

# Environmental Fitness for Sustained Population Dynamics

Nicolas Brodu

Department of Computer Science and Software Engineering  
Concordia University, Montreal, Quebec, Canada, H3G 1M8  
nicolas.brodu@free.fr

**Abstract-** This study introduces an environment in which agents evolve freely, without an explicit fitness function. The evolution is mainly directed by environmental conditions, and their coupling with the agents. The goal consists in identifying preconditions for a sustainable environment allowing diversity. A method is proposed to analyze which configurations lead to rich population dynamics. The system global behavior is then tuned accordingly. Visualization is also a key component of this project: it offers a qualitative understanding of the simulations.

## 1 Introduction

The aim of this study is to control the evolution of agents on a global scale, without introducing an explicit local fitness value. An environment is set up in which agents can reproduce freely, and their AI is controlled by a genetic algorithm. The goal is to identify which preconditions lead to a sustainable environment, allowing a diversity of agent behaviors.

Genetic algorithms are often used as machine learning tools, and have proved to be good random search optimizers. Their learning bias is usually mainly determined by a fitness function. This approach performs particularly well for some class of problems, and is not appropriate for others. This study precisely avoids defining a fitness function: agents are free to live as long as they can, and reproduce if they both wish to and can do so. In this case, the standard optimizer framework is not directly applicable.

In this project genetic algorithms are instead considered in a different way. The conditions for agents survival and reproduction are set up locally, in both space and time. Agents evolve, and phenomena are observed on larger scales: periodic patterns (time), covering the whole environment (space). Ideally, the system could be designed by using preconditions and combinations of factors, in order to induce a specific global evolution. Unfortunately this reverse engineering problem is tough, and not specific to this study. Emergence [5] is both a promising and challenging field of research.

Nevertheless, it is still possible to determine the effect of the most influential parameters. Evolution is then controlled indirectly, without using fitness functions. Re-using the machine learning perspective, the system bias is determined by: The nature of the agents' AI, the nature of the environment, and their coupling via the agents sensory inputs and effectors. The physics of the world, and the way agents interact, define the conditions for survival and

reproduction. Agents live as long as they can, and reproduce if they both wish to and can do so.

The challenge is then to generate or maintain diversity. In the real-world as well as in artificial setups, the reasons for biological diversity and richness of behaviors are still subject to controversies. Multiple factors have been investigated: Mutualism [9], Competition [7], and its interaction with Predation [1]. For example, plants competing for the sun light in the canopy of tropical forests may explain part of their observed diversity. But this cannot be the whole story. Complex behaviors appear even when there is no scarcity of resources. For example, bees may arguably compete against other insects for access to pollen. But this resource is sufficiently abundant so their organization is based more around collecting it, than around fighting for it. This study introduces resource competition and will test its influence in a digital setup, but does not address the more general and interesting case.

Predatory relationships are also used as an attempt to generate and maintain diversity. The goal is not to take position in the aforementioned debate, but rather to set up a possibility that can be used by the agents, or not. Agents can get resources by hunting other agents in addition to grazing. Evolution acts on the AI decision to do either one, or both. Similarly the resource competition aspect is also parameterized. Agents are allowed to reproduce only after they have acquired enough resources, but this threshold is considered as an environment parameter. Determining its influence is part of the experiments presented in the next sections.

The following approach is proposed:

- An environment is created for the agents to evolve, but no explicit fitness function is defined. Competition and predation are introduced in an attempt to maintain and generate diversity. Agents live as long as they can, and reproduce if they both wish to and can do so.
- The influence of the environmental parameters on the population dynamics is evaluated. The objective is to determine what preconditions lead to a sustainable diversity in terms of agent behaviors. Statistical metrics are associated to the system global response.
- This environment is used to drive evolution on a global scale: Variation along each dimension is measured in the high dimensional space of environment parameters. By analogy with gradient descent algorithms, the parameters are updated using these variations to bring the system closer to the desired behavior.
- Visualization is another main component of this project. Too abstract worlds may be hard to analyze

without an intuition on their internal workings and relationships: Direct observations may provide a qualitative understanding of complex systems. To date, some patterns are only recognizable by human eyes. The ability to restart the simulation in graphics mode has proved to be an invaluable investigation tool.

The next section introduces the simulated world and its features. Section 3 describes the Agents' AI and the genetic algorithm. Section 4 presents the experiments and the metrics used to evaluate the quality of the environment. The results are discussed in Section 5, and used to control the system global behavior in a final experiment. Section 6 concludes on the issues encountered in this study.

## 2 Description of the world

The world considered in this study is a continuous 3D environment. Agents are embodied as solid mobiles, on which steering forces are applied. Classical physics is then used to numerically integrate the motion of the agents in time. This model was inspired by [10] and extended.

