MINIMEME: Of Life and Death in the Noosphere

Stéphane BURA

www.stephanebura.com

Abstract

This paper deals with artificial animals able to communicate beliefs about their environment's properties to each other. In order to study the relations existing between the information exchanged and the emergence of behaviors or organizations, we propose a model called MINIMEME. This model exploits Dawkins' paradigm stating that ideas, or *memes*, can be compared to parasites infecting their host and trying to duplicate themselves in other hosts' memories.

MINIMEME models the interactions occurring between the world of the memes and the animal societies. It is shown that some ecological concepts can be applied to memes, as if they were some kind of *animats*, and how these concepts relate to the simulated animals. The way this model can be used as a tool to help objectively qualifying emergent behaviors in simulated societies is then discussed.

1. Introduction

How does adaptive behavior in animal societies relate to the information exchanged between its members? What is the impact of a communication protocol on the organization of such societies? If we make the hypothesis that the animals we study have a very limited cognitive capacity, how do new rules governing the working of their societies emerge? That is, what can be considered a "good idea" by a society already following a set of rules and how does change occur?

What is a good idea? Rather that computing the efficiency of a given set of rules, we will adopt in this paper the approach proposed by Richard Dawkins in [Dawkins 76]. According to Dawkins, a good idea is a successful idea, an idea that is believed by a great number of individuals. These ideas, or *memes*, are reproduced (or reproduce themselves, according to Dawkins's anthropomorphic formulation[1]) from one memory to another, acting like parasites in their hosts.

This paper will propose a simple model of the interactions occurring between the world of the memes and the animal societies. We will see that we can apply some ecological concepts to memes, as if they were some *animats*, and examine the kind of relations existing between the different levels of the simulation. Then, we'll consider how this model can be a tool to help qualifying emergent behaviors in simulated societies.

2. MINIMEME

The MINIMEME system [Bura 93][Bura 94] is composed of two parts : the environment and the animats which may change with each simulation and the meme level, or noosphere, which keeps the same structure and rules in all the simulations.

2.1. The noosphere

Dawkins defines a meme as a transmittable cultural unit or an imitation unit. If we generalize, he says that any idea capable of transmitting itself from one person to another, of replicating itself, is a meme. In MINIMEME, we'll only consider memes that define a behavior. These memes can be "executed" by their hosts to produce an effect (e.g., movement, sustenance, nest building

[1] Even if this anthropomorphic approach has been much debated in the past, some of its terminology will be used in this paper for the sake of simplicity. Memes have no real *will* to reproduce themselves, their hosts just communicate them to new hosts, following strict rules.

techniques…). In order to be successful and continue to exist, a meme must satisfy three conditions:

- It must find at least one host, that is an individual who keeps it in his memory.
- As the meme defines a behavior, the execution of this behavior must not endanger the host's life, at least not before the meme has been able to reproduce itself.
- The meme must be able to resist the attack of concurrent memes in the meantime.

There are two kinds of concurrent memes for a given meme. Either a concurrent meme contains information pertaining to the same behavior as the attacked meme and attempts to replace it (because a host can't believe two incompatible memes), or the meme is about another behavior but takes up enough memory space to prevent the acquisition of new memes. For MINIMEME's animats, memory is a finite resource and each meme has a certain size, thus limiting the number of memes a given animat can hold.

The sum of the memories of all the animals in the system constitutes a space called the *noosphere* [Morin 91]. Memes inhabit the noosphere as animats inhabit the simulated environment.

2.2. How do memes evolve

To simulate the ability of the memes to conquer a part of the noosphere, we'll use two parameters for each meme : *change*, which is a measure of the meme's propensity to mutate or to succumb to other memes' attacks, and *proselytism*, which quantifies the meme's aggressiveness, i.e., the probability that it will try to reproduce itself. These parameters take real values in [0,1], a new meme receiving random values. Thus, a successful meme has, for all its instances in the noosphere, a high mean *proselytism* and a low mean *change*.

As we'll see later, it is noteworthy that these parameters don't take into account the ability of the meme to keep his host alive.

The change and proselytism parameters evolve according to simple rules. This evolution takes place at the end of a system cycle, which is described in details in the following section.

