Norm Selection Through Simulation in a Resource-Gathering Society

Daniel Villatoro
Jordi Sabater-Mir
Artificial Intelligence Research Institute (IIIA)
Spanish Scientific Research Council (CSIC)
Bellatera, Barcelona, Spain
email: {dvillatoro,jsabater}@iiia.csic.es

September 12, 2007

KEYWORDS

AI-supported simulation, AI in simulation, Intelligent simulation environments, Model design, Parameter Identification

ABSTRACT

In this work we present a mechanism designed for the selection of a suitable set of social norms that would regulate a simulated virtual society. The approach taken for this selection problem is a combination of Genetic algorithms and Simulation. A resource-gathering society has been built using Repast, and on top of it, the genetic algorithm that finds the social norms.

Introduction

Normative systems control the behaviour of the individuals in a certain society by means of norms. These norms, when enforced by the system, formally have the status of *laws* and is the system itself (by means of authorized representatives) who takes care of their observance and fulfilment. However *laws* are only a subset of the norms that work in a society of intelligent and autonomous entities. The system cannot control the behaviour of the individuals all the time and for all the situations and to fill this gap, *laws* are complemented by *social norms*. As Axelrod noticed Axelrod (1986), social norms and laws are mutually supporting: *"Social norms are often best at preventing numerous small defections where the cost of enforcement is low. Laws, on the other hand, often function best to prevent rare defections because substantial resources are available for enforcement"*.

A *social norm* is a rule that is socially enforced by the own non-institutional individuals that are part of the society. In other words, every individual is a potential norm enforcer for the others. The sanctioning applied to the breaking of *social norms* is guaranteed by some kind of social punishment like ostracism or decrease in the reputation instead of the more direct sanctioning associated to *laws*. The nature of *social norms* allow their effects to arrive where *laws* cannot and both mechanisms are complementary. An example of this in every day life is the *supermarket line*: imagine your common supermarket where you do your groceries every time you need a product. There exists a *law*, imposed by the institution and enforced by the cashiers, that tells us that we have to pay before leaving. However nothing is said by the institution about how the queue has to be formed in the cash desks. How have we (as intelligent agents) solved the absence of rules that should guide our behaviour in that situation? Faith in reciprocity have made us come out with a set of *social norms*: we have to follow a FIFO approach in the queue formation and every time you are carrying a heavy loaded cart, and some one behinds you in the queue just carries a light basket, you will let this person pass forward you, hoping that in the exact contrary situation (you carrying a light basket, and the person forward carrying a heavy loaded cart), you will be allowed to pass forward. People that do not share these *social norms* will not receive the benefits that it reports, and no one will let this person go first, because he did not allowed anyone before. Both sets of rules (*laws* and *social norms*) are complementary in order to make shopping in a supermarket a satisfactory experience for all the participants.

In this work, we use an evolutionary approach to find the set of norms (grounded in the idea of altruism and the notion of *image*) that works well in a scenario where autonomous agents have to survive given a limited number of resources. In order to able the testing of these norms in an scenario, we will make use of simulation techniques to build it.

The document is organized as follows. Section presents the scenario that we have chosen as the metaphor for our proof-of-concept. In section gives a deep explanation of all the computational tools that we will use. Section describes all the experiments that have been carried out with their obtained results. Then, in section we do a summary of the work that have been done extracting some conclusions. Last, but not least, in section we explain the future work planned to follow this research.

Description of the Scenario

In this section we will present the scenario we will use in the experiments. We have used a metaphor of a problem already treated by Paolucci et al., based on the methods a society develops for their survival and perpetuation through the exchange of resources.
In the work of Paolucci et al. Paolucci et al. (2006), the authors study the role of groups in the evolution of the concept of altruism. Using a real example from nature, they make simulations of a model of food sharing in a vampire bat society in order to reproduce the same behaviour observed in the real bats. Paolucci et al. show that the concept of altruism in this scenario is essential for the survival of the species as demonstrated by the study of real colonies of bats, and confirmed by the simulations. We adopt the notion of altruism also in our scenario that is based (with some modifications to fit our purpose) in the bats scenario. We are inspired by the "Altruist Vampire Bat Problem" metaphor to build an scenario where we can study the norm selection process. The scenario where we will test the norm selection process is the following:

