

**Computational Genetics, Physiology, Metabolism,
Neural Systems, Learning, Vision, and Behavior
or
PolyWorld: Life in a New Context**

**Larry Yaeger
Apple Computer, Inc.
One Infinite Loop, MS 301-3M
Cupertino, CA 95014
larryy@apple.com**

1. Introduction

The study of living systems has taken many forms, from research into fundamental physical processes to ethological studies of animal behavior on a global scale. Traditionally these investigations have focused exclusively on “real” biological systems existing in our world’s ecological system. Only recently have investigations of living systems begun to occur in “artificial” systems in computers and robotic hardware.

The potential benefits of an enhanced understanding of living systems are tremendous. Some are of a grand scale and are intuitively obvious, such as improvements in our ability to manage our own real ecosystems, the development of true machine intelligence, and the possibility of understanding our own mental and physiological processes. Some are of a more prosaic scale, but more accessible thereby, and perhaps of more immediate utility, such as simple learning systems, robust pattern classifiers, general purpose optimization schemes, robotic controllers, and evolvable software algorithms. The technological issues of the study of Artificial Life (ALife) are well laid out by Langton [27] in the proceedings of the first ALife workshop; the societal and philosophical implications of ALife are well presented by Farmer and Belin [16] in the proceedings of the second ALife workshop.

This paper discusses a computer model of living organisms and the ecology they exist in called PolyWorld. PolyWorld attempts to bring together all the principle components of real living systems into a single artificial (man-made) living system. PolyWorld brings 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 inter-species competition. Predation, mimicry, 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.

In many ways, PolyWorld may be thought of as a sort of electronic primordial soup experiment, in the vein of Urey and Miller’s [33] classic experiment, only commencing at a much higher level of organization. While one could claim that Urey and Miller really just threw a bunch of ingredients in a pot and watched to see what happened, the reason these men made a contribution to science rather than *ratatouille* is that they put the *right* ingredients in the *right* pot ... and watched to see what happened. Here we start with software-coded genetics and various simple nerve cells (light-sensitive, motor, and unspecified neuronal) as the ingredients, and place them in a competitive ecological crucible which subjects them to an internally consistent physics and the process of natural selection. And watch to see what happens.

Due especially to its biological verisimilitude, PolyWorld may serve as a tool for investigating issues relevant to evolutionary biology, behavioral ecology, ethology, and neurophysiology. The original motivations for its design and implementation, however, were three-fold: (1) To determine if it is possible to evoke complex ethological-level survival strategies and behaviors as emergent phenomena (without their being programmed in), (2) To create *artificial* life that is as close as possible to *real* life, by combining as many critical components of *real* life as possible in an *artificial* system, and (3) To begin exploring Artificial Life as a path toward Artificial Intelligence, utilizing the same key elements that led to natural intelligence: the *evolution of nervous systems* in an *ecology*.

This paper will discuss the design principles employed in PolyWorld, the “species” that have evolved in various simulations, and the group and individual behaviors observed in these simulations.

2. Background

This work owes much in terms of inspiration to the work of W. Grey Walter [46,47,48], Valentino Braitenberg [3], Richard Dawkins [9,10,11], John Holland [21], Ralph Linsker [29,30,31], and John Pearson [37].

Walter’s early work with simple electronic “turtle” nervous systems, and Braitenberg’s “vehicles” suggested the whole *bottom-up*, synthesis-before-analysis approach, along with the viability of making somewhat arbitrary connections between simple sensory inputs and simple motor controls. PolyWorld (PW) diverges from these works by encapsulating its synthetic organisms in a simulated world, by employing neural systems and learning rules from the world of computational neurophysiology, and by supporting a range of interactions between organisms. And though a number of other researchers (Travers [45]; Wharton and Koball [49]; and even a commercial product from Bascom software, the author of which is not known) have built simple Braitenberg vehicle simulators (or actual physical models in the case of Wharton and Koball), these typically concentrated on a wiring-diagram user interface, and implemented vehicles through only level number 2 (of 14). PW, on the other hand, takes note of the fact that by as early as Vehicle 4, Braitenberg invoked a form of natural selection, and supports the *evolution* of its organisms’ “wiring diagrams”, rather than having them specified by hand. The neural systems of PW also utilize Hebbian learning during the lifetime of an individual, which is undoubtedly purposefully similar to Braitenberg’s “mnemotrix wire”.

Richard Dawkins’s writings communicate both the beauty and the effectiveness of evolutionary dynamics. In personal communications, he has also brought out key issues in speciation, such as the isolation of populations and the reduced viability of divergent species interbreeding, that have become important elements of this simulator.

The artificial neural systems employed in PW are based on Hebbian learning, and a novel approach to network architecture specification. Besides the obvious importance of Donald Hebb’s [18] research and speculations, their instantiation in the work of Ralph Linsker and John Pearson has guided the selection of these particular techniques for use in PW. Linsker’s work demonstrated that Hebbian learning, as employed in PW, can and will self-organize important types of neural response patterns observed in early visual systems of real organisms. John Pearson, working with Gerald Edelman, utilized a variant of Hebbian learning and successfully demonstrated important principles of neuronal and synaptic self-organization—cooperation and competition (for representing their observed inputs)—that again correspond well to phenomena observed in real living systems. PolyWorld takes this unsupervised learning technique, and embeds it in arbitrary, evolving neural architectures, and then confronts the simulated neural system with survival tasks in a simulated ecology.

In the last couple of decades, a number of researchers have developed computational ecologies targeted at various scientific issues. Conrad [7,8] and Packard [35] have built systems to explore fundamental principles of evolutionary dynamics. Jefferson *et al* [23], and Collins and Jefferson [6] have constructed systems dealing with evolutionary mechanisms, behavioral goals, and learning architectures (Finite State Automata vs. Neural Networks). Taylor *et al* [42] developed a system to investigate the relationship between individual behavior and population dynamics. Ackley & Littman [1] built such a simulator to demonstrate a novel mechanism by which evolution can guide learning. Peter Todd and Geoffrey Miller [32,43,44] have explored evolutionary selection for different learning algorithms in organisms with simple vision systems and an innate sense of “smell” that functions with varying degrees of accuracy. Danny Hillis [19,20] has used simple computational ecologies to evolve “ramps”, and exchange-sort algorithms. Core Wars [13,14,15] is a non-evolving ecology of code fragments, and Rasmussen’s VENUS [38] is an evolving system based largely on Core Wars. Thomas Ray [39] has also developed a computational ecology, Tierra, based on evolving code fragments. And John Koza [26] has developed a system for evolving LISP functions that he terms “Genetic Programming”. PolyWorld, in its original conception, was targeted principally at the evolution of neural architectures for systems faced with complex behavioral tasks; however, its biologically motivated behavioral and reproductive strategies, and the evolutionary mechanisms employed also make it suitable for use in behavioral ecology and evolutionary biology. The extent of PW’s fidelity to biological systems, together with its unique use of a naturalistic visual perceptive

system to ground its inhabitants in their environment distinguish it significantly from previous ecological simulators.

John Holland's ECHO system explicitly models a form of predation, involving "offense" and "defense" genes that determine the outcome of violent encounters. Holland notes that in his system, this form of predation was essential to the evolution of complex genomes. Though not as crucial to PW's genetic complexity, predation was also designed into PW from the beginning. In PW, genes also affect the outcome of violent encounters between organisms, but more indirectly through their "physiological" characteristics (strength and size). There is also a behavioral component to the outcome of these encounters in PW, namely the degree of "volition" associated with the "fighting" behavior (the activation level of a predefined "fight" neuron), that differs from ECHO's handling of predation.

Belew *et al* [2] give an excellent overview of recent work in the area of evolving neural networks. Reviewed briefly there, and presented in detail in their own paper, Harp *et al* [17] have developed a scheme for evolving neural architectures that has an element of ontogenetic development. Their approach involves a set of synaptic projection radii between neuronal "areas". PW's scheme for evolving architectures relies on the specification of connection densities and topological distortion of connections between neuronal groups. These architectural criteria are represented in the genome, and then expressed as an organism's neural architecture at "birth". This technique, though perhaps not quite as *developmental* as Harp's approach, or the non-neural yet very biologically motivated cellular growth work of de Boer *et al* [12], has the strengths of being much more developmental (and representationally compact) than a simple encoding of synaptic efficacy in the genes, and being computationally very efficient. It attempts to capture the statistical results of development, without the necessity of modeling the developmental process itself.

David Chalmers [4] has experimented with evolving supervised neural network learning algorithms, successfully evolving the classic "delta rule" for a linear, single layer perceptron, and speculated on applying this "genetic connectionism" approach to other architectures and learning algorithms. He also varies the diversity of his learning tasks, and demonstrates a correlation between this diversity and the generality of the evolved learning algorithm, similar to the correlation observed between amount of training data and generalization in supervised, "Back-Prop" neural networks. Though the evolution of unsupervised learning algorithms is an area of special interest to the author, the current version of PW has the classic "Hebb rule" built in. Neural architectures are, however, evolved in PW. Interestingly, by permitting free movement in a simulated environment, PW effectively can generate an unlimited amount of diverse input for the neural mechanisms employed by its denizens.

Nolfi *et al* [34] and Parisi *et al* [36] have explored evolving the connection strengths in small, fixed-architecture feed-forward neural networks controlling simple movement strategies in organisms evolved to seek food. The organisms are directly provided with angle and distance to food items, and are alone in their environment. Nolfi, Parisi, *et al* also introduce a "self-supervised" learning technique, using the traditional back-propagation-of-error algorithm, and demonstrate an improvement in *evolved* foraging efficiency associated with a *learned* ability to predict the sensory consequences of motor activity. PW employs an unsupervised learning algorithm and arbitrary neural architectures, with a more biologically-motivated vision mechanism, as well as a competitive ecology.

For the purpose of computer graphics animation, Renault *et al* [41] have experimented with visual systems for controlling computer-generated characters. Their system goes beyond visual processing, however, to include unique object identification and distances to objects as part of the input to the character control programs. These control programs are rule-based and completely hand crafted, specifically to provide obstacle avoidance. In contrast, PW uses only the pixel colors associated with visual processing, and provides these as input to the non-rule-based neural systems of evolving organisms, without specifying the meaning or use of this information.

Dave Cliff [5] has implemented a neural visual system for a simulated fly, and states that it is only by a grounding perceptive mechanism such as vision that neural models can be made sense of. For the purposes of his simulation, the model fly is attached to a rotating, but otherwise unmoving test-stand similar to real experimental setups. Organisms in PW use vision as their primary sense mechanism, but are free to explore their environment, and must do so effectively—using their vision to guide a suite of primitive behaviors—in order to survive and reproduce.

