

Gaia: An Artificial Life Environment for Ecological Systems Simulation

Nuno Gracias²
ngracias@isr.ist.utl.pt

Henrique Pereira^{1,2}
hpereira@isr.ist.utl.pt

José Allen Lima²
jalima@isr.ist.utl.pt

Agostinho Rosa²
acrosa@isr.ist.utl.pt

¹EST - Politechnical Institute of Setubal

²Systems and Robotics Institute

Instituto Superior Técnico

Av. Rovisco Pais 1096 Lisboa Codex

Portugal

Abstract

This paper presents an ecology simulator for the study of certain aspects of ecology and biology, such as learning, evolution and population dynamics.

The simulator is an artificial world, where two kinds of species can evolve: autotrophs and heterotrophs. Heterotrophic individuals, or critters, are capable of moving, eating, fighting and mating. They have a simple nervous system, a neural network, with a retina input. Associative Hebbian learning is used in the modification of synapses. Nervous system structure and physiological characteristics are coded in the critter's genome. Autotrophs are static. They are born and grow according to a definable geographic distribution and rate.

Simulations were carried out to study learning and behaviour evolution, in an approach as close as possible to biological reality. It was found that critter learning was essentially phylogenetic, i.e., Hebbian learning was limited to develop genetically defined connections, and did not perform any significant correlation between inputs and outputs. This suggests that, for animals with very simple nervous systems, behaviours evolve mainly through genetics, as opposed to a possible emergent scheme of reinforcement learning.

Interspecific resource competition was also studied. It was found that if two similar species were sowed in the world, one of them won the competition. Several explanations for this phenomena are presented. A strong relation was found, between the geographic distribution of autotrophs and heterotrophs. Limit cycles of the two populations were observed.

1. Introduction

Over the last decades, two apparently independent fields of Biology, Evolution and Neurobiology became increasingly closer. While Evolution studies the appearance and evolution of species, on a time scale of hundreds of thousand of years, Neurobiology studies the learning process at the cellular level, on a time scale from seconds to a few years, i.e. an individual's life time.

However, these two fields are unquestionably close. Without Evolution the nervous system would not appear. Without nervous system, evolution would not have taken

the path it took. Nonetheless, neither is evolution at the level of the nervous system well known, nor is the role played by the nervous system in natural evolution. And although natural selection, the basic mechanism of evolution, is nowadays known, we only have a few clues about how the structure and the physiology of the nervous system is transmitted between generations.

The main goal of this work was to test the present knowledge about learning and evolution. "Can the known biological mechanisms explain the evolution of organisms with a very simple nervous system?", was the question we tried to answer.

This work was inspired on a similar one, *PolyWorld* (PW) by Larry Yaeger [1994]. Yaeger created a two dimensional artificial world to simulate a simple ecosystem. He inoculated it with growing food cells and *critters*. The critters had a neural network to simulate the nervous system, and were able to see, to move, to eat, to fight, and most important, to reproduce. After some simulation time, the critters were seeking their own food and assuring the species continuity by mating.

The reproduction of Yaeger's results was one of our initial aims, but in our ecological simulator some features not present in PW were added, in order to study population dynamics and autotroph-heterotroph relationships. We gave the name *Gaia* to our ecological simulator. *Gaia* is an ancient Greek goddess, representing Earth as a whole organism [Margulis & Sagan, 1986]. This does not necessarily mean that we totally agree with Lovelock's Gaia Hypothesis, but we think that some self regulation features should appear in our artificial world.

2. Background

In the last decades several computer-based ecology systems have been developed. The fundamental principles of evolution dynamics were explored in systems built by Conrad [1985] and Packard [1989]. Strategy evolution on a group of artificial ants was studied by Collins and Jefferson [1992] in *AntFarm*, using Finite State Automata and Neural Networks. Population dynamics as the result of interactions between individuals was studied by Taylor [Taylor *et al*, 1989]. New mechanisms on how evolution can guide learning were studied by [Ackley & Littman, 1992]. Todd & Miller [1991] explored evolutionary selection on learning algorithms, for controlling organisms with simple vision and scent sensory devices.

An example of computational ecology is *Tierra*, developed by Ray [1992]. Based on evolutionary programming. *Tierra* is a virtual computer adapted for synthetic life, with its own set of instructions and operating system. Ray's approach was to try to recreate the complexity and diversity of the Cambrian explosion, starting with hand-coded evolving organisms.