In particular, steering forces fall into 2 categories:

- External environmental forces and constraints, including gravity. An agent has no control over these forces. They are applied whatever it does.
- Internal agent's decisions, based on its capabilities. The agent is limited by a maximum speed and acceleration. Within these limits, the agent is free to decide on which force to apply to achieve its goals.

A random terrain is generated, using spatially consistent noise as a height field. This provides a full 3D scenario, compared to using a flat floor. By construction the terrain is made cyclic along the X and Y dimensions, so it is perfectly tileable without gaps. Along the Z dimension, the world is limited by the floor, and has no upper bound. Obstacles are scattered randomly, and collision avoidance will have to be accounted for in the agents' AI.

A physically based notion of energy is the driving resource. Storing energy can only be done up to a certain limit, and incurs a small mass penalty (similar to storing fuel, or food in the stomach). This same footprint ratio is also used the other way around, for digesting a mass into equivalent energy. These maximum energy, and energy footprint ratio, are environmental parameters.


The system is *dissipative*: Agents use energy when applying forces, for reproduction, for fights, and more generally for anything in which they actively participate, including self-maintenance.

The system is *open*: The external, and only, source of energy comes from the grass. It grows on the floor with a constant rate and up to a maximum density.

An agent can get energy either by grazing this grass, or hunting other agents. Both actions are fully parametrized by the agent genome, and part of its AI (described in the next section). Agents can graze only if they move slowly, below a speed threshold.

Hunting is done on a fixed prey-predator relationship.

The only discrimination between agents belonging to distinct species is their prey-predator relationship. Agents of a species prey on agents of 2 other species, forming a cyclic predation graph. In a 5-species graph, up to 2



Agents of each species predate on agents of 2 other species, cyclically. This is actually the only difference between species: all agents start with genomes drawn from the same uniform distribution. In particular, no bias is introduced toward grazers or hunters for distinct species. This differentiation may appear, or not, with evolution.

**Figure 1: Species predation graph.**

species may go extinct without breaking the cyclic predation feedback.

The last component of the world is the genome representation. An initial problem while defining this simulation was making the agents exchange genetic material without favoring collisions. Another problem came from the absence of a fitness selection / reproduction explicit cycle: agents do not reproduce synchronously. These problems were both solved by introducing spores: Genomes are decoupled from the agents, in space and time. Agents can emit spores in the environment, and in a separate step use spores from other agents when they wish to reproduce. Reproduction is asexual, but self-pollination is not allowed. Spores have a half-life time determined by their decay rate, like radioactive elements. They are modeled by a spatial concentration, so as to avoid creating one object per spore for efficiency reasons. The concentrations are updated by numerically integrating classical diffusion equations.

Diffusion is done on a per-genome basis: A list of available spores is maintained at each point, together with their concentrations. Agents pick spores for reproduction at their current location by choosing one compatible genome at random, if any. Each genome has a probability to be picked proportional to its concentration.

Figure 2 shows the result of a run in which the population exploded after many prey/predator cycles. It is a contrast-enhanced, black and white, capture of the live program. In this example a few hills define a valley (in the middle) and a higher altitude plateau (on the left). White patches of floor where the grass was completely eaten are visible. Spore clouds fill the valleys. The populations graph for the whole run is overlaid on top of the scene; It can be switched on and off in real-time.



**Figure 2: A random world example.**

The time on the bottom right is expressed in simulation time units, as used by the physics engine: It is consistent with the other units for the steering forces, acceleration, mass, length, energy, etc. Thanks to an event scheduler architecture, the simulation time is completely uncorrelated from the real time. In particular, pausing the simulation for visual introspection has no effect on the numerical integration results. The system is completely determined by the chosen random seed. All experiments mentioned in this study are therefore perfectly reproducible (see Appendix).

Agents are given the occasion to update their steering force according to their AI periodically (in simulated time), at a given frequency. Physics are integrated at a higher frequency, and constraints are then applied independently of the agents (like being above the terrain). Though mainly introduced for computational reasons, this frequency separation also reflects a natural fact: One can only think about so much actions in a given time, but the physics inexorably unroll regardless of these actions. A third integration step is used for the spore diffusion equations.

### 3 Agents' AI and genetic algorithm

Agents have access to their local neighborhood, but see only in front of them. This was set up to give an advantage to predators, according to preliminary tests. The agents can detect the amount of grass available, the position and the nature of their neighbors (obstacles and other agents), as well as their apparent speed and direction, and the local spore concentrations. They also have access to relevant environmental parameters, like for example the values of the gravity and their own physical speed limit. All these represent the sensory inputs of an agent.

Basic steering forces [10] are provided as building blocks for the AI. These correspond to common situations: force to apply to avoid an obstacle, to flee a predator, to chase a prey, etc. However, no precomputation involves data the agent wouldn't have access to otherwise: these basic forces should really be seen as useful subroutines.