"Find and Share" Game

Suppose a given society of agents which feed themselves by collecting resources. The resources are distributed randomly in the environment and are easily detectable by the agents. Once an agent finds a resource, it collects it automatically and in this way gains a certain amount of energy, that lengthens its life for a certain period of time and makes the agent seem fatter and bigger. The agents can detect if any other agent around has eaten or is hungry, because the obvious bigger size of the satisfied agent and the smaller size of the hungry one. The agents also have the ability of transferring some of their own energy to another agent.
If a couple of agents meet in the environment, they interact. When interacting, they have two options: transfer some energy to the other agent, or, they can decide to do nothing and walk away.
We suspect that resource distribution is correlated to the emergence of altruism when a certain parameter, such as the maximizing average life expectancy of agents or minimizing the number of death agents, want to be optimized. In other words, we expect that the set of norms becomes useful in very specific configurations of the scenario.
In our case, we will study the set of norms that extend the average life expectancy of the agents, and that, in the same way, minimizes the number of deaths in the society. To achieve this goal, we expect that the set of norms to be selected, by the norm selecting mechanism, promotes altruism and cooperation.
Summing up what have been explained so far: we will have n agents, located in a bi dimensional grid, where the agents will move every time step to an adjacent cell trying to find resources or another agents with whom interact. Every step an agent makes, it consumes one of its energy units, that are restored when resources are found. In case of an interaction with another agent, it can transfer a certain amount of energy. Therefore, a norm will be a set of combined observables, which will define a possible situation, with an action attached to that situation. The ethic code will be the set of all the norms.

Elements of the Simulation

In order to be able to represent the explained problem we need to use several techniques that have been combined in order to accomplish the desired experiments. In this section we will explain how the Repast model has been built, as well as we explain the structure and operation of the Image System. Next, as one of the most important parts of this research, we explain the Norm Selection process done with a genetic algorithm.

Population and Environment: a Repast Model

One of the basic ingredients for our experiment is the population that will assume and follow the norms, and the environment where they will interact. Repast is an Agent Based Simulation Platform that allows to quickly build complete models of simulation. All the needed tools such as graphical interface, plotting tools, and basic skeleton of a model of simulation, are provided.
The Repast model is based upon three main basic components:

- *Space*: The space has all the characteristics of the "physical" environment where the agents are to act.
- *Model*: The way in which how the world works, and how the simulation will be run are explicitly defined in the model.
- *Agents*: In this component of the Repast Model is where the behaviour and capabilities of the agent are explicitly defined.

Image System

Each of our agents is provided with a memory of "known" agents. Basically, each agent is able to keep track of the result of the interactions with all the other agents present in the simulation. Agents have an image ("own believed evaluation of a target" Conte and Paolucci (2002)) of each one of the other agents. In our particular example of the recollecting agents, we want the agents to know the moments that make the image of another agent increase or decrease: An agent x will increase the image of another agent y if y makes energy transference to x, enlarging then x life's expectancy. On the other hand, the image that an agent x has of another agent y will decrease in two situations:

1. When x, knowing y, transfers energy to y. This way we simulate a basic credit system of *image points*, and ensure that an agent that receives energy from another will have to "return" that energy back to the other to be even, in "image" terms.

2. When x (being starving) meets y (not being starving) and B_y does not offer an energy transference.

At the beginning, we have considered that being unknown is the same that being known with good image, in other words, at the beginning our agents trust.

The Norm Selection Process: a Genetic Approach

We choose an evolutionary approach to find the ethic code that suites better to the given society. We would like to remark that the process of norm selection is totally dependent to the configuration of the world. It seems obvious that a given set of rational agents will not behave in a small populated environment in the same way that they would do in a larger less populated environment, as well as in an environment rich of resources or in a poor environment. For that reason, we will maintain fix all the parameters of the simulation and then study how the norm selection process behaves when changing one of these parameters. Genetic algorithms seem to be the most suitable technique to find this ethic code.