There are few references to ecological systems, intended to come close to their biological counterparts. PolyWorld was originally directed towards the evolution of neural architectures for complex behaviours, but the biologically based reproduction and evolutionary mechanisms used made it also a tool for studying ecology and evolutionary biology [Yaeger 1994]. Lindgren *et al* [1993] studied food webs resulting from artificial ecologies with external resource flows, and modelled the interactions between species using game theory. Johnson [1994] created an ecological community with organism body-size constraints and studied the resulting community structure and dynamics.

3. Approach

We consider that an ecology simulator should work at four levels, similarly to a natural ecosystem:

1. Apply **physical laws** to organisms, like physical restrictions to movement and energy use.
2. Apply **biological laws** to organisms: select in an implicit or explicit way the most fit organisms, allow interactions between individuals (reproduction or predation, etc.)
3. Implement the organisms **nervous system** and its learning mechanisms.
4. Implement **genome** decoding mechanisms, to simulate embryo development, as well as genome recombination mechanisms for reproduction.

These artificial ecosystem laws, mechanisms and structure should be as close as possible to biological reality. *Gaia* tries to satisfy these requirements.

3.1 Overview

Gaia is a two dimensional rectangular world delimited by walls. It can also have other obstacles spread on the world surface, that may work as an allopatric speciation mechanism [Curtis & Barnes, 1994]. *Gaia* uses a meaningful colour code, so that inhabitants that are able to see can easily distinguish its elements. Obstacles and walls are blue. Autotrophs are red and heterotrophs are green.

Two type of species can be sowed into *Gaia*: **autotroph** and **heterotroph**. Autotroph individuals work as **food cells** for heterotroph beings and are the only energy source on *Gaia*. Autotrophs are red coloured.

Heterotrophs, or **critters**, play the most important role in this artificial world, and logically, have the most complex structure. They have a nervous system and carry a genome which defines several structural and physiological characteristics. When a critter dies its remaining energy is converted into a food cell. This is not just a energy conservation principle, but also a way to allow the emergence of

predator-prey and cannibalism relationships. Killing another critter is a way of getting food. Critters have a triangular shape, a genetically defined green colour intensity and a genetically defined size. Their main **functional blocks** are:

- **Locomotion**: this block controls angular and linear speed.
- **Vision**: a linear colour sensor array which is the only source of information from the outside world.
- **Genome**: the genome has five chromosomes, it is decoded at *birth* according to an ontogenetic developmental program and is recombined in reproduction.
- **Nervous system**: the nervous system is a neural network, with sensory neurons in a retina, and with motor neurons deciding what kind of actions the critter will engage in; synaptic efficacies are modified through associative Hebbian learning.

Critters also have internal variables, the **state variables** and the **will variables**. The state variables are:

- **Health**: represents the degree of damage, decreasing in each collision or fight.
- **Food Value**: represents the amount of energy reserves that a critter has.
- **Can Fight, Can Mate**: expresses if a critter is ready to fight or mate; this implements a time latency in mating and fighting interactions; a predefined amount of time is required until one can mate (fight) again; new born critters can not mate until they reach maturity.

The will variables are the outputs of the critters' nervous system motor neurons:

- **Want to Eat, Want to Fight, Want to Mate**: express critter desire of eating, fighting and mating; when a critter wants to mate, it changes its colour by adding a red component to its green natural colour.
- **Moving, Turning**: commands the Locomotion block.

The simulator was developed for a IBM/ AT based personal computer (PC), running the MS-Windows operating system. For simulations with 40 critters, each time step takes about 5 seconds on a 100 MHz Pentium based PC.

3.2 *Gaia* Physics

When designing *Gaia* we adapted Newton physical laws to a simple two dimensional world. But the energy conservation principle is not verified because *Gaia* is not a closed system. Autotrophs get their energy from a virtual energy source¹.

Mass M , was defined as being proportional to the squared size. The **moment of inertia**, J , was defined as being proportional to size. Finally **kinetic energy** was defined by

¹ In *Gaia* only heterotroph beings dissipate energy, through collision, locomotion and metabolism.

$$E_c = k_1 \cdot M \cdot v^2 + k_2 \cdot J \cdot \left(\frac{d\theta}{dt} \right)^2$$

This formula is used for computing critter energy loss through collision and movement.

3.3 Autotroph beings

Autotrophs do not have a nervous system nor a genome. They are born at a constant birth rate (until saturation) and with a gaussian geographical distribution, with definable mean and variance. Because *Gaia* heterotroph beings are monophagous predators, this acts as a **limiting distribution factor** [Krebs, 1994].