Part of the AI consists in combining these helpers in a consistent way. In a presence of a prey, a predator, and an obstacle, which of seek, flee, and avoid should take precedence? In the scenario for the current study, an influence inversely proportional to the squared distance of each source is used. This was done for practical reasons, and because this distance weighting scheme favors handling the closest objects. The resulting factors are combined, with an additional weight coming from the genome: some agents will thus prefer fleeing, other obstacle avoidance, etc, and evolution acts on the weights.

The other part of the AI consists in deciding on a general activity: grazing, hunting, or seeking a reproduction place. Which activity is chosen depends on how much grass is present at the current location, how many preys and predators are visible, whether the agent is ready for reproduction or not, and how hungry it is. Each factor is once again weighted by a gene, and so are the decision thresholds. Each activity has its own specific set of weights for basic forces, as previously described.

Grazing also introduces weights for deciding when the current location lacks grass, then for exploring and moving to a new place.

Hunting mainly corresponds to the aforementioned weighted combination of seek, flee, and avoid, with the restriction of pursuing only one prey at a time. Attention is also given to the movements of mate agents belonging to the same species.

Seeking a reproduction place consists in wandering and looking for compatible spores. Final decision to reproduce is taken on prey/predator/mate/obstacle ratios, computed by taking in account genome weights as before.

Fights occur in any of these activities, as part of the collision detection mechanism. Since agents see only in front of them, a predator attacking a prey from behind is considered to have a surprise advantage, and it wins. Otherwise, the predator decides on how much energy to engage in the fight (genome dependent). According to its own energy, the prey may then escape, or be eaten. In this last case digestion is perfect: all the prey mass is converted into energy according to the energy footprint ratio previously introduced.

Each time an agent gains energy, whether by hunting or grazing, it may decide to emit a spore. This incurs an emission cost. The gain/cost ratio is once again weighted and thresholded by genes to take the decision. Agents may also emit spores when they are ready to spawn but found no compatible spore at their current location. This second way also indirectly implies having acquired enough energy. As said in introduction, restriction for emission on energy gain was set up as an attempt for directing evolution through resource competition.

Reproduction occurs when an agent has enough mass and energy, has found a compatible spore, and it likes its current location. It then spawns a child by giving it part of its mass and part of its energy, the proportions of each are genome-dependent. A minimum mass constraint is set up for the agents. A spawning agent must be at least twice this minimum to be able to reproduce, so that both itself and its child respect the constraint. Thanks to this mechanism, no energy is indirectly introduced during reproduction.

Consequently, this mass transfer implies a development phase. The children are born with an initial mass, at least the minimum, and must metabolize part of all the energy they gain to grow and be able to reproduce themselves. The energy gain amount to convert is a genome-dependent ratio. A maximum mass is also set up: agents have finite bodies. The development phase stops when reaching this maximum mass, and at this stage the energy is fully accumulated. Any extra energy gained above the energy maximum is lost. Together with the minimum mass to give to a child, these bounds indirectly limit how much children an agent can spawn without having to gain energy again, once it has found a good reproduction place.

## 4 Experiments

### 4.1 Setup

As previously mentioned, the goal is to study the influence of the environment parameters on the population dynamics. Some parameters, like the integration rates and

the gravity, are not investigated. More attention is given to the following parameters:

- Agent maximum mass
- Agent minimum mass
- Energy footprint ratio
- Agents maximum energy
- Spore emission energy threshold
- Spore emission energy cost
- Self-maintenance energy cost, per simulator time unit
- Maximum grass density (energy/surface)
- Grass growth rate (density per simulator time unit)

It is computationally too expensive to study all combinations of the environment factors for different values. Each of them is thus studied separately, with the others kept fixed. This is equivalent to exploring the surroundings of a chosen point in environment parameter space, along each dimension.

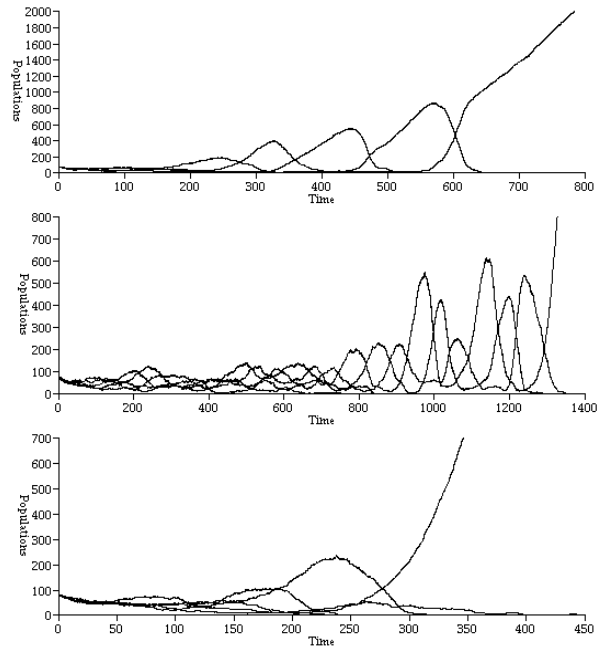
Preliminary experiments were used to choose this initial point. Qualitative feedback from the visualization of the system in graphics mode directed the random search. Figure 2 shows the result of a run in these conditions: The population explodes, but many prey-predator cycles are observed before one species takes the lead, sometimes as in this example with an intermediate reduction in total population size. This tendency is generally observed for other runs with these initial environment parameters. However, with non-negligible probability, some experiments show flatter scenarii, and the transient period before explosion greatly differs from one run to another, as shown in Figure 3.

