On the Emergence of Possession Norms in Agent Societies

Felix Flentge[1] and **Daniel Polani**[1,2] and **Thomas Uthmann**[1]

[1] Institut für Informatik, Johannes Gutenberg-Universität, D-55099 Mainz, Germany
[2]Institut für Neuro- und Bioinformatik, Medizinische Universität Lübeck, D-23569 Lübeck, Germany

Abstract

Our paper studies the emergence of social norms and their subsequent influence in a simulated society of artificial agents. These norms and their propagation in a society is put in close relation with the meme concept of Dawkins. In particular, the norms studied in this paper are concerned with the possession of goods. Here a global norm regarding possession of goods may result from the collective dynamics of the society and arise from purely local agent interactions. The role of sanctions and costs of enforcing sanctions and their relationship to the establishment of a possession norm are studied.

Introduction

An important part of the research in the field of Artificial Life is directed towards the modeling and simulation of fundamental properties of living beings. Because of their high complexity, human societies which are the topic central to social sciences pose a particular challenge to this kind of modeling. A simulation of qualitative properties of such systems can help to test and to improve models and theories from social sciences. In turn, these theories can be useful in constructing multiagent systems, which poses difficulties typically found in human societies.

A prominent aspect of human social systems are social norms. They have various functions, like the control of cooperation and the reduction of aggression. Perhaps the most important aspect of norms is that they make the individuals' behavior more predictable. By determining the behavior of an individual in a given situation, norms give other individuals the possibility to adapt themselves to the situation more easily and to act with respect to the expected behavior.

Thus norms are extremely important factors in reducing the complexity of social situations, which is of particular relevance for individuals that have typically limited resources. Furthermore, since the validity of norms extends beyond the concrete situation, they allow a coordinated behavior without tedious agreement or coordination protocols.

The goal of our paper is on the one hand the simulation of the emergence of norms and, on the other hand, the study of the effects of a norm on a society. The norms result from the dynamics of an agent society where the agents have the possibility to follow or not to follow a certain norm and to sanction or not to sanction behavior that deviates from the norm.

In existing approaches for the simulation of norms two separate aspects are studied: either the possibility that a norm emerges or not (Axelrod, 1986; Coleman, 1986/87) or the effects of an explicitly given norm (Conte/Castelfranchi, 1995; Saam/Harrer, 1999). Our model combines both approaches: it remains open whether a norm emerges or not and, in addition, the effects of the norm on the agent society are studied.

First, we will create a connection between norms and the concept of memes as put forward by Dawkins (Dawkins, 1996) that shall play an important role in our agent society model. The model itself is an extension of the sugarscape model of Epstein and Axtell (Epstein/Axtell, 1996) and will be described in the subsequent section. We will present our results and give some conclusions.

Social Norms and Memes

A *norm* can be considered as a general behavior code that is more or less compulsory and whose nonobservance is punished by sanctions. Norms often result from regularities in behavior and are enforced by institutions only after some time. We shall see that such an "unscheduled" emergence of norms can be smoothly combined with Dawkins' concept of memes. Dawkins suggests to apply the theory of evolution to cultural tradition.

Dawkins calls the "building blocks" of cultural tradition *memes* (which correspond to the *genes*, the building blocks of genetic information transfer). Memes thus encode cultural traits. They propagate from brain to brain by imitation. Behaviors are imitated and ideas are adopted. Some memes can propagate better than others and establish themselves with time. Many memes are passed on by education. Children assume many of their parents' behavioral patterns. Genetic and memetic inheritance is tightly coupled. The main difference between them, however, is that the memes can change during life while the genes can not. Norms in this sense can also be regarded as memes. They can propagate from individual to individual and are connected with certain behaviors. An individual follows a certain norm if it carries the corresponding meme.

The Model

The norm which we model in our studies is a norm concerned with the possession of goods, in short a *possession norm*. We view possession norms as regular behaviors that relate to the act of acquisition, to the state of possession itself and particularly to the respect of the possessions of others. Our model allows the development of possession norms for the individual agents. Our concept of possession relates to the possession of a good (a plot of land) producing a good important for survival (food).