An autotroph energy value increases with its age by the expression :

$$(FOOD_FIN_VALUE - FOOD_INI_VALUE) \cdot \left(1 - e^{-\frac{age}{FOOD_TIME_CONST}} \right) + FOOD_INI_VALUE$$

3.4 Heterotroph beings

Several similar species of heterotrophs can be sowed into *Gaia*. At the beginning they all share the same genetic pool, i.e., they are phenotypically and genetically identical except for a random variation. However members of different species do not interbreed, i.e., species are genetically isolated. This allows the study of **interspecific competition** and **sympatric speciation** [Krebs 1994].

3.4.1 Vision

In *Gaia*, as in PW, vision is used as the only external sensory mechanism and the most important for the guidance of the basic behaviours required for survival.

The approach for the development of a vision system in *Gaia* was to make it as close as possible to a biological retina on a functional level, while taking great care with the computational expenses. Each critter has a simulated retina which feeds directly to the neural network. All individuals have equal retinas and therefore equal visual resolution. Their visual evolution is carried out by an evolving, critter-specific number of vision neurons that sample the retina cells, and feed the information to the internal layers.

As stated before, *Gaia* is a 2-dimensional world. All its elements can be pictured as bidimensional planar shapes. For this reason, it is not reasonable to use a vision mechanism more complex than simple linear vision, where the rays of incoming light are projected along a simple line.

In our simulation a pin-hole camera model was adapted for circular retinas, which allows wide angle vision. Although not very realistic, the critters can sense light over 360 degrees. The images on the retinas are created using raytracing. Each light ray that hits the sensing cells is traced to its origin in order to determine its intensity and colour.

One of the biologically inspired feature is the critters ability to adjust their vision according to their will. They can concentrate or expand the raytracing scanning. This feature is a simple implementation of foveal vision and can be roughly compared with the zooming of an optical cam-

era. This provides a way for the critters to minimise the effects of low resolution retinas. By concentrating the scan lines on the front of their “heads” they can be more sensible to the presence of objects in the area they are heading to, and can more easily identify them. On the other hand, if they distribute them equally around their bodies, they have uniform vision and can detect any incoming attacker. The vision “spread” can be changed dynamically during the critters lifetime². This could be used in a cyclic manner although such a complex behaviour is highly unlikely to emerge.

A good distance perception of the surrounding objects can be an important competitive advantage for a critter. With it, one can guide itself efficiently, and be able to collect food near the obstacles without collision penalties. However, low resolution retinas do not provide means to accurately measure the angular size of the nearby objects, which is the first step towards a correct evaluation of its distance. Therefore an additional source of distance information was added to the vision system, which consists on the attenuation of the light rays with increasing distances. Dim stimuli can easily be associated with distant objects, and a growing stimulus with incoming critters or obstacles. This attenuation makes *Gaia* a foggy world, where vision decays with distance.

3.4.2 Nervous system

Gaia critters’ nervous system is similar to the one used in Yaeger’s PW. It is a **neural network**, with a genetic encoded structure³. Genetic representation of the neural network is based on Harp *et al* [1990] biological inspired schema, which seems to capture much of the architectural regularity of nervous systems in vertebrates.

The neural network is organized by **areas** (neuron layers) with an unidimensional structure. It has five **sensory areas** (input layers), three starting from the retina (one dedicated to each vision colour component), one receiving the value of critter’s energy, and another (optional) receiving a random value. It has six **motor areas** (output layers), one for each possible behaviour: moving, turning left, turning right, eating, mating and fighting. Finally, it has a genetically defined number of **internal areas**. Each **internal area** can be **inhibitory** or **excitatory**, meaning that synapses starting from that area are excitatory (synaptic efficacy greater than zero) or inhibitory (synaptic efficacy less than zero). Each area can have **projections** (a group of synapses) to any other area (the network can be fully recurrent), including to itself.

Ontogenetic developmental program

To translate the genome representation to a “living” neural network, *Gaia* uses an ontogenetic developmental program that is applied to a critter at the time of birth.

The genes that make up the neural network representa-

² At this moment this feature is disabled.

³ Structure is not modified during the critter lifetime, only synaptic efficacies, and these are not genetically coded. In *Gaia* there is no inheritance of acquired characteristics (Lamarckism).

tion are: number of internal neuron areas ($NINA$), number of neurons in each internal area (NN_i), initial bias in each neuron area (IB_i), bias learning rate for each neuron area (BLR_i), connection density between each pair of neuron areas (CD_{ij}), topological distortion between each pair of neuron areas (TD_{ij}), learning rate between each pair of areas (LR_{ij}).

