of stimulation she was deprived of any inducement to act. Exploration of her genotype revealed that a brain lobe gene had mutated, cutting off sensory inputs. Grand managed to modify the gene and cause it to switch on, so that the necessary lobe grew properly. After a period of rest and recuperation, the norn was sent back to her guardians in Australia, from where Grand later received a Christmas card updating him on her progress.

6 Conclusion

Although there are other entertainment software products that are (or claim to be) based on artificial life technologies, we know of none with a scope comparable to that of *Creatures*. In *Creatures*, groups of autonomous software agents interact with the human user, with each other, and with objects in the environment. In each agent, a heterogeneous continuous-time recurrent neural network with Hebbian learning is used for sensorimotor coordination. The activity of the agent's neural network may be affected by "hormones" in the agent's biochemistry, which also determines its energy metabolism. Details of each agent's network and biochemistry are genetically determined, using a marker-based encoding that permits growth in the genotype lengths and hence places no restrictions on the complexity-increasing evolution. There is a primitive linguistic capability offering the opportunity for cross-generational transmission of information through cultural channels.

Although developed purely as an entertainment product, the release and subsequent commercial success of *Creatures* inadvertently created a global digital ecosystem with around 500,000 computing nodes loosely coupled via Internet/Web connections, and up to around 5 million complex digital organisms active at any one time. In sheer scale this clearly supercedes the prior work of Ray or Yaeger and is on the same scale as Ray's plans for NetTierra. We do not see *Creatures* and NetTierra as rivals: Given their manifest differences, the two endeavors complement each other neatly.

We are conscious of the fact that there is a sense in which the launches of *Creatures* and *Creatures 2* represent, with hindsight, missed opportunities: unrecorded data concerning the spread of genes (or memes) in the early history of the product may now be lost forever. Nevertheless, we hope that the experiences we report here will draw the attention of other alife practitioners to issues arising in the development of global digital ecosystems. Although *Creatures* and *Creatures 2* are now completed products, development of *Creatures 3* is underway at CTL. Suggestions for how that, and future products, could be of service to the alife science community are welcome.[8]

References

1. Anark (1998). Website URL: http://www.anark.com/.

2. Bagley, R. J., & Farmer, J. D. (1992). Spontaneous emergence of a metabolism. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Vol. X of *Santa Fe Institute Studies in the Sciences of Complexity* (pp. 93–140). Redwood City, CA: Addison-Wesley.

3. Bagley, R. J., Farmer, J. D., & Fontana, W. (1992). Evolution of a metabolism. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: Proceedings*

8 Suggestions should be sent to the authors or directly to Toby Simpson (producer of *Creatures 3*) at toby@cyberlife.co.uk. Dave Cliff's current address is LTV/Agency, Hewlett-Packard Laboratories, Bristol BS34 8QZ, U.K.; cliffd@hplb.hpl.hp.com.

of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems, Vol. X of *Santa Fe Institute Studies in the Sciences of Complexity* (pp. 141–158). Redwood City, CA: Addison-Wesley.