This variability was thought to be a good indicator of the richness of this environment. Especially when considering qualitative phase transitions by comparison with other systems [3]: A mixture of long transients, great variability, and persistent qualitative behavior is usually an indicator of what's called the *edge of chaos* [8]. There is unfortunately no way in this study to check on this hypothesis: Sampling the system behavior at the surrounding points in the high-dimensional environmental landscape is too computationally expensive. The initial point was chosen arbitrarily, and may not correspond to an optimum in terms of environmental exploration: some other setups may exhibit richer behaviors.

## 4.2 Evaluation

The main difficulty is the definition of relevant metrics for the analysis of the experiments: What defines the quality of the environment parameters? How to estimate the population dynamics? Is it possible to synthesize the system behavior and then monitor its potential phase transitions [3, 8], supposing this is the case?

The quality of the environment comprises its ability to ensure sustained populations, but not only that: In an environment with very energetic grass and high grass growth rate, for example, agents would have no incentive to move to get energy. Up to an implicit population limit, defined by the size and shape of the finite world: the agents would then start competing for this resource. Dynamic environments are more interesting, not only long-lasting ones.



These graphs show three different population dynamics with the same initial set of environment parameters. The only difference between these runs is the random seed. Species population numbers are plotted through time, one curve per species. The right part of the second and third plots is an increase of the only remaining species until the simulation stopped. A zoom on the interesting features was preferred, and the scale is different in each graph: the bump in the third graph is about the same size as the flatter one in the first graph. Prey-predator cycles are visible, and population extinctions. In the first plot, there is a clear cut for the only surviving species when the last preys go extinct, and the agents must rely on grazing.

**Figure 3: Population dynamics diversity with the initial environment parameters.**

The first criterion chosen concerns the population dynamics: number of cycles or bumps in the species population graph. For the graphs visible in figure 3, the counts are respectively 4, 14, and 1. These numbers are computed, not visually estimated, so as to ensure the same definition for all runs: An amplitude variation, up and down, above a threshold in each direction. The threshold chosen is the initial population size: at least this number of descendants must be created to count for a bump. With this definition, the final explosion of a species is not counted (the “bump” is not complete). The number of cycles metric may thus be safely applied to all species, whether they go extinct or not. The average number of cycles/bumps per species is an activity indicator for a particular run. Mean and variance of this measure on a batch of runs, with the same parameters but different random seeds, reflect the potential of activity for these environmental conditions.

Time before species extinctions and explosions is also monitored. Preliminary tests were used to choose a population count of 2000 individuals as the definition of explosion. This threshold does not correspond to the maximum sustainable population, especially since this maximum is not constant: It depends on the environmental parameters and the shape of the terrain, thus on the random

seed. In some runs though, there is a clear slowdown in population increase. An hypothesis to explain this slowdown would be that the population maximum is close for these particular runs.

Long transients before extinction or explosion are assumed to represent how sustainable an environment is. The limit case would be a stable ecosystem. They are averaged over a batch of runs, the same way as for the number of bumps/cycles metric.

Another measure of activity for a particular environment is the number of generations that appeared during the run, per time unit. This indicator complements the cycle metric aforementioned. Unfortunately, unlike the number of cycles, the number of generations is very sensitive to the last final explosion phase. For this reason, only generations before extinction are monitored, and species remaining at the end of a run are discarded. Mean and variance for a batch of runs are then computed, but taking in account the effective number of species retained for the whole batch, whatever the number of runs.

Per-generation statistics are also collected, but should be interpreted with caution. A problem arises from the asynchronous generation definition: Generation 0 is assigned to the initial agents, and this number is incremented for the children when an agent spawns. However, in this study, agents live as long as they can, and may live together with their descendants. A child of generation  $N+1$  may have as parents an agent of generation  $N$  and another of potentially any  $M \leq N$ .

This problem can be turned into an opportunity to check for the persistence of agents in time: the average life time for initial agents is also monitored. Statistics about agents from further generations are also recorded, but are less significant: the asynchronous problem makes it hard to justify comparing life times. Especially, the qualitative nature of the system changes near the end of the run, in case of a population explosion.