Connection density determines the number of connections between two areas in the following way:

$$NC_{i,j} = CD_{i,j} \times NN_i \times NN_j$$

where NC_{ij} is the number of synapses between the source area i and the target area j .

Topological distortion determines the degree of disorder of a projection, i.e., the average receptive radius of the connections. That is, for a TD of 100%, synapses are mapped in a completely random fashion. For a TD of 0%, synapses connect contiguous stretches of neurons.

The ontogenetic developmental program can be stated as follows. We start by building a matrix of the number of synapses between each area. This matrix is used in the verification of network blueprint abnormalities. There are two kind of abnormalities [Harp *et al*, 1990]:

- fatal abnormalities: there is no pathway from input to output
- minor abnormalities: some areas don't have input or output projections

Critters with fatal abnormalities are culled out from the population, and in critters with minor abnormalities the abnormal areas are ignored in the network instantiation. In this case, genetic information about abnormal areas is preserved, acting like "introns" in the genome⁴.

The following step in the developmental program is the synaptic mapping. For each source area and for each target area the projections are built according to the TD parameter. Finally the synaptic efficacies are initialized between a definable minimum and maximum value, and neuron bias is initialized with the values specified by the critter genome.

Network dynamics and learning

During the critter lifetime, and at each pattern presentation, activations are calculated as follows, if \mathbf{j} is a sensory neuron

$$a_i^{k+1} = F_{sensory} \left(\sum_{j \in \{\text{afferent neurons}\}} a_j^k \cdot w_{ij} + \xi_i \right)$$

otherwise

$$a_i^{k+1} = F \left(\sum_{j \in \{\text{afferent neurons}\}} a_j^k \cdot w_{ij} \right)$$

where a_i^k is the activation of neuron i at local time step k , ξ_i is the input at unit i (matched retina value, health value

⁴ It should be noted that these introns can result from an undesired mutation. Their preservation provides a way for the later recovering of the original information.

or random value), w_{ij} is the connection weight⁵ between neurons j and i , $F_{sensory}$ and F are the following functions⁶:

$$F_{sensory}(x) = \frac{\tan^{-1}(x)}{\pi/2} \quad F(x) = \frac{1}{1 + e^{-\alpha \cdot x}}$$

where α is a specifiable logistic slope (typically 0.5).

Activations are calculated by areas, going from the sensory areas through the motor areas. At each area, activations are updated synchronously. Each time a new area is calculated a new time step is considered ($\mathbf{k}=\mathbf{k}+1$). That way, we assure that at each global time step⁷ (at each pattern presentation) the signal presented at network input propagates to the output. It should be noted that for each area, the activations a_i of afferent connections coming from the above areas or from the area itself, are the activations of the previous global time step, which gives "memory" to the network.

At the end of each global time step \mathbf{t} , the synaptic efficacies are updated according to a Hebb rule:

$$s_{ij}^{t+1} = s_{ij}^t + \eta_{sl,tl} (a_i^t - 0.5)(a_j^t - 0.5)$$

where $\eta_{sl,tl}$ is the learning rate between the area of the source neuron and the area of the target neuron.

As Yaeger states, this simple "summing and squashing" neuron and Hebbian update rule are coarse abstractions of the complexities present in natural nervous system, but they may capture their main information processing attributes. Linsker's [1988] work on self-organisation in the visual system gives us a good hope on this subject.

3.4.3 Reproduction & Genetics

Gaia, like other artificial life simulators, is based on Genetic Algorithms (GA). The main difference to a traditional GA is that reproduction is not based on an explicit fitness function. Reproduction needs two critters at the same time, at the same place and both expressing the reproduction will. Thus it is almost certain that initial populations, created at random, will quickly be led to extinction. To perform an adaptation of this population we introduced the Steady State Genetic Algorithm (SSGA). The SSGA assigns an explicit fitness to each living critter, rewarding the number of eatings, matings and age.

When the population size drops below a predefined critical limit, SSGA reproduces *artificially* according to the fitness function, to restore the population size. This mode should eventually be turned off during simulation. If a simulation leaves SSGA then it is considered that a Successful Behaviour Strategy (SBS) has been reached [Yaeger 1994].

⁵ The connection weight is equal to *synaptic efficacy* if the synapse is excitatory, and equal to $-$ *synaptic efficacy* if the synapse is inhibitory.

⁶ Sensory neurons are excitatory. The only way to provide negative output to other layers is using an output range from -1 to +1.