Representation
One of the most important decisions to make before applying a genetic algorithm to a problem is to decide how a solution is represented. Traditionally, a chromosome is encoded as an array of bits where each bit represents a feature of the solution, a.k.a. gene. This boolean representation fits our case where each gene position represents a situation, which in our formal definition of the problem in section was defined. The genetic algorithm will assign a boolean value, that will correspond to the action to take in that situation (`false` means *do nothing*, `true` means *share food*). As a result of assigning to every different possible situation a certain action, we are defining all the norms that will be used.

Fitness function
Evaluation of chromosomes is as important as their representation. The fitness function evaluates a chromosome assigning it a fitness value that tells how good the solution is represented by the chromosome. In our specific problem, the chromosome will have, as genes, all the possible situations that a couple of agents could find themselves (defined by a combination of the observables described in section). The genetic operators will be in charge to change the action to take in a certain situation.
However we have to deal with a very important problem directly related to the nature of our scenario. As we are simulating, many random parameters are used, so, even the same simulation with the same individuals and parameters, executed twice, would obtain different results, similar but different. A non - deterministic fitness function decreases the effectiveness of the GA. To solve the problem of the non-determinism we have adopted the following solution: every time a chromosome is evaluated, the fitness value associated to it is the average of the results of 10 simulations. This value is stored and in case the same chromosome has to be evaluated we recover the value previously calculated. This approach, from our point of view, is an elegant solution to the problem of the non-determinism of the fitness function of the GA, and also, reduces the number of cases being evaluated, which would require higher computational times. In this way, we have provided to the basic GA the capability to deal with non deterministic fitness functions.

Cycle of evolution
The GA starts with an initial population of randomly generated chromosomes and enters into a cycle of evolution. In each iteration, a new generation is produced from within the old one by using evolutionary methods. The fittest chromosomes have more chances to be selected as parents and/or survive in the next generation. The loop is as follows: We first select a number of chromosomes from the existing population to be passed directly to the next generation (a.k.a. survivors). Then from the same population we select an even number of chromosomes to be used as parents. We apply crossover operation to each pair of parents to produce offspring chromosomes. Depending on the probability rate of mutation, we mutate a fraction of the newly produced offspring. Finally, together with the selected survivors, the offspring form the next generation for the cycle.
The evolution continues till a termination criterion is satisfied (i.e. satisfying a fitness threshold or exceeding a number of maximum generations allowed)
We also incorporate the concept of elitism in our GA by passing the best chromosome in the population to the next generation, thus we prevent losing the best solution found. We have observed that using elitism increases the convergence of our GA.

Genetic Operators

- **Selection:** We use the Fitness-Proportionate Selection (a.k.a. Roulette Wheel) method to select survivors and parents. In this method, the fitness value of the chromosome marks its probability to be selected.

- **Crossover:** There are a handful of crossover implementations that would work for us. We opted for 1-point crossover since it is the simplest and we have not reached any better convergence with other crossover implementations.
- **Mutation:**The mutation operation helps us to avoid falling into local optima in the search space. Mutation takes place by flipping the value of a gene in an offspring chromosome.

Experimental Setting and Results

In this chapter we present all the experiments that have been executed and analyzed with the platform described in subsection . In order to make easier the task for the reader, we have divided the adjustments of the parameters in three groups, parameters related to: the genetic algorithm, the world configuration, and the simulation.

Setting Up the Genetic Algorithm

The genetic algorithm has been fixed with a specific set of parameters DeJong and Spears (1990) which in the literature have been probed to suit in problems with similar characteristics: **Population size** = 50 ; **Number of generations** = 20. De Jong and Spears recommend to set the number of generations to 1000, which we did at the first tests. After observing that the convergence in the worst case was produced below 20th generation, we decided to fix it to 20. ; **Crossover type** = one point ; **Crossover rate** = 0.6 ; **Mutation rate** = 0.001. We have checked this parameters to work fine for our purpose obtaining satisfying results, that will be shown later.

Setting Up the World

As explained in subsection , we have developed an altered Repast model to combine the simulation power of Repast with the genetic algorithm. We have assumed that all the agents will behave in the same way, in other words, they will all follow the same set of norms.