When their life time standard deviation is the same order or more than the mean, first agents are widely spread in time: some live very long, and others a shorter time. We can thus safely say agents from different generations live side by side. On the contrary, this is unlikely to be the case in the presence of prey/predator cycles: most agents of a given species die between the maximum and minimum of a cycle. The chance of an agent from the initial generation to be alive thus reduces at each cycle. This is reflected by a mean (and standard deviation) much smaller than the duration of the run: we see a turnover in agents. Both ratios are monitored: standard deviation over mean, and mean over extinction time. They are used as other indicators for a given environment. For the same reason as before, surviving species are not taken into account.

A measure for the distribution of grazers/hunters within each species is deduced and monitored, by using other per-generation statistics. It uses the observed variance for the AI weights related to taking the grazing or hunting decision (these weights were introduced in section 3). As previously said, this measure should be interpreted with caution, due to the generations definition. However, what's interesting in this case is not so much the values themselves, but how they evolve. This behavior span measure is computed for the initial agents and for the last generation counting more than 30 individuals.

The average of the measure indicates a global behavior span for the whole batch of run. The ratio of the averages between the end of a run and at the beginning, indicates how this range of behaviors evolves. The variance of this span is more interesting: it is a direct measure of how much agents of distinct species differ. The ratio between the variances at the end and at the beginning of the run indicates how much species differentiation is going on. Both ratios are computed only for species that goes extinct, for the same reasons as previously mentioned.

Differentiation is expressed for the purpose of this study in terms of the general behavior of agents of distinct species. This is not to be confused with a more interesting emergent speciation process, as considered in other studies [2]. Species are here explicitly predefined by prey-predator relations. Agents of the same species can only reproduce among themselves. It's thus logical to expect separate evolution for different species, and separate values for genome statistics, including a bias toward grazing or hunting.

However all agents, whatever their species, start with a genome drawn from the same uniform distribution. Within a single species, significant variation from this initial distribution at the end of the run is worth noting. This corresponds to the behavior span mean ratio.

Inter-species differences are expected as previously mentioned: The behavior span variance ratio is more than 1, due to the separate evolutions. But it is not trivial whether agents of each species keep a range of behaviors, or strongly specialize toward hunting or grazing only. For some environmental conditions, the ratio between the variance at the end and beginning of the runs is the order of 1 or 2. But for other environments, the ratio reaches higher counts, with observed values of 15 or more. In these environments, species are much more specialized.

More generally speaking, we can consider the agents behavior to be their phenotype, since their genotype consists in AI parameters. The difficulty is to interpret the influence of environmental changes in terms of genome values. The grazing/hunting behavior span ratios are indicators that can be derived from the genomes, but these are not the only ones.

## 5 Results and Discussion

### 5.1 Environment parameters landscape map

Each parameter is modified in turn, and a batch of runs is executed for the new environment conditions. Figure 4 summarizes the influence of each parameter around the initial conditions. It can be seen as a snapshot of the environment parameters landscape in higher dimensions, at the point chosen for this study.

Many results correspond to expected variations. For example, increasing the self-maintenance cost causes more and faster extinctions. Similarly, increasing the maximum grass density both reduces the number of extinctions, and increases the time to get there for a species. Given that grass is the system energy input, this is not surprising.

Other results are less immediate. For example, does more energy mean more population cycles? One could imagine less energy means more competition, thus more