One of the first decisions necessary when commencing an investigation into artificial living systems is that of scale: At what level of detail is it desirable to specify the parameters and underlying models of the simulation, and at what level does one wish to observe the resultant behaviors? The study of real living systems has spanned many physical and temporal scales: from molecular level

biochemical processes that take place in nanoseconds, through cellular level neural processes with time scales of a few milliseconds, to global evolutionary processes occurring over geological time scales. Given current constraints on compute power, it is simply not feasible to begin computation with sub-atomic physics and expect to observe ethological behaviors. Since ecology-level dynamics were the desired output level of the system being designed, it was clear that behavior models for PW's individual organisms could not be *too* complex. However, a desire to avoid rule-based behavior specification led to a decision to model the organisms' behaviors at the neuronal level. Since even natural evolutionary forces are constrained by their previous successes, the real world has filled up with organisms exhibiting a wide range of variations on assemblages of neuronal cells (in addition to other cell types, of course). Modeling PW's organisms at this level permits us to sidestep millions of years of evolution, while still taking advantage of its results to date.

3. Overview

PolyWorld is an ecological simulator of a simple flat world, possibly divided up by a few impassable barriers, and inhabited by a variety of organisms and freely growing "food". The inhabiting organisms use vision as input to a neural network brain that employs Hebbian learning at its synapses. The outputs of this brain fully determine the organisms' behaviors. These organisms and all other visible constituents of the world are represented by simple polygonal shapes. Vision is provided by rendering an image of the world from each organism's point of view, and using the resulting pixel map as input to the organism's brain, as if it were light falling on a retina.

A small number of an organism's neurons are predetermined to activate a suite of possible primitive behaviors, including eating, mating, fighting, moving forward, turning, controlling their field of view, and controlling the brightness of a few of the polygons on their bodies. Organisms expend energy with each action, including neural activity. They must replenish this energy in order to survive. They may do so by eating the food that grows around the environment. When an organism dies, its carcass turns into food. Because one of the possible primitive behaviors is *fighting*, organisms can potentially damage other organisms. So they may also replenish their energies by killing and eating each other. Predation is thus modeled quite naturally.

The organisms' simulated physiologies and metabolic rates are determined from an underlying genome, as are their neural architectures. When two spatially overlapping organisms *both* express their *mating* behavior, reproduction occurs by taking the genetic material from the two haploid individuals, subjecting it to crossover and mutation, and then expressing the new genome as a child organism.

One way to look at this artificial world is as a somewhat complex energy balancing problem. The fittest organism will be the one that best learns to replenish its energies by eating, and to pass on its genes by mating. The particular patterns of activity that a successful organism engages in—the methods by which it sustains and reproduces itself—will be optimal for some particular fitness landscape. But since that fitness landscape depends upon the behavior of the world's other inhabitants, it must, per force, be a dynamic landscape. Since there is considerable variation in the placement and behavior of food and other organisms in the world, that fitness landscape is also fundamentally stochastic. Indeed, if the "fittest organism in the world" fails to find a suitable mate in order to pass on the important bits of its genetic material, then those genes will be lost... possibly for all time. Accordingly, every world has the potential to be quite different from every other world.

Throughout this paper, the term *created* is applied to organisms spontaneously generated by the system (like the initial seed population), while *born* is used to refer to organisms resulting from the mating behaviors of the organisms. Populations of organisms that have evolved a set of behaviors which allow them to replenish their numbers through *births*, with no further *creations* (after some point in time), are said to exhibit a Successful Behavior Strategy (SBS), or simply to be *successful*. Once an SBS has emerged, there is *no fitness function except survival*. Until an SBS has emerged, PW is run in a sort of "on-line Genetic Algorithm (GA)" mode (also known as a Steady State GA), with an *ad hoc* fitness function. During this stage, a minimum number of organisms may be guaranteed to populate the world. If the number of deaths causes the number of organisms extant in the world to drop below this minimum, either another *random* organism may be *created* by the system, or the offspring of two organisms from a table of the N fittest may be *created*, or, rarely, the best organism ever may be returned to the world unchanged (known as an *elitist* strategy in traditional GAs). This *ad hoc* fitness function rewards organisms for eating, mating, living their full life span, dying with reserve energies, and simply moving. Each reward category is normalized by the maximum possible reward in each

category, and has a specifiable scale factor to permit easy tuning of the fitness function. Some simulation runs acquire an SBS in the first seed population and never require this on-line GA stage. Others never acquire an SBS, and are considered unsuccessful simulations.

Current high end simulations typically involve over 300 organisms, with up to approximately 200 neurons each, and require about 13 seconds per time-step on a Silicon Graphics Iris 4D/240-GTX. With an average life span of about 500 time-steps, and a time-to-first-offspring of about 100 time-steps, this means that 500 generations can be run at this complexity in about 1 week. More modest simulations with around 100 comparable organisms require about 4 seconds per frame, and take a day or two for the same task. And at the low complexity end, simple demonstration worlds can be run in “real time”, at a few frames per second, and allow a more interactive experience for learning the system. Figure 1 shows a sample view of the PolyWorld environment.

4. Genetics

An organism’s genes completely encode both its “physiology” and its neural architecture. Table 1 lists the full complement of genes present in the organisms of PolyWorld.

- size
- strength
- maximum speed
- ID (green coloration)
- mutation rate
- number of crossover points
- life span
- fraction of energy to offspring
- number of neurons devoted to red component of vision
- number of neurons devoted to green component of vision
- number of neurons devoted to blue component of vision
- number of internal neuronal groups
- number of excitatory neurons in each internal neuronal group
- number of inhibitory neurons in each internal neuronal group
- initial bias of neurons in each non-input neuronal group
- bias learning rate for each non-input neuronal group
- connection density between all pairs of neuronal groups and neuron types
- topological distortion between all pairs of neuronal groups and neuron types
- learning rate between all pairs of neuronal groups and neuron types

Table 1. List of genes in organisms of PolyWorld.

All genes are 8 bits in length, and may be Gray-coded or binary-coded. All but the ID gene are used to provide 8 bits of precision between a specifiable minimum and maximum value for the corresponding attribute. For example, if the minimum possible size is `minSize`, and the maximum possible size is `maxSize`, and the value of the size gene (scaled by 255 to lie between 0.0 and 1.0) is `valSizeGene`, then the size of the organism with this gene will be:

$$\text{size} = \text{minSize} + \text{valSizeGene} * (\text{maxSize} - \text{minSize})$$

These extrema values, along with a variety of other controlling parameters for the simulation, are contained in a “worldfile” that is read by the simulator at startup.

The first 8 genes control the organism’s simulated physiology. Its size and strength affect both the rate at which it expends energy and the outcome of “fights” with other organisms. In addition, its size is related directly to the maximum energy that it can store internally. The next gene, maximum speed, also affects its “metabolic” rate.

The ID gene’s only function is to provide the green component of the organism’s coloration at display time. Since organisms can actually see each other, this could, in principle, support mimicry. For example, a completely passive species could evolve to display the green coloration of a very aggressive species if it were of selective advantage. It might also be possible to attract potential mates

by displaying the green coloration of food, though this might be of limited survival value. (In practice, however, neither of these somewhat sophisticated evolutionary responses has yet been observed.)

Mutation rate, the number of crossover points used during reproduction, and maximum life span were placed in the genes in order to permit a kind of meta-level genetics, and in recognition of the fact that these parameters were themselves evolved in natural systems. They are, however, typically constrained to operate within “reasonable” limits; 0.01 to 0.1 for mutation rate, 2 to 8 for number of crossover points, and a few hundred to a few thousand “time-steps” for life span.

The final *physiology* gene controls the fraction of an organism’s remaining energy that it will donate to its offspring upon birth. The offspring’s total available energy on birth is the sum of these contributions from the two parents. Accordingly, at least one aspect of sexual reproduction may be captured by PW’s evolutionary “biology”: it is entirely possible for two interbreeding sub-species to be almost identical genetically, differing only in the amount of personal energy devoted to the reproductive process. PW has not yet been instrumented to observe for this phenomenon.

The remaining genes are used to define the organism’s neural architecture. These control parameters will be discussed in the section on Neurons and Learning. It should be noted here, however, that one of the motivations for this method of specifying the neural architecture was to reduce the number of genes necessary to specify the neural system. Early versions of PW used a simpler, fully recurrent neural architecture, and maintained a complete matrix of synaptic efficacies between all pairs of neurons in the genes. For 200 (N_N) neurons, this older model required 40,000 (N_N^2) genes. The current scheme supports evolving neural architectures which are fully specified by $12N_G^2 + 232N_G + 1026$, where N_G is the number of *internal* neuronal “groups” or clusters (*output* group sizes are fixed to 1, and *input* groups do not need biases, bias learning rates, or incoming synaptic connections). Thus, for 4 internal groups, with up to 32 neurons per group, plus up to 16 neurons per vision group (of which there are 3, one for each color component: red, green, blue), plus 2 other input groups (one neuron per group), plus the standard 7 output groups (one neuron each), a network of up to 185 neurons can be fully specified by just 2,146 genes. The large constants in this equation (232 and 1026) are due to the fixed set of input and output groups, and, especially, the desire to maintain each output neuron as a distinct group. Though the number of specifications are significantly reduced from a full crossbar matrix, this number still heavily outweighs the number of genes devoted to physiology. To permit a more robust exploration of the space of possible physiologies, then, one crossover during genetic reproduction is always forced to occur somewhere within the set of physiology genes. Note that since the minimum number of crossover points is typically set to 2, crossover will also be employed at some point(s) in the neurophysiology genes.

An organism’s genome is allocated and interpreted such that space is available for the maximum possible number of neuronal groups. That is, one of the parameters specified per pair of groups in a network with 3 groups out of a maximum of 5 groups would be accessed as:

1,1 1,2 1,3 -- -- 2,1 2,2 2,3 -- -- 3,1 3,2 3,3 -- --

where the entries marked “--” serve simply as place holders. This is as opposed to an access scheme looking like:

1,1 1,2 1,3 2,1 2,2 2,3 3,1 3,2 3,3

where the entries are contiguous. The reason for this is to permit a smoother evolution of these neural architectures. The addition of a fourth group would leave the old connections intact in the first representation, but not in the second. It is even possible for a useful sub-component of the architecture to ride along dormant in the genes to be expressed at a later time.

Though learning is supported in the neural network model employed in PW, only the architecture and some initial values are encoded in the genes; hence evolution in PW is purely Darwinian, not Lamarckian.