First, a *satisfaction* function is evaluated for each host. This function depends on the kind of simulation ran. It may involve an estimation of the correct accomplishment of a task, state variables in the host (is it hungry, ill…), constraints applying to the host, etc. If the host is satisfied, it increases the proselytism of each of its memes by 25% and decreases their change by the same amount. Conversely, if the host is not satisfied, it decreases its memes' proselytism and increases their change.

Then the memes may mutate and reproduce themselves. A mutation occurs when a random draw in [0,1] gives a number lower than the meme's change. The nature of this mutation is simulation dependent. Both change and proselytism are assigned random values for the new meme. If the meme didn't mutate and if another random draw is lower than its proselytism, replication may take place. A random number of individuals are chosen among the host's neighbors (i.e., the ones it can communicate with) and the meme is proposed to each of them. A potential host can resist taking the meme only if either it has not enough memory left or it already has an incompatible meme. In the latter case, a new random draw is made and if the result is higher than the attacked meme's change, it stays in the host's memory, repelling the attacking meme. If a meme tries to "infect" a host that already possesses it, the host's meme is reinforced (its change is decreased and its proselytism increased).

It is a easy way to model that satisfied hosts tend to hold on their ideas and spread them around, while unsatisfied hosts are more prone to change theirs or to accept new ones.

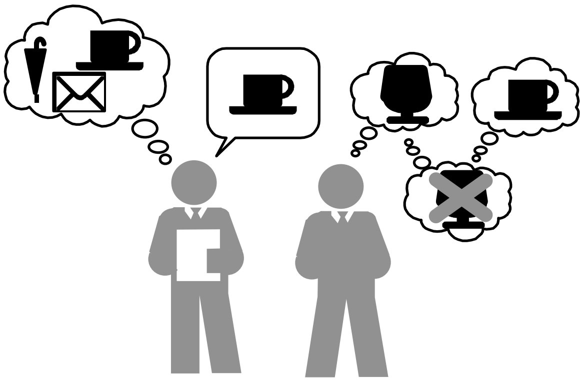

Figure 1 - Replication of the “Drink coffee” meme. As it is incompatible with the “Drink liqueur” meme (in this example), it replaces the attacked meme.

This mechanism governing the evolution of the memes is the same for all the simulations made with MINIMEME. The only characteristics to be defined for a given simulation are :

- The satisfaction criterion for the hosts;
- The nature of the mutations each meme can undergo;
- The “range” of the communication between hosts or, more precisely, how to find the “neighbors” of a given host (e.g., in the same room/cell, along a pheromone trail…). This range may be infinite if there are no limitations regarding communication.

This last characteristic is very important because of the relation between the two levels of the simulation. In this first version of MINIMEME, the hosts can learn new memes only by interacting with each other. The reproduction of memes is thus limited to a “conversion” process (as shown in figure 1). Many other ways of transferring a meme exist (imitation, coding a meme in the environment…) and they’ll be the object of future studies.

3. Grazers

The “grazers” system will help to understand the relationship between the population of animats and its noosphere and how these animats adapt their behavior to environmental constraints using memes.

As shown in the previous section as in [Dawkins 76] and [Morin 91], the evolution in the noosphere involves a positive feedback mechanism. The more instances a meme has in the noosphere, the better it is able to reproduce itself. As acquired behaviors can only be chosen among existing memes, the replication process will assure the durability of the dominant meme (or group of memes). Likewise, a mutated “deviant” meme (incompatible with the dominant memes) will find it very difficult to spread in the noosphere, because the majority memes reinforce themselves.

From a systemic point of view, it means that the noosphere stops evolving and has reached a dynamic equilibrium. The capacity to recognize such states is fundamental when one works with Artificial Life systems. The study of the noosphere’s population gives precious information concerning the system’s global behavior.

The grazers example has been chosen because it is simple, using few animats and only one family of memes, but nonetheless produces complex trajectories. Furthermore, it allows us to observe the relationship between animats’ survival and memes’. Lastly, the simulations made with this system converge very quickly toward a dynamic equilibrium (several hundred cycles at most).