Parameters Metrics	Max. mass	Min. mass	Energy footprint	Max. energy	Spore threshold	Spore cost	Self-maintain	Grass density	Grass growth	Gradient descent
Cycle count $\mu$	$\approx$	$\approx$	$\uparrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\cap$	$\uparrow$	$\uparrow$	$\nearrow$
Cycle count $\sigma$	$\cap$	$\cap$	$\nearrow$	$\cap$	$\cup$	$\downarrow$	$\downarrow$	$\uparrow$	$\uparrow$	$\downarrow$
Generations/Time $\mu$	$\cap$	$\cap$	$\nearrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\cap$	$\nearrow$	$\nearrow$	$\nearrow$
Generations/Time $\sigma$	$\cap$	$\cap$	$\cap$	$\cap$	$\cap$	$\cup$	$\cap$	$\cap$	$\cap$	$\downarrow$
Extinction count $\mu$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\nearrow$	$\nearrow$	$\downarrow$	$\downarrow$	$\approx$
Extinction count $\sigma$	$\cup$	$\cup$	$\approx$	$\nearrow$	$\cup$	$\downarrow$	$\cup$	$\cap$	$\cup$	$\nearrow$
Time to extinction $\mu$	$\cap$	$\nearrow$	$\cap$	$\cap$	$\cap$	$\nearrow$	$\downarrow$	$\nearrow$	$\downarrow$	$\downarrow$
Time to extinction $\sigma$	$\cap$	$\cup$	$\cap$	$\cap$	$\cap$	$\cup$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$
Time to explosion $\mu$	$\nearrow$	$\cup$	$\downarrow$	$\cup$	$\downarrow$	$\cup$	$\cup$	$\cup$	$\downarrow$	$\nearrow$
Time to explosion $\sigma$	$\nearrow$	$\cup$	$\downarrow$	$\downarrow$	$\cup$	$\nearrow$	$\cup$	$\cup$	$\cup$	$\nearrow$
Life time span $\mu$	$\approx$	$\approx$	$\cup$	$\nearrow$	$\approx$	$\nearrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$
Life time span $\sigma$	$\downarrow$	$\nearrow$	$\downarrow$	$\nearrow$	$\downarrow$	$\nearrow$	$\downarrow$	$\cup$	$\downarrow$	$\downarrow$
Life time duration $\mu$	$\nearrow$	$\approx$	$\nearrow$	$\nearrow$	$\nearrow$	$\approx$	$\approx$	$\downarrow$	$\nearrow$	$\downarrow$
Life time duration $\sigma$	$\approx$	$\approx$	$\approx$	$\approx$	$\nearrow$	$\approx$	$\approx$	$\downarrow$	$\downarrow$	$\downarrow$
Behavior span $\mu$ ratio	$\nearrow$	$\cap$	$\nearrow$	$\cap$	$\downarrow$	$\cap$	$\cap$	$\cap$	$\nearrow$	$\nearrow$
Behavior span $\sigma$ ratio	$\cap$	$\cap$	$\cap$	$\downarrow$	$\cap$	$\downarrow$	$\downarrow$	$\uparrow$	$\downarrow$	$\nearrow$

Legend:

- $\nearrow$  Direct variation
- $\searrow$  Inverse variation
- $\uparrow$  Strong direct variation
- $\downarrow$  Strong inverse variation
- $\approx$  No significant variation
- $\cap$  A maximum was observed at the initial point
- $\cup$  A minimum was observed at the initial point
- $\mu$  Mean of the corresponding indicator
- $\sigma$  Standard deviation of the corresponding indicator

The gradient descent column is the result of the experiment described in section 5.3.

**Figure 4: Environment parameters landscape snapshot, around the chosen initial conditions.**

influence for prey-predator relations. In fact, the reverse happens. It turns out grass factors are the most important ones in terms of cycle numbers. An average of more than 3.5 per species, with standard deviation 2, was observed at high grass density. This is much more than the examples from Figure 2 sampled at the initial point in parameter space: The average number of cycles at the initial point is 1<sup>1</sup>.

Among other unexpected results, the maximum mass is the most influential parameter for the behavior span mean ratio. The ratio increases with the mass limit, which means species are less specialized. A convincing hypothesis is missing: Why would allowing more mass lead to more versatile agents? The relation with the number of children explained in the previous section is probably not the only

effect. Moreover, no significant variation was observed on the cycle count and life time span for the mass parameters. The previous section introduced an hypothesis about the number of cycles, and the agents' life time span: supposedly, more cycles would mean less chances for agents of the first generation to survive a long time. This hypothesis is partially confirmed: In the cases were significant variation was observed, and the parameters landscape is not an optima, there is an inverse correlation between the cycle count and the life time span means. No significant variation is observed for both metrics in the case of the mass parameters, so this does not invalidate the hypothesis. The environment landscape optima are more difficult to interpret. In the case of a high energy footprint ratio, the cycles minimum were inspected graphically and are actually quite high: The agents of the first generation still have a good chance to survive each cycle. This possible explanation is consistent with the observed decrease in life time span standard deviation.

Some parameters landscape consist mainly of minimum and maximum. The generations/time ratio standard deviation is one of them, for example. An hypothesis would be these criteria are too "erratic": the corresponding parameter landscapes are jagged. The observed optima would be local, and not relevant to the quality of the chosen point. More experiments would be needed to validate this hypothesis.

## 5.2 Interpretation

Along some dimensions the parameters landscape form a saddle point, whereas variation is monotonous along other dimensions. A reason for this diversity may be the influence of the observer during the initial directed search: Environments with a broad range of global responses and behaviors will probably be estimated the most interesting ones.

Another reason for diversity might be structural and inherent to the system. Spatial sensitivity is high, and geographical features may change the qualitative behavior of the simulation. Figure 2 shows agents grazing on a sanctuary on a shallow plateau in the left of the picture, above a valley where other agents fight intensely. Until a predator finds about this wonderful prey reservoir, of course<sup>2</sup>... This second reason for results diversity is captured by the statistics, and partly explains the observed high variance in each of the metric chosen: the random seed, determining the terrain shape, has a high relative influence compared to the environmental parameters themselves.