As in most GA’s, when an organism is *created* from scratch, the bits in its genes are first zeroed, and then turned on with a certain bit probability. Unlike most GA’s, it is possible to specify a range of legal bit probabilities, rather than always using 0.5. The bit probability for an individual organism is then randomly selected from this range and used to initialize the organism’s bit-string genome. So the probability of a bit being on in a particular organism will depend on the value

randomly selected from the specified range, while the probability of a bit being on in the population as a whole will just be the mean of the specified range (0.5 if the range is 0.0 to 1.0). This permits a wider variance in local, organism-specific bit probabilities in early populations, rather than depending entirely on mutation and cross-over to so shuffle the bits. Whether this is of any real value should be tested in a simpler GA system, and may be problem-specific in any event. Here it was felt that both the older fully-recurrent neural network architecture and the later evolving neural architectures were more likely to have behaviorally/evolutionarily useful solutions with lower bit densities; this provided a mechanism for so biasing the initial seed population without ruling out selection towards the unexpected end of the spectrum.

There is an optional “miscegenation function” (so dubbed by Richard Dawkins), that may be used to probabilistically influence the likelihood of genetically dissimilar organisms producing viable offspring; the greater the dissimilarity, the lower the probability of their successfully reproducing. This function is not typically invoked until after a (specifiable) “significant” number of births without an intervening creation in order to allow the early stages of the simulation to explore as many genetic recombinations as possible. It can also be turned off entirely.

5. Physiology and Metabolism

As discussed above, the simulated physiology of PolyWorld’s organisms is determined by their genes. The *size* of the organism directly affects the maximum amount of energy that the organism can store. If an organism’s size is allowed to range between *minSize* and *maxSize*, and its energy capacity ranges between *minECap* and *maxECap*, then a given organism’s actual energy capacity, *ECap* is given as:

$$ECap = minECap + (size - minSize) * (maxECap - minECap) / (maxSize - minSize)$$

Similar linear relations are used to determine the influence of an organism’s size on the rate at which it expends energy during forward or turning movement (relative to a specifiable maximum-size-penalty), and a size-advantage it will have during a fight with another organism (relative to a specifiable maximum-size-advantage).

An organism’s *strength* also affects both its energy expenditure and its advantage in a fight. Strength directly scales the total energy used in a given time step, and thus usually ranges around 1.0 (typically 0.5 to 2.0). An attacker’s strength also scales the effect on the victim’s energy loss (fighting is discussed in more detail below in the section on Behavior).

The energy expended by an organism’s neural processing is determined linearly from the number of neurons and the number of synapses it has. A maximum number of neurons and synapses is determined from the control parameters for the entire world, then each individual’s neural energy expenditure is computed relative to these maxima. Globally applied “neuron-to-energy” and “synapse-to-energy” conversion factors then multiply these scaled neuron and synapse counts to determine the actual energy expended per time step.

There are similar *behavior-to-energy* conversion factors for each of the primitive behaviors (eating, mating, fighting, moving, turning, focusing, and lighting). The total energy expended in a time step is then the activation (0. to 1.) of the corresponding output/behavior neuron multiplied by that behavior’s energy-conversion factor, summed over all behaviors, plus the neural energy expenditure, plus a specifiable fixed energy drain, with this sum finally scaled by the organism’s strength.

As should be evident, there are clear energy conservation benefits to being small and weak, yet there are clear predatory advantages to being large and strong. Size also permits an overall greater capacity to store energy, thus making energy available for additional behavioral activity, including reproduction. The interplay between these opposing advantages is intended to produce niches in the fitness landscape, which may change over time. There are similar opposing pressures between energy expenditure and visual acuity on the number of input neurons devoted to vision.

There are two classes of energy storage in each organism: health-energy, and food-value-energy. Both are replenished by eating food. Both are depleted by neural activity and by engaging in the various behaviors. But when an organism is attacked, only its health-energy is depleted by the attack. If this health-energy reaches zero, the organism dies. When an organism dies it is converted into a piece of food containing an amount of energy equal to the organism’s food-value-energy. This separation of health-energy from food-value-energy makes the predator-prey interactions quite natural;

i.e., it is possible for an organism to be killed by having its health-energy driven to zero, while still maintaining a relatively high food value for the attacker.

An organism's food-value-energy will always be greater than or equal to its health-energy, yet both classes of energy have the same maximum capacity. Accordingly, an organism may continue to eat to replenish its health-energy after its food-value-energy has reached capacity. It is the health-energy that is provided as input to the neural network (see next section), and that is used to determine the amount of energy to be transferred to offspring.

Purely for the purposes of display, an organism's length and width are scaled by (the square root of) its maximum speed, length being multiplied, width being divided. Thus faster individuals will appear longer and sleeker, while slower individuals will appear shorter and bulkier. Since an organism's visual acuity is subject to evolutionary pressures, it is conceivable that an organism might emerge that was able to ascertain another organism's maximum speed purely from its shape, if there was a great enough advantage to the acquisition of this information.

6. Neural Systems and Learning

The *inputs* to an organism's neural network "brain" are its "vision", the current normalized level of its internal health-energy store, and a random value. The *outputs* are the suite of 7 possible primitive behaviors (eating, mating, fighting, moving, turning, focusing, and lighting). The *internal* neurons and all of the synaptic connections have *no* prespecified functionality; their utility is determined entirely by genetics and natural selection.

The form of an organism's brain, or neural system, is fully characterized by a set of parameters that are encoded in its genes. Referring back to Table 1, notice that the number of neurons devoted to each color component of vision is specified separately, permitting a specialization for more resolution in the most effective color, should this be of selective advantage. These numbers typically range between 1 and 16 neurons per color.

Next is a parameter that specifies the number of internal neuronal groups or clusters. This typically ranges from 1 to 5. In addition, there are 5 input groups (red vision, green vision, blue vision, energy level, and random), plus 7 output groups (the behaviors listed above).

Each neural group may have distinct populations of excitatory (e-) and inhibitory (i-) neurons. The number of e- and i-neurons are specified on a per group basis, and typically range between 1 and 16 neurons of each type. Synaptic connections from e-neurons are always excitatory (ranging from 0.0 to a specifiable maximum efficacy). Synaptic connections from i-neurons are always inhibitory (ranging from -1.e-10 to the negative of the maximum efficacy).

Though the bias on each of the non-input neurons varies during the simulation, the initial values for these biases and their learning rates are specified on a per group basis, for each of the non-input neural groups. Biases are updated by a Hebbian learning rule, as if it were a synaptic connection to a neuron that was always fully activated, but unlike other synapses in this network, the bias may change sign. Biases typically range from -1.0 to 1.0, and bias learning rates typically range from 0.0 to 0.2.

The remaining parameters—connection density (CD), topological distortion (TD), and learning rate (LR)—are all specified for each pair of neuronal groups *and* neuron types. That is, separate values for each of these parameters are specified for the excitatory-to-excitatory (e-e), excitatory-to-inhibitory (e-i), inhibitory-to-inhibitory (i-i), and inhibitory-to-excitatory (i-e) synaptic connections between group i and group j, for each pair of groups i and j.

Connection density, as the name suggests, is used to determine the extent of the connectivity between neuronal groups. The number of e-e synapses between group i and group j is given by the nearest integer to $CD_{e-e}(i,j) * N_e(i) * N_e(j)$, where $CD_{e-e}(i,j)$ is the e-e CD from group j to group i, $N_e(i)$ is the number of e-neurons in group i, and $N_e(j)$ is the number of e-neurons in group j. Similar expressions hold for the other types of connections between all pairs of groups. CD can range from 0.0 to 1.0.

Topological distortion is used to determine the degree of disorder in the mapping of synaptic connections from one group to the next. That is, for a TD of 0.0, synapses are mapped to perfectly contiguous stretches of neurons in the adjacent layer; for a TD of 1.0, synapses are mapped in a completely random fashion between adjacent layers. Thus retinotopic maps such as are observed in natural organisms can be enforced (or not) at the architectural level (as well as resulting from the learning process). TD typically ranges from 0.0 to 1.0.

Learning rate controls the Hebbian learning process at the synapses between each pair of neuronal groups. This permits natural selection to favor hardwired, “instinctive” connections for some neural pathways, while supporting learning in other pathways. LR typically ranges from 0.0 to 0.2.

This method of specifying the neural architecture is fairly general, and is not biased for any particular neural organization. Possibly, one might expect to evolve a preponderance of inhibitory connections, especially locally, if the simulated neural architectures evolve to match real neural systems; yet the possibility exists for establishing local excitatory connections (such as are found in CA3 in the hippocampus). The technique does not, however, explicitly model architectures whose characteristics are heavily based upon spatial organization (such as the parallel fibers originating from the granule cells in the cerebellum). A straightforward extension to the current method, that allowed unique specifications of the same parameters along multiple spatial dimensions, could account for such organizational schemes. However, with the limited compute resources currently being applied to PW simulations, and thus the limited number of neurons permitted in each brain, it was not deemed worthwhile to further decompose the groups into these spatial subcategories.

When an organism’s brain is “grown” from its underlying genome, the synaptic efficacies are randomly distributed between specifiable minimum and maximum values. The brain is then exposed to a sequence of mock visual inputs consisting of random noise, for a specifiable number of cycles. This is all pre-birth. In this fashion, it is unnecessary to store *any* synaptic efficacies in the genes. This approach was inspired by Linsker’s simulations of visual cortex, which gave rise to on-center-off-surround cells, orientation-selective cells, and so on, when exposed only to noise. The crucial aspects of the networks in this case are their architecture—layered receptive fields in Linsker’s case, evolved arbitrary topology in PW—and the learning rule—Hebbian learning in both cases.

It was debated whether to update all the organisms’ brains *synchronously* or not. That is, whether each organism’s neural network should be allowed to make a complete neural activation and synaptic learning pass with each time step. Even though it was desired to penalize organisms that evolved additional neurons and synapses, synchronous updating was ultimately selected, primarily because the corresponding structures in nature are executed in parallel, and penalties based on their serial implementation would be excessive. The penalty is more properly derived from the additional energy use associated with these additional neural structures.

At each time step, the input neurons are set to the appropriate values, corresponding to the organism’s visual field, its current health-energy level, and a random number. New neuronal activations are computed by the simple formulae:

$$x_i = a_j^t s_{ij}^t$$

$$a_i^{t+1} = 1 / (1 + e^{-x_i})$$

where a_j^t is the neuronal activation of neuron j at time t (the beginning of this time step), s_{ij}^t is the synaptic efficacy from neuron j to neuron i at time t , a_i^{t+1} is the neuronal activation of neuron i at time $t+1$ (the end of this time step), and c_{kl} is a specifiable logistic slope.

The synaptic efficacies are then updated according to a Hebb rule, as:

$$s_{ij}^{t+1} = s_{ij}^t + c_{kl} (a_i^{t+1} - 0.5) (a_j^t - 0.5)$$

where s_{ij}^{t+1} is the synaptic efficacy from neuron j to neuron i at time $t+1$, and c_{kl} is the learning rate for connections of type c (e-e, e-i, i-i, or i-e) from group l to group k . An optional multiplicative decay term may also be applied to the synaptic efficacy.