Setting Up the Simulation

Directly related to the parameters of the world configuration, some other parameters have been defined to describe the simulation. All the experiments have been executed during 10000 steps of simulation, making all the agents follow the same set of rules.
In order to observe the effectiveness of the ethic code provided by the genetic algorithm, we have prefixed three different behaviours that we consider basic behaviours for the agents, and hence being able to compare those with the obtained by the Genetic Algorithm:

1. **Random**: Expected to be the worst performing behaviour, we allow the agents to choose a random action every time they meet another agent.
2. **Egoist**: Even though it can be considered just like a very general rule, this behaviour implies that no matter what the situation is, no agent will transfer energy to another.
3. **Sharing**: One agent will transfer energy to another, in case the second one is weaker than the first and the fact of transferring energy does not put the former on risk to die.

We have programmed the Repast model to return us several measures that will be useful for the analysis of the results from the simulation when any of the previously defined parameters are modified. Some of these measures are *Average Life Expectancy* of the agents, and the number of *Living Agents*.

Results

We set up three different experiments that have been run in order to test our initial hypothesis: a genetic norm selector is suitable to find an appropriate subset of norms in order to optimize a certain parameter, which in our scenario means minimizing the number of death agents in the simulation and therefore make the society perpetuates (in some experiments we might be interested instead in lengthen the average expected life of agents). At the very first experiments, we thought that most of the parameters were not important, and that the genetic norm selector would just perfectly work for all the situations. This assumption was wrong. We realized that in some configurations of the world, any of the behaviours (genetic, sharing, random or egoist) adopted by the agents do not make a difference, as they will all obtain similar results. The main characteristics of these types of environments are the following: rich in resources spread in small amounts, widely distributed all over the environment, and with a low transference rate. The reasons why in these kind of environments we will not be able to extract any interesting conclusion is because as the agents are likely to find resources easily, it is not really important for them to share any energy. Consequently, they will all behave equally bad. On the other hand, in resource lacking environments, in case they decided to share resources, this sharing would mean only a lengthen of their death throes, instead of helping this dying agent transforming him in a healthy one, with more opportunities to find resources that would lengthen its life.

Experiment 1: A Poor Resource Environment
In Figure 1, we have represented the number of surviving agents along the time from an initial population

ID	Action	Situation		
1	Do Nothing	You generous or Unknown	You Starving	Me Starving
2	Do Nothing	You Mean	You Starving	Me Starving
3	Do Nothing	You generous or Unknown	You Plenty	Me Starving
4	Transfer Energy	You Mean	You Plenty	Me Starving
5	Transfer Energy	You generous or Unknown	You Normal	Me Starving
6	Do Nothing	You Mean	You Normal	Me Starving
7	Transfer Energy	You generous or Unknown	You Starving	Me Plenty
8	Transfer Energy	You Mean	You Starving	Me Plenty
9	Do Nothing	You generous or Unknown	You Plenty	Me Plenty
10	Do Nothing	You Mean	You Plenty	Me Plenty
11	Transfer Energy	You generous or Unknown	You Normal	Me Plenty
12	Transfer Energy	You Mean	You Normal	Me Plenty
13	Do Nothing	You generous or Unknown	You Starving	Me Normal
14	Transfer Energy	You Mean	You Starving	Me Normal
15	Do Nothing	You generous or Unknown	You Plenty	Me Normal
16	Do Nothing	You Mean	You Plenty	Me Normal
17	Do Nothing	You generous or Unknown	You Normal	Me Normal
18	Do Nothing	You Mean	You Normal	Me Normal

Table 1: *Genetic Norms for 1st Experiment*

of 25 agents when adopting their different behavioural rules. At the initial state, the agents have 100 time steps of life (in case they do not find any resources which will lengthen their life). The simulation was preset with the following parameters: *Resources Appearance Rate*: 20 step; *Heap Number*: 5 Heaps; *Heap Size*: 1 unit of resources per heap; *Energy Transference Rate*: 25 units of energy. The scenario defined by this parameters represents a resource-lacking world, and also, the fact of an energy transfer do not give an important advantage to the "poor" agent when energy is transferred to them. Once the simulation was ready, we launched the Genetic Mechanism that would find the most appropriate set of norms. These norms can be seen in Table 1.