Some metrics, like the life time duration, are defined at the agents level. Results for these measures show the greatest number of no significant variations. The analysis is more pertinent for higher-level metrics: notions and relations like the population cycle count, defined on the world scale rather than on the individual agents scale. This is a characteristic feature commonly found in studies involving emergent phenomena. In this example, population cycles can only be defined on larger time scales, and

<sup>1</sup> Remember this value is an average on all species. In figure 3, some species contribute with a count of 3 or 4, and others have no cycle. The cycle count standard deviation is about 0.7 for the initial conditions.

<sup>2</sup> In this particular example, the plateau was first occupied by two other species. The flat population curve below the cycles corresponds to the species on the plateau, and the sharp drop near the end corresponds to the coming of a predator.

so are the notions of population extinction and explosion. Some parameters, like the spore emission cost and the grass-related variables, have a strong influence globally. These parameters only act indirectly on the agents' AI. For example, a grazing agent will move more frequently if the grass density is lower. On the contrary, some parameters have a direct impact on the AI, like the maximum mass and energy limits. The results are more significant when both a global relation and a global parameter are related.

One of the initial goal was investigating on which preconditions lead to a sustainable diversity in terms of agent behaviors and species. Using statistical higher-level metrics has proved to be more efficient than local measures on agents. Moreover, the shape of the corresponding environmental landscape is easier to interpret.

Unfortunately, so far all configurations lead to a population explosion. Some exhibit longer transients than other, but with the same result nonetheless. Is this fact due to the system lacking a necessary precondition? Or is it just an artifact of the chosen point in environmental space. In this last case, could we apply some kind of gradient descent rule, using the observed variations, to tune the system toward a desired behavior?

### 5.3 Tuning the environment

The previous experiments explored the variations of the different metrics along each dimension (corresponding to each parameter). A final experiment was set up using the observed variations. Several parameters were modified simultaneously, using small changes in the directions of interest. Specifically, more cycles and longer transients are the main desired global features. Specialization and specie differentiation are also considered, as secondary points.

The maximum energy limit is decreased: it has a negative influence on the number of cycles and the behavior span  $\sigma$  ratio. The chosen initial value for the maximum energy also corresponds to a minimum in terms of time to population explosion. Spore costs are both decreased for similar reasons. The self-maintenance cost is kept unchanged: it has opposite influences on some of the desired features, and corresponds to a maximum for the number of cycles anyway.

The grass density is increased. It has a strong influence on the number of cycles, and the chosen initial point was at a minimum for the time to explosion. It also has a beneficial effect on the species differentiation (behavior span  $\sigma$  ratio). On the contrary, the grass growth rate is decreased: it has a too negative impact on the time to explosion.

Mass parameters have no noticeable impact on cycles, but they have a negative influence on the number of extinctions: they are both slightly increased.

The energy footprint is slightly decreased despite its influence on the cycles count. The variations it induces on the life time spans, the behavior spans, and the number of extinctions justify this choice.

Results for this joint variation experiment are provided in the last column of the array of the figure 4. The number of extinctions remains the same as for the initial point, and the average time to extinction decreased. Both these negative results may be explained by the contradictory effects of the parameter changes that were introduced. The variations obtained for the cycle count, the time to explosion,

and the behavior span ratios are in agreement with the expectations.

Quantitatively, the system reacted as planned. The experiment is thus globally successful, and the system overall behavior was tuned toward the desired goal.

Qualitatively, the difference is visible in graphical mode. The most noticeable features of the new point are more grazed patches on the floor, and smaller groups of agents. The new floor appearance is a logical consequence of decreasing the grass growth rate. The original conditions lead to vast zones controlled by one species of agents. With the new conditions, smaller zones are observed. The reason for this is not known.

Only one step was realized along the gradient descent lines. Ideally this whole study would be done again at the new point, for a second step of gradient descent, and so on, until an optimum is found. This could not be done due to the necessary amount of computations involved, but the current results are encouraging.

## 6 Conclusions & future work

### 6.1 Summary and review

This study has focused on the following aspects:

- Implicit fitness. Agents are constrained by their coupling with the environment and their AI. Qualitative and quantitative effects of environmental changes were investigated on a global scale. The notion of an environment parameters landscape was introduced.
- Free evolution. Agent selection was induced by resource competition by design. However, the effects of changing the amount of resources available could not be easily explained in terms of competition, as shown by the grass density example aforementioned.
- Visualization. This aspect should not be neglected: It allows for a better understanding of the system nature, and to identify patterns not easily detectable otherwise. The spatial sanctuary effect aforementioned is one of them.

The initial goal was to investigate the preconditions for a sustainable environment, allowing diversity. Statistical indicators were derived to measure the general influence of the environment on the behaviors of the agents.

An initial set of environmental parameters was chosen arbitrarily. The system behavior at this point had some features of the desired goal, thanks to qualitative feedback provided by direct visualization in graphics mode. The observer can develop an intuitive feeling of the system responses and direct the search consequently.