This simple “summing and squashing” neuron and Hebbian update rule are certainly coarse abstractions of the complexities observed in real neural systems. Credence is lent to these particular abstractions by the previously quoted simulation work of Linsker, Pearson, and others, and by Linsker’s and others’ information-theoretic analytical work on such systems, which suggest that they may capture the information-processing attributes of real neural systems, if not their precise method of action. These neuronal and learning models were selected for use in PW based on these results and the models’ computational tractability.

During the course of a simulation, neural and synaptic activities may be monitored for a number of organisms in the world (the top five “fittest”, according to the *ad hoc* fitness function discussed earlier, even if it is not being used to create new organisms). A few examples are shown in Figure 2.

Early simulations with PW had much simpler, fully recurrent neural architectures. Though not particularly representative of real biological neural architectures, acceptable behavior strategies were evolved, and some of the results being presented are from organisms using these early networks.

7. Vision

The color vision supplied as input to the organism is first rendered at the minimum window size permitted on the Iris + 1 (because even-sized buffers can be accessed faster), or 22 x 22 pixels. The pixel row just above vertical center is then properly anti-aliased into whatever number of visual neurons an organism has. Even though organisms and the environment of PW are three-dimensional, the organisms’ vision consists of just this one-dimensional strip of pixels, rather than the complete pixel map. Since the organisms are confined to motion on a ground-plane, it was felt that the benefit derived from computational efficiency outweighed the small loss of information resulting from this restriction.

As was discussed above in the Genetics section, the number of neurons devoted to each of the color components is evolved independently (though they are adjacent on the genome, and so may tend to crossover together).

As is indicated in Figure 2, an organism’s vision is shown in the display of the brain internals that may be invoked interactively for some of the “fittest” individuals. In addition, the full 22 x 22 pixel map for each of the organisms is usually displayed at the top of the screen. This is mostly for a “reality check”—visual reassurance that the organisms are seeing what they would be expected to see, and may be disabled for a slight speed gain.

The vertical field of view of the organisms is fixed at 10°, since they only see a strip of pixels just above the center of the image. Their horizontal field of view, however, is under their own “volitional”, neural control. That is, the activation of the *focusing* neuron is mapped between a minimum and maximum field of view (typically 20° to 120°). In principle, this might permit some depth of field determinations based on cyclic focusing operations, though it’s highly doubtful that anything so sophisticated could emerge in the limited neural systems employed by the organisms so far.

This type of direct perception of the environment should answer one of cognitive psychology’s most frequently sounded complaints against traditional AI: The organisms of PW are “grounded” in their environment by their sense of vision.

8. Behavior

A suite of primitive behaviors is made available to all organisms in PolyWorld, namely:

- eating
- mating
- fighting
- moving
- turning
- focusing
- lighting

All of these behaviors are expressed by raising the activation level of a prespecified neuron in the brain. Given computational constraints, it was felt that a minimum number of cycles should be devoted to motor activity, hence this simple one-neuron-one-behavior mapping. The first three behaviors, eating, mating, and fighting, all have some associated threshold that must be exceeded before the activity is initiated. Energy is expended by each of the behaviors, including eating. The energy expenditure rates are controllable by scale factors (see Physiology and Metabolism) in the “worldfile” (see The New Context).

Eating is an organism’s method for replenishing depleted energy stores. In order to eat, an organism’s position must cause it to overlap a piece of food. The amount of energy consumed is proportional to the activation of the *eating* neuron, once that activation exceeds a specifiable threshold.

Mating is an organism's method for reproducing. In order to reproduce, an organism's position must cause it to overlap another organism, and *both* organisms must express their mating behavior in excess of a specifiable threshold. The outcome of the reproductive attempt may be affected by the miscegenation function (see Genetics), or by the maximum number of organisms permitted in the world (see The New Context). The organism's "desire" to mate (the activation level of its *mating* neuron) is mapped onto its blue color component for display purposes; this coloration is visible to other organisms as well as to human observers.

Fighting is an organism's method for attacking another organism. In order to successfully attack the other organism, the attacker's position must cause it to overlap the attackee. Only *one* organism need express its fighting behavior to successfully attack another. The energy that is depleted from the prey is a function of the volitional degree of the attack (the activation of the predator's *fight* neuron), the predator's current health-energy level, the predator's strength, and the predator's size. The product of these contributing factors from the predator is scaled by a global attack-to-energy conversion factor to make the final determination of the amount of energy actually depleted from the prey. If both organisms are expressing their fight behavior, the same computation is carried out reversing the roles of predator and prey. Note that while this system of fighting/predation will permit every extreme of interaction, including the equivalent of an ant attacking an elephant (or the equally pointless act of an elephant attacking an ant for food), such actions will be of an evolutionarily-useful, survivability-enhancing value comparable to those same ridiculous examples of real world behavior. In addition, since the level of expression of this behavior is under the volitional control of the organisms' nervous systems, a full spectrum from complete pacificity to uninterrupted fighting is possible (and exhibited in the "dervishes" species discussed in the Results section below). Each organism's desire to fight is mapped onto its red color component for display purposes; this coloration is visible to other organisms as well as to human observers.

Moving refers to an organism's forward motion. Unless an organism encounters a barrier, or the edge of the world, it will move forward by an amount proportional to the activation of its *moving* neuron.

Turning refers to a change in an organism's orientation on the ground-plane (yaw). An organism will turn about its vertical axis by an amount proportional to the activation of its *turning* neuron.

Focusing refers to an organism's control over its horizontal field of view. As discussed in the Vision section, the activation of an organism's *focusing* neuron will be linearly mapped onto a range of possible angles to provide its horizontal field of view. This makes it possible for an organism to use its vision to survey most of the world in front of it or to focus closely on smaller regions of the world.

Lighting refers to an organism's control over the brightness of a cap of several polygons on the front face of its "body". The activation of an organism's *lighting* neuron is linearly mapped onto the full 0 to 255 brightness range in all color components of these front polygons. Accordingly, a simple form of visual communication is possible, in principle, for the organisms inhabiting PW. (No evidence of their use of this form of communication has yet been found nor sought to date, though evidence of the organisms' use of vision for controlling locomotion has been observed.)

9. The New Context

The "world" of PolyWorld is a flat ground-plane, possibly divided up by a few impassable barriers, filled with randomly grown pieces of food, and inhabited by the organisms previously described.

The number of organisms in the world is controllable by several means. First, a maximum number of organisms is specifiable, in order to keep the problem computationally tractable. Second, a minimum number of organisms is specifiable to keep the world populated during the early on-line GA stage (see Genetics). Finally, an initial number of organisms is specifiable to determine how many individuals to seed the world with at the start of the simulation.

Food is grown at a specifiable rate up to a specifiable maximum number of *grown* food items. The number of food items may be guaranteed to be kept between a specifiable minimum and maximum food count. Subject to this maximum, food is also generated as the result of an organism's death. The amount of energy in a piece of food that is grown is randomly determined between a specifiable minimum and maximum food energy. The amount of energy in a piece of food resulting from the death

of an organism is that organism's food-value-energy (see Physiology and Metabolism) at death, or a specifiable minimum-food-energy-at-death.

An arbitrary number of barriers may be placed in the world, which inhibit movement of the organisms. These can serve to partially or completely isolate populations of organisms, and as such can contribute significantly to speciation (genetic diversity). For reasons of computational efficiency, they are typically placed parallel to the z (depth) axis, though this is not strictly necessary.

It is possible to manage the minimum, maximum, and initial numbers of organisms and food items, along with the *ad hoc* fitness statistics, simultaneously for a number of different independent "domains". These domains *must* be aligned parallel to the z (depth) axis, and typically, though not necessarily, coincide with the divisions imposed on the world by the barriers. This permits the simultaneous "culturing" of completely independent populations when barriers extend the full length of the world, or limits the spread of genes between domains to those resulting from actual movement of organisms when the barriers are arranged so as to leave gaps for organisms to travel through. If the domain fitness statistics were not kept separately, then genes from one domain could migrate to another domain by virtue of their global fitness during the start-up on-line GA phase.

It is possible to set a flag such that the edges of the world act as barriers (the usual), wrap around, or aren't there at all. In this last case, PW's ground-plane acts much like Braitenberg's *table top*, with organisms that move past the edge of the world dying instantly.

Various monitoring and graphing tools exist to assist in following the progress of a simulation and in developing an understanding of the evolutionary and neural dynamics at work. As was mentioned earlier (in the section on Neural Systems and Learning), a display of the internal workings of any of the five "fittest" organisms may be called up at any time. In addition, a small window that maintains an overhead view of the world will automatically track that same organism upon request. This overhead window may also be zoomed in and out to follow the organism more closely.

Also available are graphic displays of the time histories of certain quantities of interest, including: (1) population sizes (overall and per domain), (2) the past maximum, current maximum, and current average values of the *ad hoc* fitness function, (3) the ratio of the number of organisms "born" (by mating) to the sum of the number of organisms born and created, and (4) the ratio of the difference of food-energy *in* and food-energy *out* to the sum of these two values. These last two items in particular are important gauges of the course of the simulation. Item (3) will start at 0.0 and asymptote to 1.0 for successful simulations, in which at least one species has emerged with an SBS; it will peak well below 1.0 for unsuccessful simulations. Item (4) ranges from -1.0 to 1.0, and should asymptote to 0.0, for a world where energy is conserved. Three values are actually plotted for item (4): (a) the total food-energy, including the initial seeding of the world, which starts at 1.0 and should asymptote to 0.0, (b) the average food-energy, excluding the initial seeding of the world, which starts at 0.0, and rapidly becomes negative, but should also asymptote to 0.0, and (c) the current food-energy on a time-step by time-step basis, which fluctuates rapidly, but should cluster around the average food-energy.

One additional display can graphically present the results of an analysis of the genetic variability in the population. All pairs of organisms are examined to determine the magnitude of the Hamming distance between them in gene space, normalized by the maximum possible genetic distance between two organisms. These normalized distances are divided into as many distinct histogram bins as there are vertical pixels in the graph, and pixel brightness is used to indicate how many pairs of organisms fell into each bin. A single column of pixels thus shows the distribution of "genetic separation" for the entire population at a single point in time. A new column of pixels is added each time the genetic makeup of the population changes (each birth or death). The result is a complete time-history of genetic variability in the world. This approach's strength is that it is able to show such a complete temporal evolution of population-wide genetic variability at a glance. Its weakness is that by reducing genetic differences to a single number, one cannot tell the difference between many genes that are only slightly different and a few genes that are very different; nor is it possible, of course, to tell *which* genes differ.

All of the simulation control parameters and display options are defined in a single "worldfile" that is read at the start of the simulation. In addition, some of the display options can be invoked interactively at runtime.