3.1. Defining the grazers

This example deals with the emergence of a stable territorial distribution of animats subjected to various environmental constraints. As we’ll see this distribution depends on the stability of the noosphere.

While the rules governing the noosphere are immutable, the animats and their environment can change according to the kind of simulation chosen.

Here, the environment is made up of four identical territories having the same *carrying capacity*. The carrying capacity is the maximum number of *grazers* that can find sustenance in this territory during each cycle.

A grazer is a very basic animat that can’t do much. It can only move from one territory to another and communicate with the other grazers in the same territory. There are twelve grazers in the system, each being defined by its position, its energy and the contents of its memory.

The position of a grazer is one of the four territories. Even if the territories have a carrying capacity, there is no upper limit to the number of grazers a given territory can hold. The original position of a grazer is chosen at random.

The energy of a grazer is an integer from 0 to 5. A new-born grazer gets the maximum energy. This energy is decreased by 1 when the grazer can't eat and increased by the same amount (up to the maximum) when it finds food. Movement costs energy too (1 point if the grazer changes its position) as do predators attacks (2 points; cf. last experiment). When its energy reaches 0, a grazer dies. It is then replaced by a new grazer whose memory is initialized. This allows to keep the size of the noosphere constant[2]. A new-born grazer automatically receives a new meme, either learnt from one of its neighbors or randomly generated if the grazer is alone in its territory.

The memes used in this simulation all belong to the same family. They are all beliefs about the optimal density of grazers per territory. There are ten different memes corresponding to densities from one to ten grazers per territory. Each grazer knows one of these memes and only one (they are all incompatible). At the beginning of each cycle, a grazer checks if the density in his own territory is equal to the optimal value it seeks. If this isn't the case, the grazer moves to the best territory according to its meme. For instance, a grazer with a meme whose value is 5 will seek a territory populated by 4 other grazers. A meme is mutated by randomly increasing or decreasing its value for the optimal density by one.

It is worth noting that grazers do not arbitrarily favour less populated territories when they move. Again, a grazer believing that the best density is 5 would choose randomly between territories holding 3 or 5 other grazers if none held four. Apart from the rules governing the memes evolution, there are no cognitive apriorisms.

Lastly, a grazer can only communicate with other grazers in the same territory and it is satisfied only if its energy is at its maximum.

3.2. The system's cycle

The simulations use a discreet time, each cycle consisting of four phases.

• Action phase: Each animats executes its meme. For the grazers, it means checking the position's density and possibly moving. The grazers are sorted by increasing energy, so that the "fittest" act last. This way they suffer less from the perturbations caused by the movement of the other grazers.

• Environment phase: This phase is significant only during the last experiment involving predators (q.v.).

• Feeding phase: Now sorted by decreasing energy, the grazers eat. A territory can only feed as many grazers as its carrying capacity. If there are grazers in excess, only the stronger get to eat (to emphasize the effects of overpopulation, the food is not split between them all).

• Meme phase: The evolution of the noosphere takes place during this phase. Each grazer tests its satisfaction (is its energy at its maximum?) and possibly communicate with some of the other grazers in the same territory.

4. Experimenting with the grazers

Three experiments involving different environmental conditions have been run. For each experiment, a hundred simulations have been made. A simulation was stopped when the distribution of memes in the noosphere had stayed unchanged for 100 consecutive cycles[3] or if it had run for 1000 cycles. Most of the simulations (76%) lasted less than 300 cycles and only 2%

[2] Indeed, this choice has an impact on the working of our system. The death of hosts generally means the disappearance of the noosphere. The size of the studied population being so small, this rule (akin to some kind of birth control) is used. A future version of this simulation may use memes to regulate births over a larger population.

[3] Further experiments lasting several thousand cycles have shown that once a state of dynamic equilibrium had been reached in the noosphere, the system didn't evolve anymore. A stray "deviant" meme could appear in particular instances (because of the randomness of the process), but the system would then quickly return to its previous state (in a few cycles).

failed to yield a dynamic equilibrium before the 1000th cycle.

The parameters for the three experiments were :

• Just enough food : In this experiment, the carrying capacity of the territories is 3, which means that there is just enough food in the system to feed all the grazers (4 x 3 = 12). The optimal distribution for the grazers is three individuals per territory (3-3-3-3). This distribution is said to be optimal for the grazers because, after a certain time, it leads to the satisfaction of all of them.