Each parameter was modified in turn. This allowed to build a map of the environmental landscape around the chosen point. A final experiment was set up following the lines of steepest variations, as in a gradient descent algorithm. This experiment is successful, and the system quantitative and qualitative responses were tuned toward the desired goals.

The full gradient descent method in the environment landscape could theoretically be applied (with the risk of finding a local optima); In practice, computational power is currently a limiting factor.



## 6.2 Future work

This study is about the possibility of driving evolution on a global scale, without resorting to an explicit local fitness function. It does not pretend to be a full artificial life environment.

Future work could be done in this direction. In particular, the current agents are memoryless: they are purely reactive on their sensory inputs. The possible interactions are a limiting factor too. Extensions could include more elaborated relations [6]: agent to agent messages, resource trading, and stigmergy, for example. The agents' AI could also be improved. Speciation is another issue that could be investigated. It would be interesting to monitor the impact of these new possibilities on the global system response, in comparison with the current study.

A more general problem is to decide on the trade-off between model complexity and computational requirements. How simple can the system be and still retain essential features necessary for the goal to achieve? In the present case, a long-term sustainable setup could not be found despite the environment complexity. Is this due to limitations on the agents' AI, the nature of the world, and the interactions of the different parts of the system (including the spore clouds)? Or is there a region in environmental parameter space that fulfills the desired goals? Some results reflect the system limitations, but others are more difficult to interpret. To what extent are these results expected, and predictable? How could we design the system so as to induce a desired global property? The approach proposed in this study allows some control over the system behavior, but does not answer these questions fundamentally.

The real solution to this problem would be a comprehensive theory on multi-agent systems and the emergence of global phenomena. Progress is made in this direction [5], but the task is far from complete. Such a theory would include the identification of the necessary preconditions, for the emergence of a desired target behavior.

## Appendix: Software setup

This project source code is available under the GNU General Public License, and is hosted by an independent third-party revision control system. See the project page at <http://gna.org/projects/crogai/> for more information.

All random seeds and parameters used in this study are included under revision control. These experiments are reproducible with revision 17. Random seeds 1112159829, 1112206792, and 1112284094 were used for figure 3.

The Mersenne Twister random number generator, by Makoto Matsumoto, was used. See <http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/emt.html>.

The neighborhood query utility was improved and originally reused from the OpenSteer library, <http://opensteer.sourceforge.net/>.

The graphics framework is based on the Open Scene Graph library <http://www.openscenegraph.org/>.

All other algorithms and routines are my personal creation, including the simulator engine.

## Acknowledgments

I'm indebted to Peter Grogono for invaluable discussions, for pointing out speciation problems, for the idea of checking on inter-generations mix using the life time as metric, and for overall support. Some ideas of this project were inspired by the "FormAL" [4] framework.

My thanks goes also to Nawwaf Kharmas for general discussions, in which we came up with the idea of using a spore system for reproduction as an alternative to agents meeting.

Financial support was provided by the EADS Corporate Research Center, in cooperation with the French Ministry of Foreign Affairs.

## Bibliography

- [1] Jonathan M. Chase, Peter A. Abrams, James P. Grover, Sebastian Dieh, Peter Chesson, Robert D. Holt, Shane A. Richards, Roger M. Nisbet and Ted J. Case (2002) *The interaction between predation and competition: a review and synthesis*. Ecology Letters 5: 302–315
- [2] Debashish Chowdhury and Dietrich Stauffer (2004) *Computer simulations of history of life: speciation, emergence of complex species from simpler organisms, and extinctions*. Physica A, vol. 340, 685.
- [3] B. Derrida and D. Stauffer (1986), *Phase transitions in two dimensional Kauffman cellular automata*, Europhysics Letters, 2:739-745
- [4] Peter Grogono, GuoRong Chen, JunFeng Song, Tao Yang and Lei Zhao (2003), *Laws and Life*, 7th IASTED Conference on Artificial Intelligence and Soft Computing (ASC 2003), p158-163
- [5] John Holland (1998), *Emergence: From chaos to order*. Addison-Wesley books, ISBN 0-201-14943-5
- [6] Peter T. Hraber, Terry Jones, Stephanie Forrest (1997) *The Ecology of Echo*, Artificial Life 3: 165-190
- [7] Jef Huisman and Franz J. Weissing (2001) *Fundamental Unpredictability in Multispecies Competition*. In The American Naturalist, vol. 157, no 5.
- [8] Chris Langton (1990) *Computation at the edge of chaos: phase transitions and emergent computation*, Physica D, 42, 12
- [9] Elizaveta Pachepsky, Tim Taylor, Stephen Jones (2002) *Mutualism promotes diversity and stability in a simple artificial ecosystem*. Artificial Life 8: 5-24.
- [10] Craig Reynolds (1999) *Steering Behaviors For Autonomous Characters*, Game Developers Conference