4. Banzhaf, W. (1994). Self-organisation in a system of binary strings. In R. Brooks & P. Maes (Eds.), *Artificial Life IV* (pp. 109–118). Cambridge, MA: MIT Press/Bradford Books.
5. Batali, J. (1994). Innate biases and critical periods: Combining evolution and learning in the acquisition of syntax. In R. Brooks & P. Maes (Eds.), *Artificial Life IV* (pp. 160–171). Cambridge, MA: MIT Press/Bradford Books.
6. Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, *3*(4), 471–511.
7. Best, M. L. (1998). An ecology of text: Using text retrieval to study alife on the net. *Artificial Life*, *3*, 261–287.
8. Blumberg, B. (1994). Action-selection in hamsterdam: Lessons from ethology. In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (Eds.), *From Animals To Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior* (pp. 108–117). Cambridge, MA: MIT Press/Bradford Books.
9. Blumberg, B. (1996). *Old tricks, new dogs: Ethology and interactive creatures*. Unpublished doctoral dissertation, MIT Media Lab.
10. Cliff, D., & Miller, G. F. (1995). Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In F. Morán, A. Moreno, J. J. Merelo, & P. Chacón (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life (ECAL95)* (pp. 200–218). Berlin: Springer-Verlag.
11. Coderre, B. (1988). Modeling behavior in petworld. In C. Langton (Ed.), *Artificial life*, *Santa Fe Institute Studies in the Sciences of Complexity* (pp. 407–420). Redwood, CA: Addison-Wesley.
12. CTL (1997). Website URL: http://www.cyberlife.co.uk.
13. Edelman, G. M. (1989). *Neural Darwinism: The theory of neuronal group selection*. Oxford: Oxford University Press.
14. Elphick, M. R., Kemenes, G., Staras, K., & O'Shea, M. (1995). Behavioural role for nitric oxide in chemosensory activation of feeding in a mollusc. *Journal of Neuroscience*, *15*(11), 7653–7664.
15. Elphick, M. R., Williams, L., & O'Shea, M. (1996). New features of the locust optic lobe: Evidence of a role for nitric oxide in insect vision. *Journal of Experimental Biology*, *199*, 2395–2407.
16. Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom up*. Cambridge MA: MIT Press/Bradford Books.
17. Fontana, W. (1992). Algorithmic chemistry. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Vol. X of *Santa Fe Institute Studies in the Sciences of Complexity* (pp. 159–210). Redwood City, CA: Addison-Wesley.
18. Fujitsu Interactive Inc. (1998). Website URL: http://www.finfin.com/.
19. Grand, S., & Cliff, D. (1998). Creatures: Entertainment software agents with artificial life. *Autonomous Agents and Multi-Agent Systems*, *1*, 39–57.
20. Grand, S., Cliff, D., & Malhotra, A. (1997). Creatures: Artificial life autonomous software agents for home entertainment. In W. L. Johnson & B. Hayes-Roth (Eds.), *Proceedings of the First International Conference on Autonomous Agents* (pp. 22–29). New York: ACM Press. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP434.
21. Harvey, I. (1992). Species adaptation genetic algorithms: A basis for a continuing SAGA. In F. Varela & P. Bourgine (Eds.), *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91)* (pp. 346–354). Cambridge, MA:

MIT Press/Bradford Books. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP221.

22. Hinton, G., & Nowlan, S. (1987). How learning can guide evolution. *Complex Systems, 1*, 495–502.
23. Karakotsios, K., & Bremmer, M. (1992). *SimLife: The official strategy guide*. Prima Publishers.
24. Kauffman, S. (1993). *The origins of order: Self-organization and selection in evolution*. Oxford: Oxford University Press.
25. Koza, J. R. (1992). *Genetic programming: On the programming of computers by means of natural selection*. Cambridge, MA: MIT Press.
26. MacLennan, B. J. (1992). Synthetic ethology: An approach to the study of communication. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Vol. X of *Santa Fe Institute Studies in the Sciences of Complexity* (pp. 631–658). Redwood City, CA: Addison-Wesley.
27. MacLennan, B. J., & Burghardt, G. M. (1994). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior, 2*, 161–188.
28. Maes, P. (1995). Artificial life meets entertainment: Lifelike autonomous agents. *Communications of the ACM, 38*(11), 108–114.
29. Matarić, M. J., & Cliff, D. (1996). Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems, 19*(1), 67–83.
30. Maxis (1998). Website URL: http://www.maxis.com/.
31. Mayley, G. (1996). Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation, 4*(3), 213–234.
32. Morowitz, H. J. (1993). Artificial biochemistry: Life before enzymes. In C. G. Langton (Ed.), *Artificial Life III: Proceedings of the Workshop on Artificial Life*, Vol. XVII of *Sante Fe Institute Studies in the Sciences of Complexity* (pp. 381–388). Reading, MA: Addison-Wesley.
33. Noble, J., & Cliff, D. (1996). On simulating the evolution of communication. In P. Maes, M. J. Matarić, J.-A. Meyer, J. Pollack, & S. W. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 608–617). Cambridge MA: MIT Press/Bradford Books.
34. Nolfi, S., & Parisi, D. (1995). Evolving artificial neural networks that develop in time. In F. Morán, A. Moreno, J. J. Merelo, & P. Chacón (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life (ECAL95)* (pp. 351–367). Berlin: Springer-Verlag.
35. Nolfi, S., & Parisi, D. (1995). 'Genotypes' for neural networks. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 431–434). Cambridge, MA: MIT Press/Bradford Books.
36. Packard, N. H. (1989). Evolving bugs in a simulated ecosystem. In C. G. Langton (Ed.), *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems, Los Alamos, NM, Sept, 1987* Vol. VI of *Santa Fe Institute Studies in the Sciences of Complexity* (pp. 141–156). Redwood City, CA: Addison-Wesley.
37. PFMagic (1998). Website URL: http://www.pfmagic.com/.
38. Ray, T. S. (1992). An approach to the synthesis of life. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Vol. X of *Santa Fe Institute Studies in the Sciences of Complexity* (pp. 371–408). Redwood City, CA: Addison-Wesley.
39. Ray, T. S. (1994). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artificial Life, 1*, 179–210.