• Too much food: The carrying capacity is raised to 4 for each territory. The environmental constraint being relaxed, an optimal distribution for the grazers consists of groups of zero to four individuals per territory.

• Too much food with predators: The carrying capacity is still 4, but territories holding less than four grazers during the *environment phase* of the cycle are attacked by predators. Each of the grazers in the attacked territories loses 2 energy points. There is only one optimal distribution consisting of four grazers per territory (4-4-4-0, one territory remaining empty).

The system being fully defined, we may try to predict its behavior[4]. Apparently, MINIMEME is controlled by a simple negative feedback loop : when a meme that is not adapted to the environmental constraints “infects” some hosts, their energy soon decreases and the meme mutates or is replaced by another meme (because its *change* increases greatly). Even if the meme resists, its hosts will die and, in the end, a more suited meme will take its place in the reinitialized memories. Thus this system seems bound to lead to the optimal distribution for the grazers.

4.1. Just enough food

The colonization of the noosphere by memes producing the optimal distribution can effectively be observed in some simulations (Figures 2 and 3).

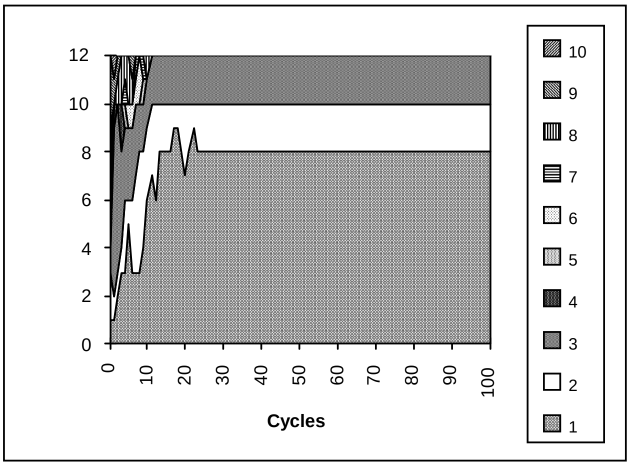

Figure 2 - Distribution of the memes in the noosphere. Just enough food. The carrying capacity is 3. (Example 1) The equilibrium is reached before the 30th cycle, lea-ving only “1”, “2” and “3” memes in the noosphere.

Figure 2 shows the distribution of the memes in the noosphere during the simulation. The vertical space allotted to each meme is a measure of how many grazers hold it in their memories. Thus, at the 50th cycle, the “1” meme has eight hosts and the “2” and “3” memes have two hosts each. The “X” meme means “The density in my territory should be X.”

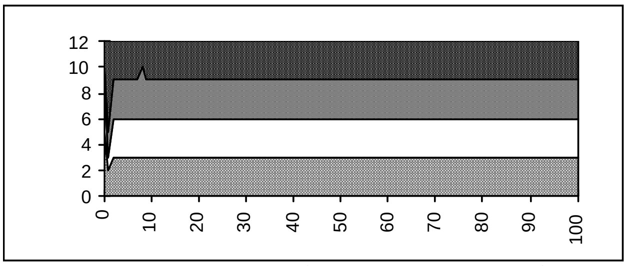

Figure 3 - Distribution of the population in the four territories. Just enough food. (Example 1) After the 10th cycle, each territory holds three grazers

Figure 3 shows the distribution of the grazers among the different territories (each color corresponding to a territory). For instance, at the 9th cycle (small peak), the first and second territories hold three grazers, the third four and the last two (a 4-3-3-2 distribution).

[4] According to Assad and Packard [Assad & Packard 92], identifying the degree of deductibility of a system is a mean to qualify emergent behavior.

Memes with a value higher than 3 disappear quickly from the noosphere. Because a grazer seeks the territory with the density closest to the value of its meme, the "1" and "2" memes have the same effect than the "3" meme, once the 3-3-3-3 distribution is reached. Grazers distribute themselves evenly in the four territories, they are all satisfied and they don't have to spend their energy to move. The mean change for all the memes diminishes quickly while the mean proselytism reaches its maximum (around the 30th cycle).