There isn't space to go into many details of the code itself. However, it may be worth noting that it consists of about 15,000 raw (not compiled) lines of C++, and is entirely object oriented, *except* for a single routine devoted to handling the organism-organism and organism-food interactions (for reasons of computational efficiency). The organisms, food, and barriers are maintained in doubly-linked lists sorted on a single dimension (x). This simple data structure has minimal maintenance

overhead, yet rules out most non-intersections very well, and permits a sorting algorithm to be used that capitalizes on the expected frame-to-frame coherency of organism positions. It runs on a Silicon Graphics Iris (to take advantage of its hardware renderer for all the vision processing), and uses a set of object oriented C++ graphics routines (included in the line count above) that wrap around the standard Iris graphics library.

The tools currently available in PW for tracking population histories, fitnesses, genetic makeup, and so on, are primarily graphical, and exist largely for the purpose of monitoring the progress of the simulation. They exist more to help create intuitions than to support detailed quantitative analysis. The only exception to this rule is one particular measure of population-wide genetic diversity (the normalized Hamming distance between all pairs of organisms, discussed above), which can be recorded to a file on demand. The current version of PW does not purport to provide a universal set of tools for all the possible evolutionary biology and behavioral ecology experiments that one might wish to perform with it; it is doubtful whether all such possible tools could even be imagined in advance. This is one of the principle reasons that PW's full source code has been made available, in the hopes that anyone wishing to apply the PW simulation environment to a particular study could add their own data-gathering and data-analysis tools. And, hopefully, the object-oriented style used for programming PW will make it relatively easy to add a new graphical analysis tool, or to put in hooks to capture the precise data required for a particular study. The author has provided his email address and is willing to help interested researchers navigate in the code to help facilitate their research needs. If time permits, coding help may also be available from the author, and tools added for other studies will be integrated into the baseline code.

10. Results: Speciation and Complex Emergent Behaviors

Despite the variability inherent in different worlds, certain recurring "species" have occurred in a number of the simulations run to date. By "species" is meant: groups of organisms carrying out a common individual behavior that results in distinctive group behaviors. Since the selection of these behaviors are derived from the activity of their neural network brains, and the success of these behaviors is partially a function of their physiologies, both of which are in turn based on the genome of the organism, the behavioral differences may generally be traced to the organism's genetic code. Hence these behavioral differences are representative of different genetic species. No effort has been made to date to quantify or uniquely identify the genetic differences between these species, due mostly to constraints on the author's time (PW is still very much a work in progress). Examining potentially subtle differences in thousands of genes (or tens of thousands in the older fully recurrent nervous systems) will require some well designed and implemented graphical analysis tools (though some thought has gone into just what such a tool might look like, no code has been written). For now, these ethological-level behaviors (one of the original motivations for building PW, recall) may be the best way to begin developing some understandings and intuitions about the evolutionary dynamics possible in such a system.

A simulation is considered "successful" if and only if some number of species emerge with a Successful Behavior Strategy (SBS); these populations are capable of sustaining their numbers through their mating behaviors, and thus organism *creations* cease. The observational reports below only refer to such "successful" simulations. Accordingly, there are typically many tens to many thousands of examples of any particular behavior, especially in the simpler, more homogeneous species, where all members of the species exhibit approximately the same behavior, and where that species has usually occurred in multiple simulations.

The first of these species has been referred to as the "frenetic joggers". In an early simulation without barriers, without a miscegenation function, and with borders that wrap around (essentially forming a torus), a population emerged that basically just ran straight ahead at full speed, always wanting to mate and always wanting to eat. That particular world happened to be benign enough, that it turned out they would run into pieces of food or each other often enough to sustain themselves and to reproduce. It was an adequate, if not particularly interesting solution for that world. And without the miscegenation function or any physical isolation due to barriers, whatever diversity was present in the early world population was quickly redistributed and blended into a single species that completely dominated the world for as long as the simulation was run.

The second recurring species has been referred to as the "indolent cannibals". These organisms "solve" the world energy and reproduction problem by turning the world into an almost zero-

dimensional point. That is, they never travel very far from either their parents or their offspring. These organisms mate with each other, fight with each other, kill each other, and eat each other when they die. They were most prevalent in simulations run before the parents were required to transfer their own energies to the offspring; the organisms of these worlds were exploiting an essentially free energy source. With proper energy balancing, this behavior was reduced to only an occasional flare-up near corners of the world, where some organisms with limited motor skills naturally end up congregating, sometimes for quite extended periods of time. It turns out that the primary evolutionary benefit associated with this behavior was the ready availability of mates, rather than the “cannibalistic” food supply. This was determined by completely eliminating the food normally left behind by an organism’s death, yet still observing the emergence of such species. Large colonies of these indolent cannibals look from above like a continuous (non-gridded) version of Conway’s game of LIFE.

The third recurring species has been referred to as the “edge runners”. These organisms take the next step up from the cannibals, and essentially reduce their world to an approximately one-dimensional curve. They mostly just run around and around the edge of the world (system constraints prevent them from falling off in most simulations). This turns out to be a fairly good strategy, since, if enough other organisms are doing it, then some will have died along the path, ensuring adequate supplies of food. And mates are easily found by simply running a little faster or a little slower, running in the opposite direction, or simply stopping at some point and waiting for other runners to arrive (all of which behaviors have been observed). A form of this behavior persists even when barriers block access to the rest of the world; organisms still sometimes congregate along any edges, including the barriers. It has been suggested [22] that this may be a form of *behavioral isolation*, permitting this species to retain its genetic identity to the exclusion of other species.

Another species recently emerged as the first evolutionarily stable solution to a Braitenberg-style, “table top” world—one with deadly, drop-off edges. These “dervishes” evolved a simple rapid-turning strategy that kept them away from the dangerous edges of the world, and yet explored enough of the world to bring them into contact with food and each other. While this basic behavioral strategy persisted for many hundreds of generations, the dervish populations continued to explore optimum degrees of predation, in a sort of continuous prisoner’s dilemma over optimum degrees of cooperation. Indeed, largely uniform behaviors (as indicated by the organisms’ coloration) in whole sub-species were suggestive of a Tit-For-Tat strategy; even though variations in predation were being introduced into the communities from their peripheries (or through mutation), whole populations soon tended to adopt the same basic behavior patterns. Waves of varying levels of expression of this fighting behavior could be observed sweeping through several distinct populations, with the greatest variation in behaviors clearly seen at the boundaries between these populations.

The most interesting species and individuals are not so easily classified. In some worlds individuals’ behaviors have been so varied as to preclude any obvious classification into distinct species. In other worlds there appear to be multiple distinct species, with no single species obviously dominating. It is especially in these simulations that a number of complex, emergent behaviors have been observed, including:

- 1) responding to visual stimuli by speeding up,
- 2) responding to an attack by running away (speeding up),
- 3) responding to an attack by fighting back,
- 4) grazing (slowing upon encountering each food patch),
- 5) expressing an attraction to food (seeking out and circling food), and
- 6) following other organisms.

The first item is important in that it implies that conditions have been found that will cause evolution to select for the use of the organisms’ vision systems. All four of the earlier, simpler species’ behaviors would be appropriate even if these vision systems did not exist. Yet PW was built on the assumption that vision would be a powerful, useful sense mechanism that evolution could not fail to employ. Even a simple speeding up in response to visual stimulation could result in reaching food or a potential mate more effectively, and this was the first observed visual response to emerge.

The second and third items both represent reasonable responses to attack by a predator. Fleeing may reduce the effect of the attack, and fighting back is an energy-efficient use of the organism’s own ability to fight (as opposed to expressing the fight behavior continuously which would expend unnecessary energy).

Strategies four and five represent efficient feeding strategies. As simple a survival skill as grazing might seem—to simply notice when one’s internal energy is going up, and cease moving until it stops going up—it was not observed in early simulations. It is still not a wide-spread phenomenon, though a few instances have now been observed. Only the most recent simulation, as of this writing, has given rise to a population of organisms that seem to be able to actively seek out food and “orbit” it while eating; such “foraging” is clearly a valuable survival trait. These organisms appear to be drawn to the food as if there were a magnet or some point attractor located in the food and controlling the organisms’ behavior, though no such mechanism exists in PW. Their attraction to the food is purely a result of selection forces acting on the neural architecture connecting their vision systems to their motor systems.

The final, “following” strategy has also emerged only in this most recent simulation. Clearly of value, whether for seeking a prey or a mate, this represents the most complex coupling of the vision sense mechanism to the organisms’ motor controls yet observed. Small “swarms” of organisms, and one example of a few organisms “chasing” each other were even suggestive of simple “flocking” behaviors. The “swarm” appears to be a fairly stable formation, persisting for as long as it could be observed in the particular simulation. Having observed the driving forces in the simpler species in PW, this can be reasonably well understood: By staying within the swarm, reproductive partners are readily found (as with the “cannibals”), yet the swarm drifts, allowing its members to find new food sources as the old ones are depleted (more like the “dervishes”). To achieve this limited “flocking” or “swarming” behavior, an attraction to other bright (or moving—the neural architectures of these organisms have not been analyzed, and most definitely should be) objects is all that is required. As with the foraging behavior, this attraction is mediated exclusively through the action of natural selection on the organisms’ nervous systems.

Except for the last two behaviors, all of the species and behaviors discussed in this section have been observed in multiple, distinct simulations. And these final two behaviors are exhibited by tens or hundreds of organisms in the single most recent simulation in which they occurred.

All of these behaviors, being inherently temporal phenomena, require some sort of temporal medium for display. Short video clips of most of the above species and behaviors should be available in a companion videotape released by the publisher of this book.

The various species and individual behaviors discussed in this section are necessarily snapshots of a moving target. The PW code has undergone almost continuous development throughout the period of time covered by the simulations discussed above. In particular, and probably significantly, the most complex “foraging” and “following” behaviors were never observed in organisms based on the older, fully recurrent networks. It will not come as a surprise to anyone with even the most limited knowledge of neurophysiology that the architecture of an organism’s nervous system is vitally important. The arbitrary neural architectures of the most recent versions of PW appear to offer significant advantages for exploring the space of “useful” nervous systems. And though it would be difficult for anyone, the author included, to reproduce precisely many of the early simulation results (though snapshots do exist of the code in various stages of development along with the “worldfiles” that generated all of the organisms discussed here), it should be quite straightforward for anyone to reproduce the more significant and interesting later results, since both the source code for PW and the worldfiles that define the starting conditions for these simulations have been made freely available (see Future Directions below). (Many hours of videotapes and sometimes detailed, sometimes sparse, scientific journals also exist for most of these simulations.)