We will say that a possession norm has developed in the agent society if a large number of agents has a behavior code respecting possessions claimed by others and if non-observance of this code is sanctioned. Possession does not exist *a priori*, and can only persist if enough individuals behave appropriately and respect it. We chose the *sugarscape* model from Epstein/Axtell, 1996 as a basis for our model since the concept includes a landscape that produces food (sugar) and it allows an extension towards the simulation of norm development in a natural way.

The Sugarscape Model

The sugarscape model by Epstein and Axtell consists of a landscape that provides sugar and agents which feed upon this sugar and can move around the landscape. In our simulations the landscape consists of a 50×50 grid of individual cells with periodic boundary conditions (torus).

For the agents and the cells of the landscape certain rules are defined which describe the interaction between the agents and between agents and cells. Each such rule can be switched on and off individually, being valid for all agents or cells. Each agent is activated once per simulation step and performs the rules valid for agents. The agents are activated in a random order each step. First the agent and then the cells are updated, according to the corresponding rules.

Each cell has a maximum sugar capacity and a current sugar level. The sugar capacity is chosen in such a way that one obtains a landscape with two central regions with the maximum capacity, around which the sugar capacity slowly drops. A cell can be *occupied* by a single agent at most.

For the landscape cells the *sugarscape growback rule* G_α is valid: in each simulation step the sugar content of a cell grows by α units until it reaches the sugar capacity of the cell (α is an integer).

The agents have a certain metabolism. It is given by an integer value and determines the number of sugar units used up from an agent's internal sugar level per simulation step. Each agent has a certain *vision range* which determines how far an agent can see in the four main directions of the landscape grid (n, s, e, w). Agents cannot see in diagonal directions. The vision range is fixed at the creation of an agent and chosen from an externally given interval.

At its turn, an agent can move to an arbitrary unoccupied cell inside its vision range and collect all the sugar on this cell. It then uses up the amount of sugar given by its metabolism. Agents die if they have no more sugar or when reaching a certain maximum age determined at their creation (a random value from a given interval). Agents perform the *Agent Movement Rule M* every simulation step.

- *Agent Movement Rule M*
  - search the unoccupied cells in vision range with a maximum current sugar level
  - if several such cells are present, choose the closest
  - go to this cell
  - collect all sugar in the cell

Interactions between agents can only take place between direct neighbors in the four main grid directions. Agents can reproduce. There are two types of agents, male and female. The agent sex is determined randomly with the probability of 1/2 at the creation (birth) of an agent. An agent is fertile if its age is in a certain age interval and if the agent has at least as much sugar as it had at its birth. Fertile agents perform the *Agent Sex Rule S*.

- *Agent Sex Rule S*
  - select a neighboring agent randomly
  - if this neighbor is fertile, if the neighbor belongs to the opposite sex, and if at least one of the agents has an unoccupied neighboring cell, a new agent is created and placed on this cell
  - if the sugar level of the current agent is high enough to reproduce again, and if there are other neighbors that were not selected this update, repeat rule S for another neighbor

The new agent inherits metabolism and vision range with a probability each of 1/2 from one of both parents. From each parent the agent obtains half of the amount of sugar that the respective parent possessed at its own birth. If an agent is "wealthy" enough, it can reproduce several times per simulation step.

Cultural Tags and Memes

Each agent has several *cultural tags* which can assume one of two different values each. The number of tags is the same for all agents. The tags represent the cultural attributes of agents. After moving, each agent performs the *Cultural Transmission Rule K*[1].

- *Cultural Transmission Rule K*
  - for each neighbor choose randomly one of its tags
  - set the neighbor's tag to the same value as the current agent's tag.

In our extended model, these tags play an extremely important role and will determine agent behavior in a very direct way. Therefore from now on we will denote these tags as *memes*. Rule K then describes the dynamics of meme propagation via imitation. Memes are inherited by children from their parents at birth and thereby simulate the influence of the parents on their children via education. Later the parents may adopt some memes from their children via rule K, though.