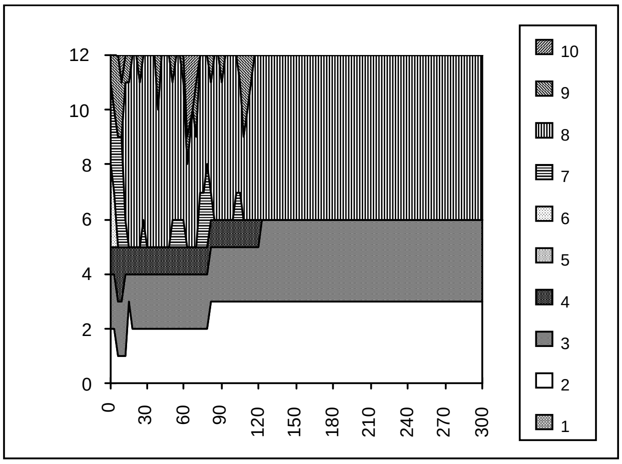

Figure 4 - Distribution of the memes in the noosphere. Just enough food. The carrying capacity is 3. (Example 2) Even if the "2" and "3" memes are still present, the "8" meme has six hosts.

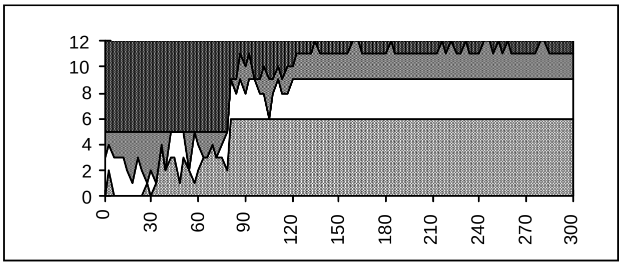

Figure 5 - Distribution of the population in the four territories. Just enough food. (Example 2)
The grazers distribution oscillates between 6-3-3-0 and 6-3-2-1.

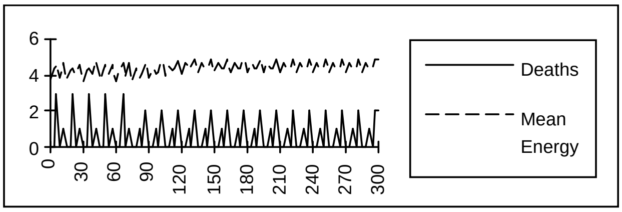

Figure 6 - Death rate (for each cycle) and mean energy. Just enough food. (Example 2)

Yet, in the same conditions, the grazers may fail to reach their optimal distribution. In the second example (Figure 4), even if the "2" and "3" memes are still present, curiously the "8" meme is persistent. This is curious because, as we can see in the figures 5 and 6, the distributions it causes (6-3-3-0 and 6-3-2-1) provokes an increase in the death rate of grazers. Nonetheless, some "8" memes have managed to get a low change and a high enough proselytism to survive and be replicated. What happens is the following self-catalytic phenomenon: memes with a high value provoke the gathering of their hosts. In such a milieu, they reinforce each other and their replication is made easier. **Thus, a meme that kills its host can survive in the noosphere and even become dominant.**

Of course, this requires special conditions (there are seven grazers in the same territory at the beginning of the simulation and a large proportion of high value memes), but it must nonetheless be taken into account.

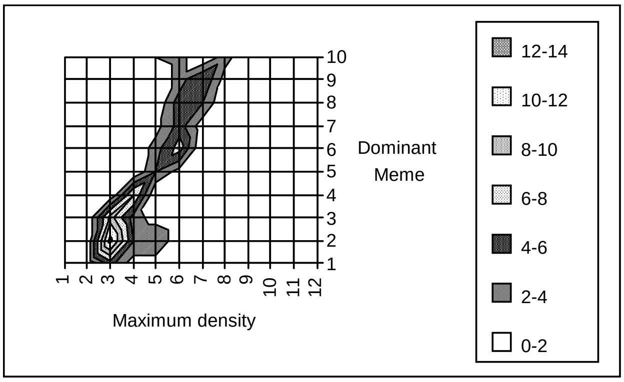

Figure 7 - "Just enough food" Synthesis. (100 simulations) The carrying capacity is 3. The values shown are the number of couples (density / meme) in all the simulations once equilibrium has been reached.