⁷ For a given critter a global time step is equal to $numAreas$ local time steps ($numAreas$ is the number of network areas).

The genetic code is composed by five chromosomes, and each chromosome can hold a variable set of eight bit genes. Each gene is scaled to use all its dynamic range and then it is converted to Gray code.

The genetic operators used are only mutation and crossover. Mutation works at bit level, while crossover works at gene level, i.e., the crossover operator can not disrupt a gene. The control parameters of these operators such as mutation rate or number of crossover points, are coded in the critter's own genetic data. Although, in nature, this is not entirely true, we hoped it would work as a self tuning GA.

The genetic code also holds physical characteristics, like size, and neural characteristics. And, being the size of cognitive structure variable, so is the genetic code. Chromosome 0 holds tuning parameters and physical aspects. Chromosome 1 holds the neural network main description. Chromosomes 2 to 4 hold area-to-area dependency matrices.

Chromosome 0	Mutation Rate Crossover Rate Crossover Points Size Secondary Size Colour ID Life Spawn Fraction of Energy to Offspring
Chromosome 1	Number of input neurons devoted to green Number of input neurons devoted to red Number of input neurons devoted to blue Number of neural areas Array of number of neurons per area Array of each area bias learning rate Array of each area initial bias
Chromosome 2	Connection density matrix
Chromosome 3	Topological distortion matrix
Chromosome 4	Areas learning rate matrix

Table 1 - Genetic structure

3.4.4 Metabolism & Interactions

One of the main characteristics of living beings is their ability to modify and interact with the environment. Another characteristic of life is metabolism, i.e., the work of all internal systems with the purpose of keeping the entire system alive. In Gaia, interactions can be expressed as energy (or health) exchanges between a critter and the environment.

At each time step a critter loses internal energy (called food value) by moving and by *thinking* (called vital metabolism). The first is computed by the kinetic energy formula, the second is proportional to the neural network size. A critter can also gain or lose energy or health by interacting with the environment. Next we describe the basic interactions. However, more complex interactions such as cannibalism may appear as combinations of these interactions

and behaviours.

Critter-critter fight interaction: Occurs when at least one wants to fight. The combat results in a decrease of health proportional to the opponent size. In this case there is an advantage for being bigger than the opponent.

Critter-critter mating interaction: This only occurs when both express mating desire and both belong to the same species. The result is the birth of one offspring close to its parents, and a decrease of the parents' food value. This energy is transferred to offspring in the amount defined in the genetic code.

Critter-food eating: In order to eat a food cell, a critter must express its will to do so. After eating, the food value and health are increased proportionally to the cell energy.

Critter-wall (collision): The critter gets its health decreased by the collision energy. Small critters lose less energy. For collisions there is an advantage in being small.

Critter's birth: When a critter is created by mating, it is placed in the parents' neighbourhood. If it is created by the SSGA, it is placed at random.

Critter's death: A critter dies when its food value or health becomes negative, or when it ages reaches its life span. If a critter dies with any food energy left, its carcass is turned into a food cell. This enables hunting and cannibalism interactions.

4. Results

It was important to verify the reproducibility of Yaeger's results. The first simulations had only one initial species like in PolyWorld and, even with very different world settings, some common individual and group behaviour patterns were present. After this, we started studying inter-specific competition, predator prey dynamics and sexual selection.

4.1 Group Behaviour

The first successful group behaviour was the **reproduction niche**. It was characterised by a large group of critters being concentrated in a small portion of the world, rotating around a patch of food, like a spiral galaxy. This group behaviour reaches a high reproduction rate, and has success when the autotroph have an high birth rate and a small variance distribution. Otherwise critters die of starvation, because the group does not explore efficiently the surrounding space. It was found that if the world dimensions were changed (for instance to take a very narrow shape) critters regulated their linear speed to be able to rotate without colliding. This group behaviour is somehow a mix of Yaeger's "frenetic joggers" and "edge runners".

In some simulations where autotrophs had large distribution variances (at least along one of the axis), the *niches* started moving like **flocks**. The motion made possible a more efficient world exploration, allowing a longer lifetime for the individuals. **Figure 1** shows an example of this behaviour.




Figure 1 - Example of flocks in Gaia. Circular objects are food cells, triangular ones are critters.

A very interesting and beautiful behaviour emerged between reproduction *niches*, the **mating synchronization** behaviour. It happened when two overlapping rotating niches, forming an eight, synchronized their movement and mating latency. This maximized their reproduction rate with a strong selective advantage.