In this set of norms we do observe some reasonable norms. For example, norms 7, 8, 11, 12, 14 favour the cooperation when another agent is in a worse situation than the first, saving starving agents from death, and making normal agents plenty. Moreover, we observe that this norms are complemented with norms 1, 2, 3, 6, 9, 10, 15, 16, 17, 18 which can be basically resumed in the behaviour that an agent will not transfer energy to another when they are in equal conditions, or the first one is in a worse situation than the second. For the rest of norms, where we have obtained strange results (in the semantic meaning of the norms) we have an easy explanation for it. There is a law inherent to the agent that forgives to it to transfer energy when the agent do not have the amount of energy required for the transference, avoiding the suicides. Therefore, it does not influence the fact that norms 4 or 5 tells the agent to give resources, agents will be told by a superior authority (the environment) that the agents are not able to produce an energy

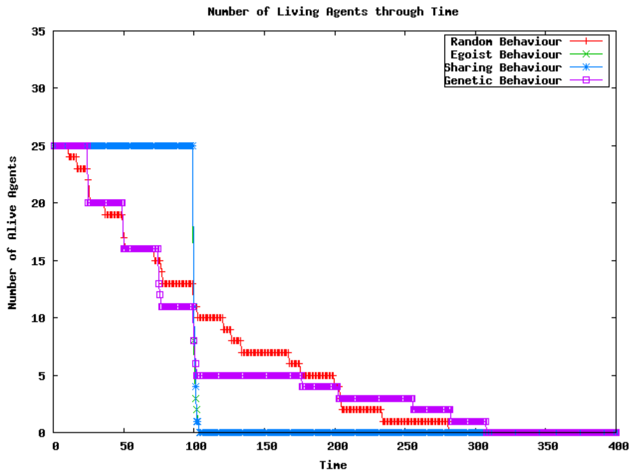

Figure 1: *Surviving Agents in a Poor Environment*

transference (Energy transference rate = 70) when they are starving, which means that their actual level of energy is below 25. As a consequence, in this specific scenario, an starving agent will always have his action imposed (*Do Nothing*) by the environment, without any consideration of the norms. We can observe the results of a comparison of the different behaviours in Figure 1.

We see that, as expected, the Genetic Behaviour (the behaviour that follows the norms established by the genetic algorithm) is the one that extends the society's life for the longest time. The explanation is that the genetic algorithm promotes the cooperation, even though if this cooperation puts in risk the agents life. Observing Fig 1, at the Genetic Behaviour, we see how the population decreases slowly, which means that some "sacrifices" (an agent donate energy even though it will leave itself with a very low level of energy, close to death) are being carried out, for the survival of the specie.
On the other hand, we observe that sharing and egoist behaviours make the society disappear around year (time step) 100, which means that agents were not able to find any resources. Moreover, due to how the sharing behaviour have been designed (sharing is only allowed if after the transference, the "source" agent does not become an starving agent), the sharing agents do have very low probability to share any energy: with an initial life of 100, a transference rate of 70, the starving level set to 25 and very low resources, agents will "not find any resources" and they will only have 5 time steps since the initial instant, to make a successful exchange. Hence we can say, that sharing and egoist in this scenario will behave in the same way. Surprisingly, we observe how the random behavioural agents do a better performance than the sharing or egoist, and the explanation for that is simple: in some situations, the random behaviour makes the agents behave close to the