Figure 7 synthesizes the results of the 100 simulations made with these parameters. It shows the number of instances of each couple (dominant meme / maximum density). The dominant meme in a simulation is the one taking up the largest space in thenoosphere, once the equilibrium is reached. The maximum density is the number of grazers in the most populated territory, once the equilibrium has been noted. The reason for the size of this area is that some distributions of the grazers are note stable (as in example 2 above). Nonetheless, the peaks allow us to identify easily the possible states of dynamic equilibrium for this system and their relative frequencies.

In this first experiment, 45% of the simulations generate a maximum density of 3 or 4, 40% of 5 or 6, the last 15% giving higher densities.

4.2. Too much food

In this second experiment, the carrying capacity is raised to four. The excess food, diminishing the death rate, has two consequences. If the memes with a value smaller than 5 predominate in the noosphere, a stable optimum distribution is soon obtained (4-4-4-0 or 3-3-3-3 distributions). If this is not the case, the self-catalytic effect is accentuated as shown by figure 8. The synthesis (Figure 9) reveals that most of the simulations (66%) lead to an equilibrium situated far from the optimum distributions for the grazers. Relaxing the environmental constraints only speeds up the action of the high value memes.

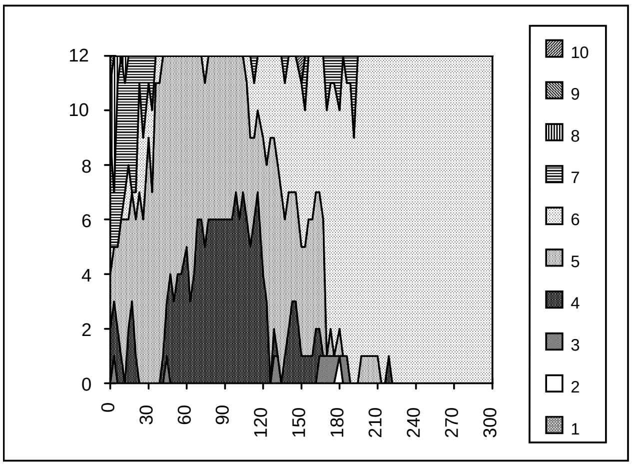

Figure 8 - Distribution of the memes in the noosphere. Too much food. The carrying capacity is 4. A majority of “5” and “6” memes soon produces a 6-6-0-0 distribution of the grazers.

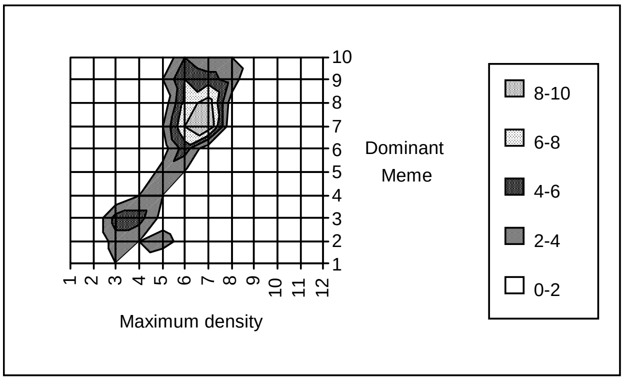

Figure 9 - “Too much food” Synthesis. (100 simulations) The carrying capacity is 4. Most of the results are far from the optimum densities (3 and 4).

4.3. Too much food with predators

In order to make life harder for both the grazers and the memes, predators are put into the system. As they attack the territories holding less than four grazers, the optimum distribution becomes 4-4-0-0. In the simulation described in figures 10 and 11, a high value meme (“8”) quickly overtakes the other memes. By the 100th cycle, it has conquered the noosphere and caused all the grazers to gather in one territory. But this meme is so “unfit” (the death rate is too high in the main territory and the few grazers that leave it are attacked by predators), that it is soon rejected by the system. The meme’s *change* raises rapidly, provoking its mutation, and its *proselytism* diminishes so that it is unable to stop his fall. This doesn’t mean that the “4” meme (which would lead to the optimum distribution) takes over. The perturbations caused by two “less unfit” memes (“5” and “6”) are not important enough to prevent them from surviving.