Model Extension towards Norm Simulation

To study the emergence and the impact of possession norms, we extend and partly modify the sugarscape model. In our model, we give the agents the option to acquire resources, namely a "plot of land" (i.e. cells). For this purpose, agents can mark cells. Only unmarked cells can be marked and only if occupied by the agent that wants to mark it. Cell marks can be seen by an agent if the cell is in its vision range. On death of an agent, all its cell marks are deleted.

An important aspect of the model is that marking a cell does not automatically mean that this cell becomes the property of the agent. For this to happen, it is necessary that also other agents respect the cell as being in possession of the first agent and behave accordingly. Here the meme model plays a central role. The first two memes of our agents determine their behavior[2].

The first meme (the *possession meme*) determines the behavior of an agent with respect to marked cells. Agents with an active possession meme consider marked cells as zero sugar cells and do not collect sugar from these cells. Under rule M they will not move to these cells unless all other cells in their vision range are zero sugar cells. Cells are only marked by agents with an active possession meme.

The second (*sanction*) meme determines the behavior when norm violations are observed. Norm violations are only registered by agents with an active possession meme and inside their range of vision. If an agent has an active sanction meme, it will sanction all norm violations observed by it. On being sanctioned, the "wrongdoer", i.e. the norm violating agent, loses a certain amount of sugar from its internal sugar level. The "avenger", i.e. the sanctioning agent, also incurs a cost and has to pay a certain amount of sugar.

These model extensions introduce a necessity to change the movement rule M since otherwise the agents' rate of reproduction is too low. The reason for this is that agents often have no incentive to move to a cell neighboring a possible mating partner because these are often already marked by some other agent. Therefore we introduce the movement rule *MF*. With this rule, agents still try to maximize the amount of sugar they collect, but their first priority is reproduction.

- *Movement Rule MF*
  - if current agent is not fertile, perform movement rule M
  - else search all unoccupied cells in vision range neighboring a fertile agent of the opposite sex
  - if no such cell is found, perform rule M, else go to a cell with a maximum amount of sugar
  - if several such cells are present, choose the closest
  - collect the sugar of the cell

In our extended model, the landscape of the original model is modified in such a way that the sugar capacity of the original landscape (ranging from 0 to 4) is multiplied by the factor of 10. This allows larger and finer variations of the values for metabolism in combination with the growback rule G_α. Among other aspects this enables us to prevent that agents with a (minimal) metabolism of 1 can survive by just sticking to a cell where in every simulation step at least a single unit of sugar grows back.

An important change w.r.t. the original model is to start the simulation with a random agent age between zero and the maximum age. Otherwise one has a very unnatural age distribution at the beginning of the run and has to wait for a large number of generations until this transient behavior normalizes.

Furthermore, in the original work metabolism and vision range evolve by time. Instead, we vary these values systematically from run to run, but keep them fixed during each individual run. This enables us to filter out the additional effects introduced by the variations of these parameters during the runs and to interpret our results more easily.

Results

In our simulations we keep one part of the parameters fixed and the other one varies systematically within cer-

[1] In the original work (Epstein/Axtell, 1996), K denotes a combination of two rules. Here we only focus on the Cultural Transmission Rule which we denote by K for simplicity.

[2] Other memes exist, following the original sugarscape model, but they are not used here and only influence the

on the values Epstein and Axtell have chosen for their simulations. So the sugar level of an agent at the time of birth lies between 50 and 100 sugar units and the maximum age between 60 and 100. As mentioned above the birth age of the first agent generation (at the beginning of a simulation) is distributed uniformly within this interval. The fertility for *males* starts between the age of 12 and 15 and ends between the age of 50 and 60. For *females*, fertility also begins between the age of 12 and 15, it ends between the age of 40 and 50.