4.2 Individual Behaviour

PolyWorld known individual food related behaviours were observed (and also the basic PolyWorld behaviour “responding to visual stimuli by speeding up”). New and more complete individual behaviour strategies were found.

The two main individual behaviour strategies observed where:

- food seeking
- obstacle avoidance

Individuals presenting good food seeking strategies were fairly common in some simulations. Typically they were also able to reproduce, thus transmitting this behaviour through generations. Food seeking strategies could be generally described as follows: increasing linear speed when seeing food, decreasing speed and turning when not. It is advantageous to eat as fast as possible in the presence of food, because in a group of critters, the faster eats the most. If there is no food in sight, an energy saving strategy is adopted, which consists of turning in small circles, and scanning the surroundings for newly grown food.

Obstacle avoidance strategies were of two kind:

- simple obstacle avoidance: critters slowed down and turned when seeing an obstacle.
- wall-following: critters moved along the walls, keeping a constant distance to the wall, and turning in the corners.

These strategies were not so common as food seeking strategies, although also being of selective advantage. Nonetheless they appeared more often in narrow shaped worlds.

Most of this behaviours can be established by simple colour associations, like increasing speed when seeing red food elements (excitatory connection from red sensory neurons to linear moving motor neuron).

4.3 Learning

It was found that emerged behaviours were not the result of associative learning during critters lifetime, but instead it was the result of “natural” selection of network structures able to develop well adapted behaviours. Hebbian learning does not seem to provide significant knowledge to the critters. It seems that in *Gaia* Hebbian learning is limited to

developing genetic pre-established connections, i.e., learning occurs over generations and not during a generation (a study of this is presented in [Pereira, Lima & Gracias, 1995]).

One of the most unexpected phenomena was the appearance (and survival) of many non-reactive critters. This happened mainly in simulations where a reproduction niche emerged. These critters were born with high values of initial bias for the several layers, and with very low connection densities. As a result, they presented constant activation values at motor neurons, determined by bias weights. So, the motor neuron activations were in practice static and inherited from the parents.

Recurrent neural networks did not perform better than non-recurrent neural networks. Simulations were carried out with non-recurrent networks, and with general recurrent networks. The same behaviours emerged in both types of simulation.

An interesting phenomena observed was some “senility” in the reactive critters. It seems that sometimes, after some age, Hebbian learning saturates synaptic efficacies. This clear deserves further study⁸.

4.4 Intraspecific competition and speciation

In Gaia, critters species compete for the same resource: food. Two competing species were studied. Two different things could be expected to happen:

- species populations reaching an equilibrium through competition: a constant ratio between the two species (for instance 1:1) could emerge, when birth rate equalled death rate.
- speciation to different energy gathering strategies: one species could speciate in a greater energy gathering (greater size), and in consequence slower exploring speed, and the other one in a smaller energy gathering and greater exploring speed.

The time scales for these two phenomena would be fairly different, the last one should be several times greater.

Only the first phenomena was observed. It could be expected that, starting from the same genetic pool, and the two species being, by this fact, almost equal, the system would have fixed points in any population ratio between the two species, as predicted by the Lotka-Volterra equations for competition [Krebbs, 1994]. But what happened was that one of the species won the competition, and the other extinguished (see **Figure 2**). This is also the most common

⁸ Oja et al [1991] proposed a stable Hebbian learning rule for non-linear neurons. Biological plausibility of this new learning rule is not yet clear, but may eventually solve this problem.

phenomena in nature [Curtis & Barnes; 1994].

The first explanation for this result comes from the fact that Gaia populations are finite (as in Nature). The system can jump from one fixed point to another as a result of a small disturbance. Because the populations are finite, the accumulation of this disturbances can lead to the extinction of one of the species.

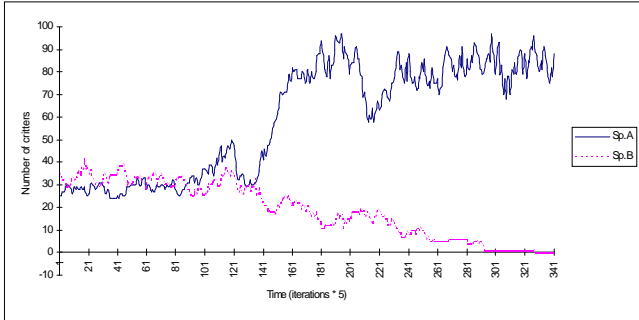


Figure 2: Competing species evolution in a typical Gaia simulation

The second explanation, and the most important is that critter reproduction is "sexual"⁹. Somehow, the reproduction rate of a species (and in consequence the birth rate) is proportional to the probability of critters meetings. In turn, this probability is proportional to the number of critters of that species. So, when one of the species suffers a casual disturbance, and its population is slightly reduced, its birth rate is also diminished. This works as a positive feedback, and the population is easily led to extinction. On the other hand, the other species gets more food and increases its population, which increases the birth rate, and also acts as a positive feedback.