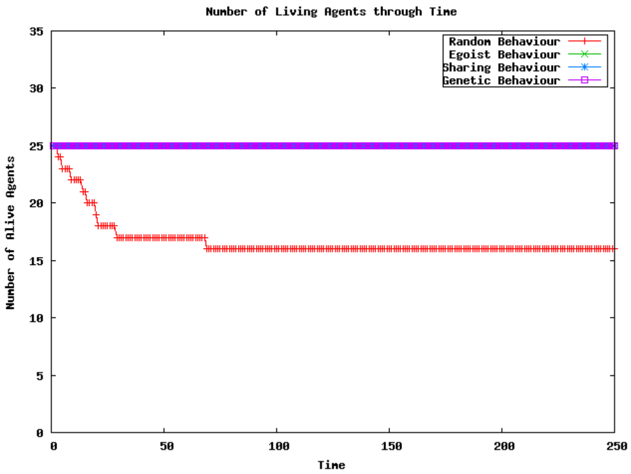

Figure 2: *Surviving Agents in a Rich Distributed Environment*

ID	Action	Situation		
1	Transfer Energy	You generous or Unknown	You Starving	Me Starving
2	Transfer Energy	You Mean	You Starving	Me Starving
3	Transfer Energy	You generous or Unknown	You Plenty	Me Starving
4	Do Nothing	You Mean	You Plenty	Me Starving
5	Transfer Energy	You generous or Unknown	You Normal	Me Starving
6	Transfer Energy	You Mean	You Normal	Me Starving
7	Transfer Energy	You generous or Unknown	You Starving	Me Plenty
8	Transfer Energy	You Mean	You Starving	Me Plenty
9	Do Nothing	You generous or Unknown	You Plenty	Me Plenty
10	Do Nothing	You Mean	You Plenty	Me Plenty
11	Transfer Energy	You generous or Unknown	You Normal	Me Plenty
12	Transfer Energy	You Mean	You Normal	Me Plenty
13	Transfer Energy	You generous or Unknown	You Starving	Me Normal
14	Do Nothing	You Mean	You Starving	Me Normal
15	Do Nothing	You generous or Unknown	You Plenty	Me Normal
16	Transfer Energy	You Mean	You Plenty	Me Normal
17	Do Nothing	You generous or Unknown	You Normal	Me Normal
18	Do Nothing	You Mean	You Normal	Me Normal

Table 2: *Genetic Norms for 3rd Experiment*

behaviour they would do in the genetic behaviour.

Experiment 2: A Rich Dispersed Environment

Another very different situation would be if the amount of resources in the environment would allow the specie to perpetuate. Therefore, to achieve this, we set up an experiment where the amount of resources would allow in some situations (depending how agents manage the resources) the specie to survive during the whole simulation. At this point of the research we realized how the distribution of resources would favour the emergence of a set of norms promoting the cooperation. Extending the experiment showed in Fig 1, we increased the amount of resources available to the agents in a simulation with the following parameters: *Resources Appearance Rate*: 1 step; *Heap Number*: 125 Heaps; *Heap Size*: 1 unit of resources per heap; *Energy Transference Rate*: 70 units of energy. This parameters makes our world to be well provided of resources, and the transference rate this time do make a difference when sharing by giving a real opportunity to the dying agent, since after the transference it state will change from *starving* to *normal*. In the same way than in the previous experiment, we launched the Genetic Algorithm to find for us the most appropriate set of norms for this scenario, returning us a set of norms. We do not consider that the analysis of the obtained rules have importance, since for the nature of the scenario, makes the fact of sharing or not sharing resources unimportant, as it can be seen in Figure 2. Nevertheless, the genetic mechanism has found a better performing behavioural norms than the random ones, equalizing the sharing and non sharing performances. We can observe the results in Figure 2.

Our suspicion about the relation of the resource distribution and the sharing importance becomes clearer after observing this figure. In such scenario, which is plenty of resources and very well distributed all over the space, agents have a high probability of finding one of these resources that will probably lengthen its life for the same time steps that took the agent to find the resources. Somehow, we have designed an independent-of-norms environment for the agents, where, no matter what they do (share or not), the perpetuation of their specie is ensured. Only in the case of random behaviours we have observed some instabilities and deaths, which we consider reasonable, due to the spontaneous and non-reasonable actions that an agent can decide; for example, a dying agent can meet a plenty agent and decide to transfer the other some energy, giving its last units of energy, remaining therefore in a very close situation to death.

Experiment 3: A Rich Concentrated Environment

Finally, and to confirm this suspicion, we decided to set up an scenario with the following characteristics: *Resources Appearance Rate*: 50 steps; *Heap Number*: 10 Heaps; *Heap Size*: 500 units of resources per heap; *Energy Transference Rate*: 70 units of energy. These parameters are defining a world where resources appears every large intervals of time (almost half the initial life of agents), and, very concentrated although in huge amounts; in other words, it is relatively hard for the agents to find resources, but when they find the resource, they find a huge amount of it. Likewise the two other experiments, the genetic algorithm this time returns us the set of norms in Table 2.