The agents have 11 memes of whom only the first two are explicitly used in our scenario. Since only one meme at a time is affected by the interactions between agents, the probability of involving one of the two memes is given by 1/11. The first meme (the possession meme) is responsible for the possession norm, the second meme (the sanction meme) for carrying out sanctions against norm violations. The rules MF, F, G_1, and K are active. At the beginning of a run, 400 agents were placed in the landscape.

Epstein and Axtell used their model only to demonstrate the occurrence of certain phenomena and therefore performed just a single run for each parameter set. In contrast to this, to attain a broad statistical basis, our experiments consist of 100 runs for each parameter set, each with 2000 steps (updates). Metabolism rate and range of vision are varied systematically in double steps each from two to ten. So it is possible to analyze the enforcement and the effects of norms under conditions differently favorable for the agents. The chosen landscape is the original landscape with the tenfold sugar capacity for each cell.

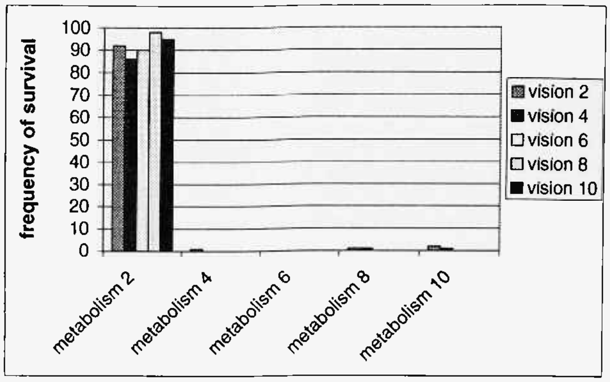

Figure 1: Frequency of survival without possession norm. The figure shows the number of runs in which the agent population survived 2000 steps versus various metabolism rates and vision ranges

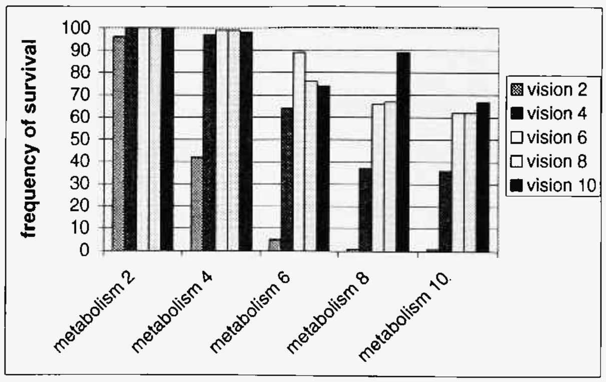

Figure 2: Frequency of survival with possession norm

Effects of Possession Norms

First we intended to clarify the effects of the possession norm "do not collect sugar on someone else's cells". For this purpose, we compare two runs: in one run the whole population has an active possession meme, in the other run the possession meme is deactivated for the whole population. The respective meme state does not change throughout the entire run, as at birth the agents adopt the memes of their parents and memes are always transmitted without defect.

The Figs. 1 and 2 show the number (frequency) of runs for which the agent populations survived the 2000 simulation steps. In all other runs the entire population became extinct. Agents respecting the possession norm can survive much better under unfavorable conditions (high rate of metabolism and small range of vision) as agents not aware of the norm. Without possession norm the agents in fact only can survive with a metabolism rate of two. The range of vision has almost no effect on the survival frequency. In the case a possession norm exists the range of vision has a great impact on the survival frequency because it is of high importance to have a far-reaching range of vision, since no sugar can be collected in many fields as they belong to other agents. The reason for the higher survival probability is that agents can usually collect more sugar per step in this case. If they possess several fields, they can wait until it is really profitable to collect sugar from the corresponding fields. In such a way, for an agent with a metabolism rate of two a possession of two fields is sufficient. It can shuttle back and forth between these two fields and collect two pieces of sugar each time.

Establishment of Possession Norms

Evolution Dynamics without Sanctions An established possession norm is a definitive advantage for the survival of an entire agent population. Thus the question arises whether a possession norm can establish itself with only genetic and memetic evolution as its driving force. For this purpose, we perform runs where only half of the agents have an active possession meme in the initial population.