4.5 Predator-prey dynamics

Autotroph and heterotroph population sizes shows a highly correlated evolution (see **Figure 3**). Heterotroph population reaches a steady state when birth rate equal death rate. Birth rate is proportional to available food density, i.e., to autotroph population size. In turn, autotroph population grows at constant rate. Equilibrium is reached when this number equals the number of autotroph individuals eaten, which is proportional to heterotroph population size. Somehow the "food birth constant" defines the environment support value for heterotroph populations. When the food birth constant was increased, heterotroph population average size naturally rose in direct proportion.

Other phenomena observed in autotroph-heterotroph dynamics was stable limit cycles of autotroph and heterotroph numbers (see **Figure 3**). A period of autotroph population growth is followed by heterotroph population growth, and in consequence autotroph population starts to decrease, which also leads to heterotroph population decrease. So, populations of autotrophs and heterotrophs pres-

ent periodic oscillations, which are relatively stable in *Gaia*.

It was also observed a kind of genetic drift. When autotroph population dropped to very low values, critters started to lose their ability to eat. This can be explained by the fact that knowing how to eat was no longer of selective advantage. These critters would die of starvation. Being this characteristic out of selection pressure, harmful mutations would not be rejected.

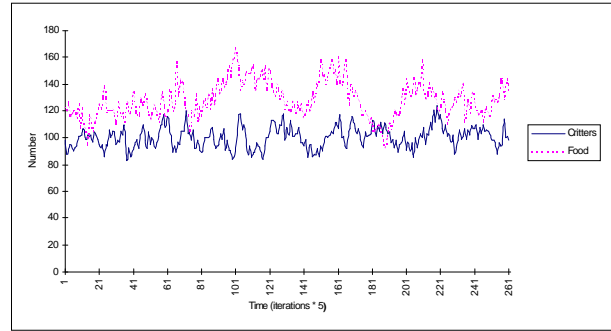


Figure 3: Autotroph heterotroph populations evolution in a typical Gaia simulation

4.6 Self-Tuning GA

Although this was a promising method, we observed that the mutation rate always converged to the minimum allowed value. This can be explained by a selfish attitude of the mutation gene [Dawkins, 1976]. If a gene is carrying a high mutation rate, it is carrying a high chance of being destroyed by a mutation. On the other hand, if it is carrying a low mutation rate, it is carrying a lower change of disruption. This leads to the lowest value of mutation, i.e., a higher chance of gene survival. In fact, this tuning structure does not optimise the GA itself, but the survival of the low mutation allele.

Crossover rate and crossover points, never seemed to converge to a particular value, making us wonder if there was any tuning at all.

4.7 Sexual selection

A kind of "sexual selection" emerged in *Gaia*. When SSGA maintained the simulation, there was no significant sexual selection and critters evolved to the minimum possible size, as a strategy of energy saving. But when the population went out of SSGA, the ability to reproduce started to be very important, and big critters become common. They reproduced easily, due to the bigger reproduction radius. It is interesting to see how the reproduction ability pressure is higher than the low metabolic rate pressure.

5. Conclusions

Results suggest that learning in animals with very simple nervous systems is mainly phylogenetic. Associative Hebbian learning, seems to play a minor role when compared to the natural selection of genetic structures of nervous systems. It also seems that, Hebbian learning in neural networks with the topologies and sizes used in *Gaia*, cannot

⁹ With "sexual" we mean that it is necessary the mating of two critters for reproduction. There's no vegetative reproduction.

explain higher associative learning mechanisms.

It was found a deep connection between the complexity of behaviours emerged and the difficulties posed by the world. In worlds where was very easy to survive, only dumb critters appeared. But if worlds were difficult (like having few food elements), intelligent beings appear. Nature always goes the simpler way.

The behaviour study showed that the most important feature was the reproduction ability. Any interesting behaviour, such as obstacle avoidance, can only "survive" if is connected with a good reproduction strategy. This differs from classical GA. In GA every new feature is expressed in fitness. In Gaia, even if one presents a good eating strategy, this feature will disappear if one cannot reproduce.