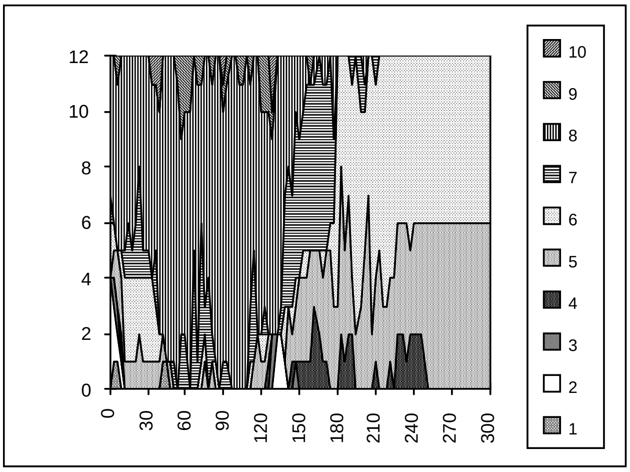

Figure 10 - Distribution of the memes in the noosphere. Too much food with predators. The carrying capacity is 4. Rise and fall of the “8” meme which fails to survive even after having conquered the whole noosphere.

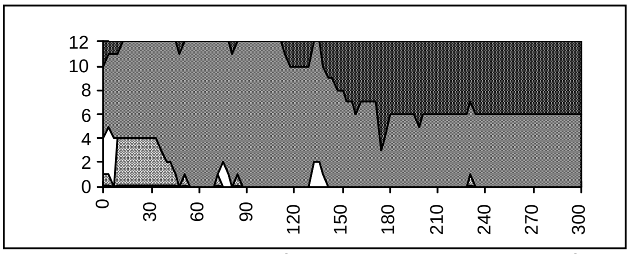

Figure 11 - Distribution of the population. Too much food with predators. A stable 6-6-0-0 distribution appears around the 240th cycle.

It is worth noting that the self-catalytic process observed in the previous experiments can be stopped. In fact, instead of a simple loop based on negative feedback, we have two intertwined processes in MINIMEME. As we’ve seen, memes’ duplication is essentially controlled by positive feedback and shapes the animats’ societies. But the environment, through the animats, can regulate the self-catalytic process, exposing “dangerous” memes.

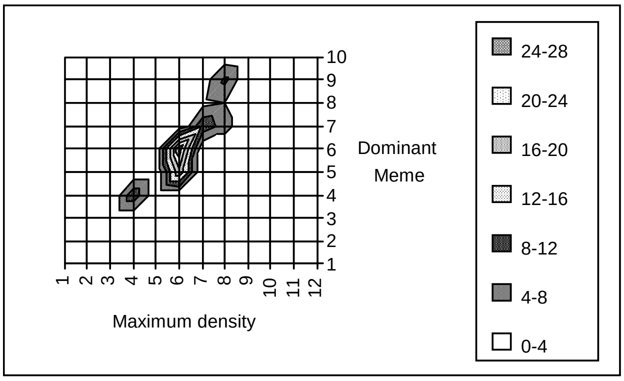

Figure 12 - “Too much food with predators” Synthesis. (100 simulations) The carrying capacity is 4. Notice the co-evolution of high value memes and the “4” meme (the absolute maximum density is eight).

Survival is nonetheless possible for killing memes (“9” in 10% of the simulations) as show in figure 12. However, it requires the parallel evolution of memes that keep them in check. In all the simulations of the peak in (density 8 / “9” meme), the last four grazers are hosts to the “4” meme. It is essential if they are to avoid the overpopulated territory, thus not upsetting its fragile balance. The co-evolution of two antagonistic memes allows the survival of one of them. This can be compared to the way a co-evolving parasite improves the evolution of a given species [Hillis 91].