In this set of norms in Table 2 we do observe a promotion of the cooperation. Likewise in the Table

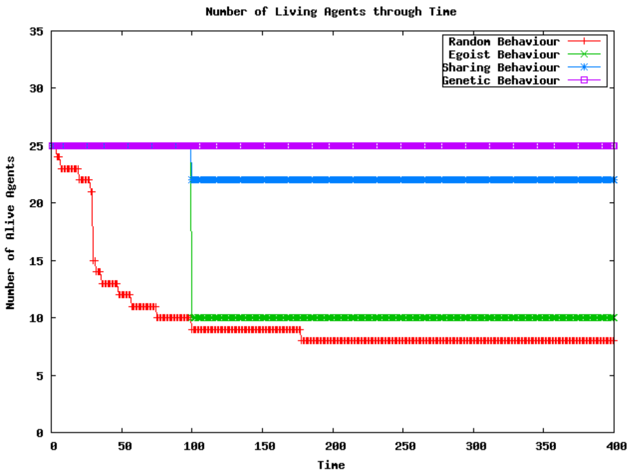

Figure 3: *Surviving Agents in a Rich Concentrated Environment*

1, we can not pay attention to the first 6 rules, that involves acting an starving agent, that, no matter what the norms say, the agent will never transfer energy in that situation, as we have the non-suicide implicit norm. Even so, the rest of norms "enforces" the cooperation. When an agent has a higher energy level than another one, the first one will make a transfer of energy, except in the situation where it can put its life in risk: if the first agent is normal and the second is starving, the first one will only do a transference of energy, if the first agent has a good image of the second. Apart from that, if two agents are in the same energy level, the norm says "do nothing".
The results of the comparison can be observed in the Figure 3

Finally, the suspicion is confirmed. We observe how this time the fact of sharing becomes really important in the society in order them to perpetuate. We observe how sharing and genetic norms make a similar performance, confirming our suspicion that sharing is crucial in some scenarios such as this one. The fact of sharing will make the agents survive for longer, having therefore more possibilities to find resources when those resources are spread again. The population of non sharing and random agents reduces, and the fact that the population reduces, makes that less agents would find food, making the society close to the collapse. It is easily noticeable in Figure 3 that the population in the four cases stabilizes after a time; this is due to after some time steps we suspect that the resources in the world are in equilibrium with the surviving agents, avoiding them to die.
To sum up, the experiments on our platform have allowed us to realize about the restrictions on the norm selection process. The restrictions are inherent to the problem to be simulated, in our case, the relation between the key parameters such as amount of agents, size of the world, food distribution, and food renewal rate. We have analyzed the three scenarios that we consider basic, extracting good results from our norm selection mechanism. Furthermore, we have found a type of scenario where the altruism and the cooperation are promoted by the norms in order to keep the society alive, avoiding the society's collapse.

Conclusions

Norm selection is a useful method in order to find a suitable set of norms that regulates a society in order to optimize a parameter. During the research we have realized that the way of representing the set of norms was a crucial decision. The chosen representation made that our norm selecting process was appropriate to handle with a genetic algorithm.
The integration of all the elements taking part in the norm selection process (Repast model and genetic algorithm) has a complexity we have to deal with. Repast, as a simulation platform, offers many functionalities that are useful during a simulation. However, when Repast is required just to return one value from a certain simulation, launch many different simulations and keep track of the values returned, the simulation platform is not capable. We have had to adapt some details in the structure and implementation of Repast to fully accomplish the task we wanted, having as a result what we could define as a "Genetic Configurable Repast Simulator".
Last but not least, we have extracted some conclusions from the performed experiments. We have been able to see that the cooperation and altruism only make sense in certain configurations of the scenario. These configurations basically are those where the resources are distributed in a very concentrated way with huge amounts and very distant in time intervals. In other scenarios, the altruism is not promoted or it performs in the same way that with an egoist behaviour, although the genetic norm selector still finds the set of norms that behaves better in every different scenario.

Future Work

In this subsection we present all the tasks that will be carried out to continue the work started in here, trying to find answers to questions that still have not been answered with this research, and to answer questions that have arisen during the work.