For competing species with sexual reproduction, the importance of reproduction ability can also explain the typical natural "winner takes all" competition result. A competing species that suffers a decrease in population size, also suffers a global decrease in reproduction ability (the probability of finding a mate decreases), which acts as a positive feedback until species extinction.

This work shown that some ecology features can be studied in an artificial system. Nonetheless, ecology simulators still have many limitations.

6. Future Work

In such a complex system as nature, it is actually impossible to produce what might be considered a complete simulator. There is a never ending number of features that can be added to make it more close to its biological counterpart. But Gaia itself is a very complex system and a more carefully analysis on the effects of some parameter variations is the first step.

For new versions of Gaia we are considering the use of alternative cognitive structures, such as reinforcement learning, conditioned leaning with modulatory synapses and systems leaving for the GA the choice of the best learning structure.

Another future direction is the improvement of biophysics by introduction of temperature. Critters metabolism loss would raise temperature, zones without critters would decrease temperature. Food should have an optimal growth temperature. With this, we expect to see self-regulation features at extra-organism level.

The inclusion of diploidism and dominance relationships is another idea for Gaia. We are now studying a new diploidism method and testing it in very simple experiences.

Acknowledgements

The author would like to thank the Portuguese Youth Institute and the Scientific Youth Association for their support on this research.

We also would like to show here our gratitude to Tiago Domingos for all his cooperation in the reviewing of this paper.

References

- Ackley D., Littman M. [1992]. *Interaction between learning and evolution*. Artificial Life II, C. Langton (Ed.), Addison-Wesley: Reading, MA.
- Conrad, M [1987]. *Computer Test Beds for Evolutionary Theory*. Oral presentation at Artificial Life I Conference, 1987
- Collins R.J., Jefferson D.R. [1992]. *Antfarm: toward simulated evolution*. Artificial Life II, C. Langton (Ed.), Addison-Wesley: Reading, MA
- Curtis, H., Barnes N.S. [1994]. *Invitation to Biology*. Worth Publisher.
- Dawkins, R. [1976] *The Selfish Gene*. Oxford: Oxford University Press.
- Harp, S., T. Samad, and A. Guha [1990]. *Towards the Genetic Synthesis of Neural Networks*, Proceedings of the Third International Conference on Genetic Algorithms, J. D. Schaffer (Ed.), Morgan Kaufman: San Mateo, CA
- Johnson A.R. [1994]. *Evolution oversize structure predator prey community*. Artificial Life III, C. Langton (Ed.), Addison-Wesley: Reading, MA
- Krebs, C.J. [1994]. *Ecology: Experimental analysis of distribution and abundance*. Harper Collins.
- Lindgren K., Nordahl M.G., [1994]. *Artificial food webs*. Artificial Life III, C. Langton (Ed.), Addison-Wesley: Reading, MA
- Linsker, R. [1988] *Self-Organization in a Perceptual Network*. Computer, March 1988, Vol.21, No. 3, 105-117.
- Margulis L., Sagan D. [1986]. *Microcosmos*. Summit Books (Ed), New York.
- Miller G.F., Todd P.M. [1991]. *Exploring adaptive agency I: theory and methods for simulating evolution of learning*. Connectionist models, Proc. 1990 Summer School. Tourestzky D.S. (Ed), Morgan Kaufmann, CA.
- Packard, N.[1989] *Intrinsic Adaptation in a Simple Model Evolution*, Artificial Life I, C. Langton (Ed.), Addison-Wesley: Reading, MA
- Pereira, H, N. Gracias, J. Lima [1995]. *Solving the XOR Problem Using Genetic Algorithms*, Proceedings of 1st Workshop of Biomedical Engineering, A. Rosa (Ed.), Instituto Superior Técnico
- Ray, T.S [1992]. *An Approach to the Synthesis of life*, Artificial Life II, C. Langton (Ed.), Addison-Wesley: Reading, MA
- Taylor, C. E., D. R. Jefferson, S. R. Turner, and S.R. Goldman [1989] *RAM: Artificial Life for the Exploration of Complex Biological Systems*, Artificial Life, C. Langton (Ed.), Sante Fe Institute Studies in Science of Complexity, Proc. Vol. VI. Redwood City, CA: Addison-Wesley
- Yaeger, L. [1994]. *Computational Genetics, Physiology, Metabolism, Neural Systems, Learning, Vision and Behaviour or PolyWorld: Life in a New Context*, Artificial Life III, C. Langton (Ed.), Addison-Wesley: Reading, MA