11. Discussion

Real benefits have already begun to accrue from the studies of artificial neural systems. Meanwhile, the study of artificial evolution—genetic algorithms—is yielding insights into problems of optimization, and into the dynamics of natural selection. One form of the study of Artificial Life is the perhaps obvious combination of these two fields of research. Adding computer graphics visualization techniques yields the basic substrate of PolyWorld.

One of the primary goals set out for PW has already been met: the evolution of complex emergent behaviors from only the simple suite of primitive behaviors built into the organisms of PW, their sense mechanisms, and the action of natural selection on their neural systems. These recognizable behavioral strategies from real living organisms, such as “fleeing”, “fighting back”, “grazing”, “foraging”, “following”, and “flocking”, are purely emergent in the PW environment. And built as

they are from simple, known primitive behaviors, in response to simple, understandable ecological pressures, they may be able to remove a little bit of the mystery, if not the wonder, at the evolution of such behaviors in natural organisms.

The simple but effective strategies evolved by organisms in the earlier, simpler simulations may be valuable as sort of “null hypotheses” about certain forms of animal behavior. In particular, aggregation and wall-following amongst these simple organisms occurs without need for elaborate behavioral strategies. It is sufficient that corners and walls obstruct the simple trajectories of limited motor skills. Yet if enough organisms occupy these locational niches, it becomes a behavioral niche as well, by providing readily available mates, and an easily achieved form of behavioral isolation.

It is, perhaps, easier to contemplate and understand these behaviors in the simulated organisms of PW than it is in natural organisms, precisely because they are simulated. The blessing and the curse of Artificial Life is that it is much more difficult for humans to anthropomorphize (zoomorphize? biomorphize?) these organisms in a machine than it is natural organisms. This frees us from prejudices and preconceptions when observing and analyzing the behaviors of artificial organisms, yet the most highly motivated of ALife researchers is going to find it difficult to look at an artificial organism and declare it unequivocally alive.

As more and more sophisticated computational models of living systems are developed, it will be only natural to ask whether they are in fact *really* alive. To answer this, however, requires a resolution to probably the greatest unanswered question posed and addressed by the study of Artificial Life: “What is life?” Farmer and Belin [16] offer an analogous question for consideration: “If we voyage to another planet, how will we know whether or not life is present?” One might also ask: If we “voyage” to an artificial world, how will we know whether or not life is present? In a tentative first step towards trying to answer such a question, Farmer and Belin offer a set of “properties that we associate with life”. Here is a brief analysis of how well the organisms of PolyWorld meet these criteria:

- “*Life is a pattern in spacetime*, rather than a specific material object.”

By this, Farmer and Belin mean to point out that even a specific living organism is really the process (that persists), rather than the nutrients or specific chemical constituents in which that process is embedded (which typically do *not* persist). Inherent in this observation is that the process may very well be taking place in something quite different than the traditional hydro-carbon chains associated with biological life (BLife). PW organisms are indeed patterns in a computer, rather than any traditional substrate. And as with BLife organisms, in order to persist they extract nutrients from the environment which, while necessary for life, are not the *process* of life. So PW organisms neither extend nor violate this first condition.

- “*Self-reproduction*.”

PW organisms certainly reproduce within the context of their world. The initiation of their act of reproduction is, akin to higher level BLife organisms, a result of the processing of a moderately complex nervous system—a “volitional”, behavioral action. The method of their reproduction is dependent on the software genetics built into PW—the pseudo-physics and -biology of PW. This pseudo-biology is not itself an emergent property, but an assumed capability. Whether this, therefore, represents *self*-reproduction or not may be arguable.

Some might argue that the step-by-step process of reproduction itself must be wholly contained within the simulated physics of an artificial world, perhaps in the manner of cellular automata models, in order to qualify as *self*-reproduction, and that an assumed capability and corresponding software mechanism for reproduction is too high a level of abstraction to support *real* artificial life (though one might counter argue that the assumed, software physics of such cellular automata systems are just as *ad hoc* a construction as these software genetics, just as *unreal*, and that even if a model of real molecular-level chemical interactions were used as the substrate, the physics underlying those interactions were *ad hoc*, and so on *ad infinitum*, until we reach levels of sub-atomic physics that cannot be modeled, since they are not known).

On the other hand, all ALife systems for the foreseeable future *must* make certain assumptions, and select levels of abstraction—levels of organization and complexity—at which they will develop their models. As Chris Langton puts it, they all must write an I.O.U. at some level of detail. The genetic code of PW’s organisms is, by design, in the system’s software architecture, and the mechanism for combining—crossing-over and mutating—those genetic codes, and then interpreting them as a new organism is also embedded in the software that

defines PW's physics and biology. It is not obvious that organisms that evolve to exploit these software reproduction mechanisms are any less *alive* than organisms that continue to evolve characteristics and behaviors which, though represented there, are otherwise unrelated to the biochemical mechanisms for the reproduction of DNA. (A longer-necked giraffe is selected for not because of the way its genes are copied, but because its longer neck offers a behavioral advantage; a more intelligent ape descendent is selected for that intelligence, not for the method of its genetic coding. Though a non-evolving genetic code does significantly limit the open-endedness of a simulation, only for organisms sufficiently primitive that the method of storing and reproducing their genetic code is still being explored and optimized are these characteristics an undeniably essential aspect.) Nor is it obvious that a PW organism's dependence on these software reproductive mechanisms is any less lifelike than the dependence that an organism in Ray's Tierra has on a built-in copy instruction, or the dependence that a software virus has on a computer system's built-in duplication functions, or a biological virus's dependence on existing mechanisms for copying its host's DNA.

Finally, a simulator may also be thought of as both the world context it creates *and* the underlying software. Even though the software was hand-crafted, rather than evolved, in a very real sense it represents the fundamental physics and, possibly, the low-level biological mechanisms of the simulated ecology. If a simulated organism is thought of as the combination of its simulated physiology, its simulated neurophysiology, and its real software, then even the high-level, abstracted form of reproduction employed in PW might be thought of as *self-reproduction*. Whether this violates the spirit of Farmer and Belin's criterion or not is unclear.

- "*Information storage of a self-representation.*"

PW organisms use an analog of the same storage mechanism Farmer and Belin mention for natural organisms: their genetic representation. Here again, that representation is stored and interpreted within PW's software context, not directly within the resulting simulated world. The representation itself cannot evolve within the current PW framework (though the author has considered a number of alternative schemes for supporting this), and is not itself an emergent property. These limitations do not violate Farmer and Belin's stated criterion, but do indicate some constraints imposed on the system by its chosen level of abstraction.

- "*A metabolism.*"

A PW organism's metabolism effectively converts food found in the environment into the energy it needs to carry out its internal processes and behavioral activities, just as is the case in natural organisms. In PW's current biochemistry, there is only one type of nutrient required to sustain life and permit reproduction of its organisms (though this needn't be the case in later versions of PW). In contrast to Holland's ECHO, for example, the single food type in PW is converted directly to an energy that is available to all of an organism's behavioral systems, rather than accumulating in a reservoir of components until all the constituent elements are available to support reproduction; given the existence of only a single nutrient type in PW, there is little point in treating it otherwise. But as with ECHO, sufficient quantities of this nutrient must be available in order to support reproduction. And in PW, the energy derived from this nutrient is then expended through every action, as well as the neural processing, of the organism; and the effects of one of those actions—fighting—is directly scaled by this available energy. So in some ways the metabolism of PW's organisms is simpler than that in ECHO, but in some ways it is more complex. In any event, the metabolism in PW organisms is certainly *much* simpler than the metabolism in biological organisms, but if the basic functionality is the same, does the complexity of the underlying process matter?

- "*Functional interactions with the environment.*"

PW organisms certainly do interact with their environment. Besides eating food and expending the resultant available energy, they interact extensively with other organisms in the environment. In fact, as in BLife, the other organisms comprise the most important element of that environment. The more sophisticated organisms in PW respond behaviorally to changes in the environment, and such responses are purely under the control of the organism.

- "*Interdependence of parts.*"

Following Farmer and Belin's reasoning, PW organisms can and would die were they somehow separated from their internal energy store. And severing an organism's brain in two would not produce two organisms with behavior anything like the original. As was discussed in the Results section above, the most sophisticated behaviors in PW only emerged once the

arbitrary neural architectures were implemented. The particular clusters of neurons in an organism's brain and the precise pattern of connections between those clusters *define* that organism's behaviors; altering those clusters and their connections produces a different (and not necessarily viable) organism. Stepping outside the bounds of the simulation, they would also "die" if their various procedures and data were destroyed or isolated. In either case, half an organism is no longer that organism, if it is any organism at all.

- "Stability under perturbations."

PW organisms can survive small changes to their environment. Indeed, whole species have reemerged in entirely different simulations. Again stepping outside the simulation, whole species have emerged with and without any of a variety of errors in the code.

- "The ability to evolve."

PW organisms clearly can and do evolve. There are undoubtedly limits to their evolution; e.g., they could not possibly evolve a sense of smell without programmer intervention. However, all natural organisms we know of have limits to their evolutionary capabilities: It is highly unlikely that humans could evolve a steel appendage; if Einstein is correct, it is absolutely impossible for them to evolve a method of personal locomotion that would exceed the speed of light. All organisms, natural or artificial, are bound by the physics of their universe. Similar to the question about metabolism, does the complexity of the underlying physics matter?

So, with the above caveats, questions, arguments and counter arguments, it would appear that the organisms of PolyWorld come surprisingly close to fulfilling Farmer and Belin's set of criteria; indeed, they *may* do so entirely. Farmer and Belin probably would not argue with the conclusion that we need to further refine our constraints on the definition of life. It is unknown which side they might argue of the knowingly contentious statement that we may already need to welcome a new *genus* to the world.

This issue of just what really defines life—what really is and isn't alive—will continue to be both a driving force and a thorn in the side of the field of ALife. Some, including Chris Langton [28], the father of the field, would argue that perhaps all of Farmer and Belin's criteria must themselves be self-organized. It might also be argued that *none* of the above criteria are especially necessary or appropriate: Consider a hypothetical, Turing-certified, artificially-intelligent computer. It might not be able to reproduce itself. It might not have a complete self-representation. It's metabolism would be no more complicated than that of a PW organism, "eating" nothing but electricity. It may not evolve. Yet if you could discuss Tolstoy and Terminator with it, if it understood the concepts of *noblesse oblige* and the gentleman farmer, and if it, perhaps, shared some of your musical tastes... could you fail to consider it *alive*? But on the other hand, perhaps this argument unnecessarily conflates *intelligence* and *life*. Or, more damningly, perhaps, like Searle's famous Chinese Room argument, it postulates an impossibility to make its (therefore invalid) point. Indeed, intelligence may only be achievable through a process of evolution—through life, natural or artificial. (Even if one could copy every nuance of a human being's nervous system into a computer, thus sidestepping computational evolution, the thing being copied is itself the product of an evolutionary process.)