5. Qualifying emergent behavior

Structures and global behaviors emerge from simulations, identified as macro-level properties of micro-level rules [Ferber & Drogoul 92][Bura & al 93]. Functionalities emerge from complex systems [Steels 92], as do functional dependencies or global properties [Bourgine & Varela 92]. It is hard to characterize emergence in Artificial Life systems when there are so many definitions for it. But it is generally accepted that emergence implies a certain degree of surprise, an intrinsic difficulty to predict the behavior of a system [Langton & al 92][Assad & Packard 92]. Does this apply to the system’s programmer or to a naïve observer? Can we only observe emergent behaviors in systems complex enough to confuse even their makers?

The main problem is objectivity for, most of the time, we know what we want to see emerge or what the system is capable of[5]. Much work has already been done concerning the objective qualification of emergent properties, either by looking for asymptotic functions [Steels 92] or identifying *clues* for cooperative behaviors [Miriad 92].

Using memes, as is MINIMEME, can take objective observation of a system a step further. In order to survive and become dominant, a meme must modify the hosts' environment so that its replication becomes easier. Conversely, we've seen that the environment constraints the kind of memes that may get a "niche" in the noosphere. To sum it up, **a successful meme is an idea about the system that became a property.** If the memes are varied and their possible combinations numerous, the evolution of the system should be hard to predict. Moreover, results don't need to be interpreted to find *what* has emerged as one has only to look at the composition of the noosphere. Thus, we can reasonably say that the behavior of the hosts, too, emerge as the memes transform the system and find niches in the noosphere.

Because the memes' evolution mechanism is very simple, it doesn't really compare with the way real ideas evolve. For instance, some ideas need only a small group of hosts and should decrease their proselytism once they've found it. Future work may involve "meta-memes" governing this evolution, as well as the notion of *schemes* or groups of memes that get replicated together.

References

[Assad & Packard 92] Andrew M. Assad & Norman H. Packard 1992. "Emergent colonization in an artificial ecology". In [Bourgine & Varela 92].

[Bourgine & Varela 92] Francisco J. Varela & Paul Bourgine 1992. "Toward a practice of autonomous systems", MIT Press.

[Bura 93] Stéphane Bura 1993. "MINIMEME, la mémoire collective d'un système multi-agents". Actes de la Journée Multi-Agents du PRC-IA, Montpellier, 1993.

[Bura 94] Stéphane Bura 1994. "De la vie et la mort dans la noosphère". Actes des Journées de Rochebrune, "Autonomie et interactions fonctionnelles". A paraître.

[Bura & al 93] Stéphane Bura, France Guérin-Pace, Hélène Mathian, Denise Pumain, Lena Sanders 1993. "Multi-agents systems and the dynamics of a settlement system". In "Proceedings of Simulating Societies '93" Nigel Gilbert.

[Dawkins 76] "The Selfish Gene", Richard Dawkins 1976, Oxford University Press.

[Ferber & Drogoul 92] Jacques Ferber & Alexis Drogoul 1992. "Using Reactive Multi-Agent Systems in Simulation and Problem Solving". In "Distributed Artificial Intelligence: Theory and Praxis" N.M. Avouris & Les Gasser.

[Hillis 91] Daniel W. Hillis, "Co-Evolving Parasites Improve Simulated Evolution as an Optimization Procedure". In [Langton & al 91].

[Langton & al 91] "Artificial Life II" Christopher G. Langton, Charles Taylor, J. Doyne Farmer, Steen Rasmussen, Addison-Wesley.

[Miriad 92] Miriad, "Approcher la Notion de Collectif", Actes de la Journée Multi-Agents du PRC-IA, Nancy, 1992. (Article Collectif de l'équipe MIRIAD, LAFORIA, Université Paris VI).

[Morin 91] Edgar Morin, "La méthode - 4. Les idées - Leur habitat, leur vie, leurs moeurs, leur organisation", 1991, Seuil.

[Steels 92] Luc Steels 1992. "Toward a theory of emergent functionality". In "Simulation of adaptive behavior: From animals to animats" Jean-Arcady Meyer & Stewart W. Wilson.

[5] This doesn't mean that collective sorting, for instance, doesn't "emerge" from robot ants' interactions. The surprise in this case comes from the lack of complexity in the animats. It's a kind of top-down approach to emergence.