Brute Force Solution Even though we consider that the Genetic Algorithm is a suitable solution to the norm selection problem, and that we consider the solution given by the GA as a reference for what we will be obtaining with the learning, do we have any reference to compare the results of the GA more than our intuition?

As we do not have any, and do not trust much our intuition in this way, we will focus our efforts in the short term in comparing and analyzing the results obtained by the Brute Force algorithm. The Brute Force algorithm will go through all the possible ethic codes evaluating their functionality and returning the best civil code that optimizes a certain parameter. In this way we could see how far our Genetic Approach was from the global maximum. The Brute Force approach is only taken due to its viability in this simple scenario (as in a bigger one it could not be possible), and to confirm the results with the genetic algorithm.

Different type of agents Another experiment that we want to design is the configuration of an environment were heterogenous agents live, and see how the proportion of normative agents (those that follows the norms) against non normative agents could affect to the variation of the parameter being studied, in our case, the perpetuation of the specie. As we have designed the norms, we give to the agents the opportunity to develop some kind of social exclusion by recognizing their images. Therefore, by introducing heterogenous agents we expect some kind of exclusion emerge for the benefit of one of the agent's groups.

Deep study on the relation of the parameters During this research and the process of experimentation, we have realized that setting up the parameters was a complicated task, due to the direct relation that exists between them when setting up an scenario. For example, it is obvious that the number of agents and the size of the world would determine the density of the population; the density of the population and the amount of food distributed are also related and are key factors for the simulations. Once we find the relation that exists between these parameters, through an exhaustive study of the space of values of those, we will be able to extract better-founded conclusions.

Re-implementation in a different platform The whole process of the norm selection mechanism has been probed to work correctly in a Repast model (Swarm and Java based). In order to be able to affirm the utility of it, we would like to test the same procedure but in different simulation platforms, and if possible, in different programming languages.

Applications

The application domain of this research is directly related to an ongoing research which is carried out by a group of archaeologists. We are presented a non - prehistoric society, already extinguished, known as "the Yámanas". This society was located in Southern Argentina and are one of the groups of the societies commonly known as "canoeros". They lived there for around 6000 years in a very hostile environment. The main success, and reason of study, of this peculiar society is their capacity of auto-organization: the Yámanas were able to auto-organize themselves as a hunter-gatherer society and survive around 6000 years. The archaeologists consider as the hypothesis that the key of the success of this society was due to their strong respect to a known set of social norms. The archaeologists think, that through the use of agent based simulation, they could study the effect of these social norms in two different ways. First of all they are interested in the study of how these social norms were created, which in our research would represent the norm emergence problem. Secondly, they want to study the necessity of all the social norms that once existed, by checking how would the society had worked without one of those norms, and therefore, find what Tennenholtz Fitoussi and Tennenholtz (1998) already called the *"Minimal Social Laws"*, which somehow would correspond to our norm selection mechanism, adding the functionality of selecting it minimal. We find this collaboration work really interesting because it gives us a real domain where to prove our research.

Acknowledgments

This work was supported by the European Community under the FP6 programme (eRep project CIT5-028575 and OpenKnowledge project FP6-027253) and by the project Autonomic Electronic Institutions (TIN2006-15662-C02-01). Jordi Sabater-Mir enjoys a Ramon y Cajal contract from the Spanish Government.

REFERENCES

Axelrod R., 1986. *An Evolutionary Approach to Norms. The American Political Science Review*, 80, no. 4, 1095–1111.

Conte R. and Paolucci M., 2002. *Reputation in artificial societies: Social beliefs for social order*. Kluwer Academic Publishers.

DeJong K.A. and Spears W.M., 1990. *An Analysis of the Interacting Roles of Population Size and Crossover. Proceedings of the International Workshop Parallel Problem Solving from Nature*, 38–47.

Fitoussi D. and Tennenholtz M., 1998. *Minimal social laws.* In *Proceedings of the fifteenth national, tenth conference on Artificial intelligence, Innovative applications of artificial intelligence.* 26 – 31.

Paolucci M.; Conte R.; and Tosto G.D., 2006. *A Model of Social Organization and the Evolution of Food Sharing in Vampire Bats. Adaptive Behavior*, 41, no. 3, 223–239.