Tom Ray [40] has suggested that perhaps all of ALife research is really just modeling *aspects* of life, rather than actually creating "capital-L" *Life*, and each effort, to a large degree, echoes the interests and biases of the individual researcher. Chris Langton's Loops met his criterion at the time—self-reproduction, even though their reproduction is that of a precise crystalline form which cannot evolve. Ray's Tierra organisms met his criteria too—self-reproduction and evolution, though their behaviors are limited exclusively to the reproductive process. PolyWorld's organisms met their creator's criteria—reproduction, evolution, and ethological-level behaviors, though the intelligence of these organisms is limited, at best, to that found on the lowest rungs of BLife organisms. Perhaps extensions to PolyWorld, or the next researcher's ALife environment will successfully evolve more intelligent, more obviously *alive* organisms. Or perhaps Langton's Swarm work—an eco-simulator based on a multi-plane CA—will reproduce the higher-level behaviors of PolyWorld, but build them on the lower-level physics of CA's. Curiously, man-made life seems almost easy to model, almost impossible to accept as truly alive.

Ultimately, the resolution to this question of real life in artificial organisms may have to be based on a consensus, as with Turing's famous test for artificial intelligence. Perhaps in this case, however, the consensus of a knowledgeable and informed jury is needed, rather than that of Turing's unspecified, presumably average group of individuals. As with the debate about the "aliveness" of

natural viruses being properly resident with biologists, the question of “aliveness” in artificial organisms is probably best argued by a combination of computer-aware biologists and biology-aware computer scientists.

Still, researchers in this field seem to have a feeling that if you have a box of stuff, there really ought to be a way to tell if something is alive in there. The best approach may yet be an information-theoretic one, harkening back to the suggestions of Von Neumann and Schrödinger that the crucial, defining aspect of living organisms is that they are information processors—information-rich islands in a sea of background information. If, say, the ability to predict the chemical composition or electrical charge or some other measure of state, in either the current location over time, or in adjacent locations at the same time, were measured—at some scale—in a bounded volume, then perhaps we could use such information measures to make an assessment of the amount, or degree, of life contained therein. Or perhaps if such a determination were made at a *variety* of scales within the contained region, living things might stand out clearly from the non-living background by the manner in which their information content scales. But for now, no such quantitative measures (or even particularly cogent theories) exist, so intelligent conjecture, argument, and opinion will have to suffice.

12. Future Directions

The various species and behaviors that have emerged in the different simulations suggest that PW may be a rich enough simulation environment to pursue further evolutionary studies. In particular, a way of sort of “benchmarking” PW—the way one compares the results of a computational fluid dynamics code to known analytical solutions for flow over a flat plate or a cylinder, or measured flows over an airfoil in a windtunnel—may be possible in the form of optimal foraging strategies as studied in the field of behavioral ecology. A simple, canonical foraging experiment has been defined and analyzed, and some preliminary simulations run in PW. Agreement or disagreement with the analytical model should be examined and understood.

The neural architectures that provide the most useful survival strategies should be analyzed and understood. It would also be fairly straightforward to encode an entire range of learning algorithms in the genes of the organisms in PW, and attempt to evolve the most effective learning algorithm, rather than assuming it to be Hebbian. (Some consideration has even been given to the possibility of having the fundamental genetic representation of information—the genetic code—evolve.) At least it might be worthwhile implementing cluster-to-cluster initial connection strengths, initial connection strength variances, and maximum connection strengths, to begin to hint at distinct cell types. Or it may be more worthwhile to jump directly to a more sophisticated cell model, capable of capturing the actual temporal dynamics of spike trains rather than average firing rates.

Though the statistical approach to the specification of neural architectures currently employed in PW can, to a certain extent, finesse the need for an ontogenetic, developmental process, this is felt to be one of the most potentially valuable directions for future ALife work. A richer, more biologically-motivated developmental process might provide as significant an improvement in the process of searching the space of evolutionarily useful neural architectures as the current scheme did over the fixed, fully-recurrent networks. An ontogenetic process for the organisms’ neurophysiologies (and physiologies) might serve both to smooth the fitness landscape and, occasionally, to introduce useful cliffs in that landscape, as it is conjectured to do in natural evolving systems. And the converse may also be true; PW may be a very effective testbed for alternative ontogenetic theories and algorithmic models of development.

More environmental interactions should be supported, including the ability for the organisms to pick up, carry, and drop pieces of food, and perhaps even pieces of barrier material. This should yield useful reasons for organisms to cooperate, other than simply to reproduce.

The simple metabolisms dictated by the current use of a single type of nutrient could readily be replaced with a more complex biochemistry. Multiple nutrients distributed amongst multiple food types would make for a considerably richer environment. (Even the potentially substantial impact of ordered, rather than random, food growth has yet to be explored.) It might even be feasible to incorporate a simple set of internally consistent chemical reactions, with appropriate energetics, catalytic relationships, and so on.

Though not discussed in the earlier parts of the paper, the gross energetics of the system have been observed to be crucial to the evolution of successful survival strategies. Mirroring the differences between energy-rich tropical zones and energy-starved polar zones in our one known, natural

ecosystem, artificial life flourishes in energy-rich simulations, and languishes in energy-starved simulations. Perhaps someday it may be possible to make useful predictions about viable ranges of energy flux for natural systems from artificial ecologies like PW.

A quantitative assessment of the degree to which the isolation of populations affects speciation may be possible with PW. Some tentative first steps have already been taken in this direction, though questions remain about the most appropriate comparisons to make and the appropriate times at which to make these comparisons. This coupled with the problems associated with assuring the emergence of an SBS in every population, and the simple magnitude of processing time required to perform the simulations has delayed a complete series of experiments of this nature.

There are thousands of other interesting experiments that one might perform with this system, including: Monitoring brain size in otherwise stable populations, such as the "dervishes"... are smaller and smaller nervous systems actually being selected for? Monitoring the frequency and magnitude of attacks on other organisms as a function of their genetic (dis)similarity. Monitoring the amount of energy given to offspring in a single species... is there any indication of an asymmetric split into different relative contributions? Hand-tailoring a good neural architecture or two and seeding the world with these engineered organisms. Providing multiple internal, neural time-cycles per external, action time-cycle. Evolving three completely independent domains of organisms, with barriers in place, and then removing the barriers to observe the interspecies dynamics. It may even be possible to model the entire population of Orca whales that frequent the waters around Vancouver, and look for an evolutionary split into pods that travel little and eat essentially stationary food sources versus pods that travel widely and feed on fish, a very mobile food source. And on and on. In hopes that others may find PolyWorld to be a useful tool for exploring these kinds of questions, it has been made available via ftp from ftp.apple.com (130.43.2.3) in /pub/polyworld. Complete source code and some sample "worldfiles" are provided.

In a more fanciful, and perhaps more visionary vein, it is hoped that, someday, one of the organisms in PolyWorld that demonstrates all the survival behaviors observed to date, plus a few others, could be transferred from its original environment to, say, a maze world, and become the subject of some classical conditioning experiments. Klopff's [24,25] success at demonstrating over 15 classical conditioning phenomena in a single neuron using differential Hebbian learning (he called it "drive-reinforcement" learning), strongly suggests that such phenomena should be demonstrable in PolyWorld's organisms.

And then, of course, there is simply "more, bigger, and longer": more organisms, with bigger neural systems, evolving longer. As a *gedanken* experiment, consider just how *much* "more, bigger, and longer" might be useful: The current 3×10^2 organisms, 3×10^2 neurons, and 10^3 generations (approximately), could be expanded to 10^6 organisms, neurons, and generations, through an increase in compute power of about 10^{10} . (Though this sounds like a tremendous increase to ask for, consider that the current simulation is running on a single, scalar workstation processor, not a vectorized, massively parallel processor, then extend today's trends in compute power, and this ceases to be such a daunting request. Also, the compute power needed to model quite complex organisms may be significantly less than this due to the greatly reduced motor and autonomic nervous systems that would be required by artificial organisms.) By coincidence, it turns out that this is a fairly reasonable amount of compute power with which to consider modeling a complete human brain—basically devoting one of today's fast computers to every neuron—but no one understands how to actually construct such an artificial brain. However, this same amount of compute power might be used to evolve the equivalent of several new species of computational lab rats every week... and *this* is how: by combining evolution, neural systems, and ecological dynamics.

If there is any question about why one would wish to pursue these research directions, it is always possible to point to the benefits to be derived in the evolutionary, ecological, biological, ethological, and even computer science fields. But it may also turn out to be the only "right" way to approach machine intelligence.

One view of intelligence is as an evolved, adaptive response to a variable environment, that due to historical constraints and opportunism on the part of nature happens to be based upon neuronal cells. One might further recognize that intelligence is actually a near-continuum—a spectrum from the simplest organisms to the most complex—rather than some singular event, unique to human beings. Then, by utilizing both the method—Natural Selection—and the tools—assemblies of neuronal cells—

used in the creation of natural intelligence, PolyWorld is an attempt to take the appropriate first steps toward modeling, understanding, and reproducing the phenomenon of intelligence.

For while one of the grand goals of science is certainly the development of a functioning human level (or greater) intelligence in the computer, it would be an only slightly less grand achievement to evolve a computational *Aplysia* that was fully knowable—fully instrumentable and, ultimately, fully understandable. And perhaps it is only through such an evolutionary approach that it will be possible to provide the important milestones and benchmarks—sea slug, rat, simian,...—that will let us know we are on the right scientific path toward that grander goal.

Acknowledgments

The author would like to thank Alan Kay and Ann Marion of Apple's Vivarium Program for their support and encouragement of this admittedly exotic research. He would also like to thank Richard Dawkins for early encouragement and suggestions, Steve Nowlan, Terry Sejnowski and Francis Crick for their interest and kind words, and Chris Langton, Luc Steels, and Tom Ray for some wonderfully thought provoking feedback and discussions. Finally and emphatically, he would like to thank his wife, Levi Thomas, for her support, understanding, and humor (especially for christening PolyWorld's organisms "C-monkeys").