40. Ray, T. S. (1995). *A proposal to create a network-wide biodiversity reserve for digital organisms*. (Tech. Rep. TR-H-133). ATR, Kyoto, Japan. URL: http://www.hip.atr.co.jp/˜ray/pubs/pubs.html.

41. Ray, T. S. (1998). Continuing report on the network Tierra experiment. Web page, ATR, Kyoto, Japan. URL: http://www.hip.atr.co.jp/˜ray/tierra/netreport/netreport.html.

42. Ridley, M. (1993). *The red queen: Sex and the evolution of human nature*. London: Viking.

43. Sharkey, A. (Ed.). (1997). Combining artificial neural nets: Modular approaches. [special issue]. *Connection Science 9* (1).

44. Simpson, T. (1998). *Creatures 2: Strategies and secrets*. Alameda, CA: Sybex.

45. Sims, K. (1995). Evolving 3D morphology and behavior by competition. *Artificial Life, 1*, 353–372.

46. Sloman, A. (1990). Motives, mechanisms, emotions. In M. A. Boden (Ed.), *The philosophy of artificial intelligence*, Oxford Readings in Philosophy (pp. 231–247). Oxford: Oxford University Press.

47. Stork, D., Jackson, B., & Walker, S. (1992). 'Non-Optimality' via pre-adaptation in simple neural systems. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Vol. X of *Santa Fe Institute Studies in the Sciences of Complexity* (pp. 409–429). Redwood City, CA: Addison-Wesley.

48. Terzopoulos, D., Tu, X., & Grzeszczuk, R. (1995). Artificial fishes: Autonomous locomotion, perception, behavior, and learning in a simulated physical world. *Artificial Life, 1*, 327–351.

49. Todd, P. M., & Miller, G. F. (1993). Parental guidance suggested: How parental imprinting evolves through sexual selection as an adaptive learning mechanism. *Adaptive Behavior, 2*, 5–47.

50. Travers, M. (1988). Animal construction kits. In C. Langton (Ed.), *Artificial Life, Santa Fe Institute Studies in the Sciences of Complexity* (pp. 421–442). Redwood City, CA: Addison-Wesley.

51. Van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory, 1*, 1–30.

52. Werner, G. M., & Dyer, M. G. (1992). Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Vol. X of *Santa Fe Institute Studies in the Sciences of Complexity* (pp. 659–688). Redwood City, CA: Addison-Wesley.

53. Yaeger, L. (1993). Computational genetics, physiology, metabolism, neural systems, learning, vision, and behavior or PolyWorld: Life in a new context. In C. G. Langton (Ed.), *Artificial Life III: Proceedings of the Workshop on Artificial Life*, Vol. XVII of *Sante Fe Institute Studies in the Sciences of Complexity* (pp. 263–298). Reading, MA: Addison-Wesley.

54. Zaera, N., Cliff, D., & Bruten, J. L. (1996). (Not) evolving collective behaviors in synthetic fish. In P. Maes, M. J. Matarić, J.-A. Meyer, J. Pollack, & S. W. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 635–644). Cambridge MA: MIT Press/Bradford Books.