References

1. Ackley, D., and M. Littman (1992), "Interactions between Learning and Evolution" In *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, and S. Rasmussen. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. X. Addison-Wesley, Redwood City, CA, 1992.
2. Belew, R. K., J. McInerney, and N. N. Schraudolph (1992), "Evolving Networks: Using the Genetic Algorithm with Connectionist Learning" In *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, and S. Rasmussen. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. X. Addison-Wesley, Redwood City, CA, 1992.
3. Braitenberg, V. (1984), *Vehicles: Experiments in Synthetic Psychology*. A Bradford Book, MIT Press, Cambridge, 1984.
4. Chalmers, D. (1991), "The Evolution of Learning: An Experiment in Genetic Connectionism" In *Connectionist Models, Proceedings of the 1990 Summer School*, edited by D. S. Touretzky, J. L. Elman, T. J. Sejnowski, G. E. Hinton, Morgan Kaufmann, San Mateo, CA, 1991.
5. Cliff, D. (1991), "The Computational Hoverfly; a Study in Computational Neuroethology" In *From Animals to Animats, Proceedings of the First International Conference on Simulation of Adaptive Behavior*, edited by J.-A. Meyer and S. Wilson. A Bradford Book, MIT Press, Cambridge and London, 1991.
6. Collins, R. J., and D. R. Jefferson (1992), "AntFarm: Towards Simulated Evolution" In *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, and S. Rasmussen. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. X. Addison-Wesley, Redwood City, CA, 1992.
7. Conrad, M., and M. Strizich (1985), "EVOLVE II: A Computer Model of an Evolving Ecosystem", *Biosystems* **17**, 245-258, 1985.
8. Conrad, M. (1987), "Computer Test Beds for Evolutionary Theory" Oral Presentation at Artificial Life I Conference, 1987.
9. Dawkins, R. (1976), *The Selfish Gene*. Oxford University Press, Oxford, 1976.
10. Dawkins, R. (1983), *The Extended Phenotype: The Gene as a Unit of Selection*. Oxford University Press, Oxford, 1983.
11. Dawkins, R. (1986), *The Blind Watchmaker*. W.W. Norton, New York, 1986.

12. de Boer, M. J. M., F. D. Fracchia, and P. Prusinkiewicz (1992), "Analysis and Simulation of the Development of Cellular Layers" In *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, and S. Rasmussen. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. X. Addison-Wesley, Redwood City, CA, 1992.
13. Dewdney, A. K. (1984), "Computer Recreations: In the Game Called Core War Hostile Programs Engage in a Battle of Bits", *Scientific American* **250(5)**, 14-22, May 1984.
14. Dewdney, A. K. (1984), "Computer Recreations: A Core War Bestiary of Viruses, Worms, and Other Threats to Computer Memories", *Scientific American* **252(3)**, 14-23, March 1985.
15. Dewdney, A. K. (1987), "A Program Called MICE Nibbles its Way to Victory at the First Core War Tournament", *Scientific American* **256(1)**, 14-20, January 1987.
16. Farmer, J. D., and A. d'A. Belin (1992), "Artificial Life: The Coming Evolution" In *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, and S. Rasmussen. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. X. Addison-Wesley, Redwood City, CA, 1992.
17. Harp, S., T. Samad, and A. Guha (1990), "Towards the Genetic Synthesis of Neural Networks" In *Proc. Third Intl. Conf. on Genetic Algorithms*, edited by J. D. Schaffer, Morgan Kaufmann, San Mateo, CA, 1990.
18. Hebb, D. O. (1949), *The Organization of Behavior*, John Wiley and Sons, Inc., New York, 1949.
19. Hillis, D. (1990), "Simulated Evolution and the Red Queen Hypothesis" Oral Presentation at Artificial Life II Conference, 1990.
20. Hillis, D. (1992), "Co-Evolving Parasites Improve Simulated Evolution as an Optimization Procedure" In *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, and S. Rasmussen. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. X. Addison-Wesley, Redwood City, CA, 1992.
21. Holland, J. (1990), "Echo: Explorations of Evolution in a Miniature World" Oral Presentation at Artificial Life II Conference, 1990.
22. Hugie, D. (1992), Behavioral Ecology Dept., Simon Fraser Univ., Vancouver, B.C., Canada, Personal Communication.
23. Jefferson, D., R. Collins, C. Cooper, M. Dyer, M. Flowers, R. Korf, C. Taylor, and A. Wang (1992), "Evolution as a Theme in Artificial Life: The Genesys/Tracker System" In *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, and S. Rasmussen. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. X. Addison-Wesley, Redwood City, CA, 1992.
24. Klopff, A. H. (1986), "A Drive-Reinforcement Model of Single Neuron Function: An Alternative to the Hebbian Neuronal Model" In *Neural Networks for Computer*, edited by J. S. Denker, AIP Conference Proceedings **151**, American Institute of Physics, New York, 1986.
25. Klopff, A. H. (1987), "A Neuronal Model of Classical Conditioning", AFWAL-TR-87-1139, Air Force Wright Aeronautical Laboratories, October, 1987.
26. Koza, J. R. (1992), "Genetic Evolution and Co-Evolution of Computer Programs" In *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, and S. Rasmussen. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. X. Addison-Wesley, Redwood City, CA, 1992.
27. Langton, C. G. (1989), ed. *Artificial Life*. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. VI. Addison-Wesley, Redwood City, CA, 1989.
28. Langton, C. G. (1993), Personal Communication.
29. Linsker, R. (1988), "Towards an Organizing Principle for a Layered Perceptual Network" In *Neural Information Processing Systems*, edited by D. Z. Anderson. American Institute of Physics, New York, 1988.
30. Linsker, R. (1988), "Self-Organization in a Perceptual Network", *Computer* **21(3)**, 105-117, March 1988.

31. Linsker, R. (1989), "An Application of the Principle of Maximum Information Preservation to Linear Systems" In *Advances in Neural Information Processing Systems 1*, edited by D. S. Touretzky. Morgan Kaufmann Publishers, San Mateo, CA, 1989.
32. Miller, G. F., and P. M. Todd (1991), "Exploring Adaptive Agency I: Theory and Methods for Simulating the Evolution of Learning" In *Connectionist Models, Proceedings of the 1990 Summer School*, edited by D. S. Touretzky, J. L. Elman, T. J. Sejnowski, G. E. Hinton, Morgan Kaufmann, San Mateo, CA, 1991.
33. Miller, S. M., and L. E. Orgel (1974), *The Origins of Life*. Prentice-Hall, Englewood Cliffs, NJ, 1974.
34. Nolfi, S., J. L. Elman, and D. Parisi (1990), "Learning and Evolution in Neural Networks", CRL Tech. Rep. 9019, Center for Research in Language, UCSD, La Jolla, CA, 1990.
35. Packard, N. (1989), "Intrinsic Adaptation in a Simple Model for Evolution" In *Artificial Life*, edited by C. Langton. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. VI. Addison-Wesley, Redwood City, CA, 1989.
36. Parisi, D., S. Nolfi, and F. Cecconi (1991), "Learning, Behavior, and Evolution", Tech. Rep. PCIA-91-14, Dept. of Cognitive Processes and Artificial Intelligence, Institute of Psychology, C.N.R. - Rome, June 1991. (To appear in *Proceedings of ECAL-91—First European Conference on Artificial Life, December 1991, Paris*).
37. Pearson, J. (1987), "Competitive/Cooperative Behavior of Neuronal Groups in Brain Function" Oral Presentation at Artificial Life I Conference, 1987. (And in Edelman, G. M. *Neural Darwinism: The Theory of Neuronal Group Selection*. Basic Books, New York, 1987.)
38. Rasmussen, S., C. Knudsen, R. Feldberg, and M. Hindsholm (1990), "The Coreworld: Emergence and Evolution of Cooperative Structures in a Computational Chemistry" In *Emergent Computation*, edited by Stephanie Forrest, North-Holland, Amsterdam, A Special Volume of *Physica D*, Vol. 42 (1990) Nos. 1-3.
39. Ray, T. S. (1992), "An Approach to the Synthesis of Life" In *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, and S. Rasmussen. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. X. Addison-Wesley, Redwood City, CA, 1992.
40. Ray, T. S. (1993), Personal Communication.
41. Renault, O., N. M. Thalmann, and D. Thalmann (1990), "A Vision-based Approach to Behavioural Animation", *J. of Visualization and Computer Animation* Vol. 1, 18-21.
42. Taylor, C. E., D. R. Jefferson, S. R. Turner, and S. R. Goldman (1989), "RAM: Artificial Life for the Exploration of Complex Biological Systems" In *Artificial Life*, edited by C. Langton. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. VI. Addison-Wesley, Redwood City, CA, 1989.
43. Todd, P. M., and G. F. Miller (1987), "A General Framework for the Evolution of Adaptive Simulated Creatures" Oral Presentation at Artificial Life I Conference, 1987.
44. Todd, P. M., and G. F. Miller (1991), "Exploring Adaptive Agency II: Simulating the Evolution of Associative Learning" In *From Animals to Animats, Proceedings of the First International Conference on Simulation of Adaptive Behavior*, edited by J.-A. Meyer and S. Wilson. A Bradford Book, MIT Press, Cambridge and London, 1991.
45. Travers, M. (1989), "Animal Construction Kits" In *Artificial Life*, edited by C. Langton. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. VI. Addison-Wesley, Redwood City, CA, 1989.
46. Walter, W. G. *The Living Brain*. W. W. Norton, New York, 1963.
47. Walter, W. G. (1950), "An Imitation of Life", *Scientific American*, **182**(5), 42-45, May 1950.
48. Walter, W. G. (1951), "A Machine that Learns", *Scientific American*, **185**(2), 60-63, August 1951.
49. Wharton, J., and B. Koball (1987), "A Test Vehicle for Braitenberg Control Algorithms" Oral Presentation at Artificial Life I Conference, 1987.




Figure 1. Screenshot of the PolyWorld ecological simulator populated by several related but distinct sub-species of “Dervishes” (see Section 10, Results: Speciation and Complex Emergent Behaviors). The largest panel shows a broad view of the world: the dark green ground plane, the brown, impassable barriers, the bright green pieces of food, and the multicolored organisms. Just above this oblique world view are four graphs of various simulation parameters. Above these, at the top of the figure, are many small views of the world drawn from the point of view of each of the organisms in the world; these are the images seen by those organisms. At the top right are a few numerical statistics describing the simulation. And in the bottom right pane is a zoomed-in, close-up view of the current “fittest” organism.




Figure 2. Three sample neural architectures evolved in PolyWorld. Gray-scale is used to denote neural activation (between 0.0 and 1.0) and synaptic efficacy (between $-\text{maxEfficacy}$ and $+\text{maxEfficacy}$). At the very bottom of the grid is the color vision buffer. The neural activations at the beginning of this time step are shown in a horizontal row just above the color vision, along with the red, green, and blue input neuron activation levels, and the energy and random input neuron activation levels. White frames are drawn around each neuronal group, except for the vision neurons which are framed in their corresponding color. Black frames are drawn around each synapse, hence the unframed areas are regions of null connectivity. Synapses that appear brighter than the neutral gray background are excitatory; those that appear darker than the background are inhibitory. The leftmost vertical bar shows neuronal biases. The non-input neural activations at the end of this time step are shown in the adjacent vertical bar; again, neuronal groups are framed in white. Hence the diagram may be read as an incomplete crossbar connecting the neuronal states at the beginning of the time step (horizontally) to those at the end of the time step (vertically), through the various synaptic connections.