The result is: the survival frequency in this case is just as bad as if no possession norm exists. A closer analysis of the individual runs shows that the meme always disappears first and then the population becomes extinct. It turns out that the agents without possession norm can reproduce significantly better on the short term, though ultimately leading the population to extinction.

Introduction of Sanctions In order to get the agents to obey norms, sanctions for norm violations have to be introduced. Fig. 3 shows the survival frequency for a punishment of four and sanction costs of zero. At the beginning, the possession meme was active in half of the population and the sanction meme in the entire population. Therefore, all agents respecting the norm sanction its violation.

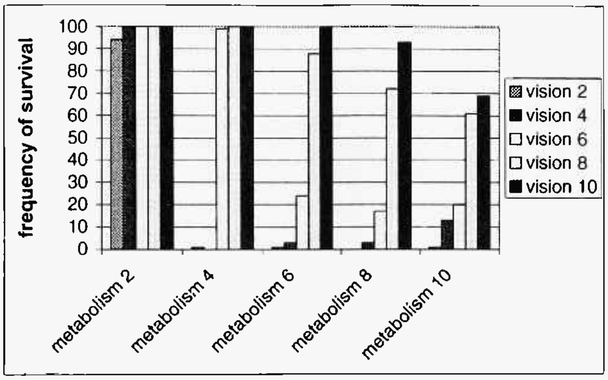

Figure 3: Frequency of survival with 50% possession meme, 100% sanction meme and a punishment of four

In this case the survival frequency is considerably better than without sanctions. Particularly for large values of the vision range very good values for the survival frequency are achieved. Because at high ranges of vision more agents can observe a norm violation and every agent knowing the norm is sanctioning a violation, overall punishment is considerably higher than at small ranges of vision. All in all it can be stated that with an increasing level of punishment the survival frequency of the population increases as well, even for small ranges of vision.

Therefore, the norm can assert itself quite well with sanctions. However, this is only valid as long as sanctions are not combined with costs. Further simulations show that, if costs exist, the norm cannot be asserted that easily because the sanction meme now is combined with disadvantages for its owner. The sanctioning agent now has to bear the costs for the sanctions, without obtaining an immediate advantage. Thus for increasing costs the percentage of the sanction meme decreases and, in turn, the probability of the possession meme to assert itself is reduced.

Conclusion

The simulation runs showed the emergence of possession norms under certain conditions. We speak of an established norm if a large number of agents behaves correspondingly. Thus, it represents an emergent quality of the agent society. Furthermore, the possession norm has been proven extremely useful for the survival of the agent population. Without the possession norm the agents become extinct much earlier in the case of unfavorable conditions than with possession norm.

However, the dilemma consists in the fact that there is a short-term advantage to agents ignoring the possession norm. The short-term interest of the agents to collect as much sugar as possible reduces the chances for the society to survive. An advantage for the individual results in a disadvantage for the society. Sanctions can offer a way out from this dilemma by reducing the short-term advantages. However, a new problem arises if sanctions are combined with costs and an individual may therefore avoid to set up sanctions. In real societies this problem is often solved by *institutions* which guarantee the obedience to certain norms and which are responsible for imposing sanctions.

In this case, costs have still to be paid. But institutions can exert a pressure on all individuals to share these costs. Thus, those may be kept lower on average and it reduces the probability that an individual can profit from the norm without contributing the cost necessary for its support. However, there is no guarantee that an institution will indeed be able to gather the relevant support from the individuals. For the future, it is intended to simulate the evolutionary emergence of institutions.

References

R. Axelrod: An Evolutionary Approach To Norms, in: American Political Science Review, Vol. 80 No. 4, December 1986, pp. 1095-1111.

J. S. Coleman: The Emergence of Norms in Varying Social Structures, in: Angewandte Sozialforschung, Jg. 14,1 , 1986/87, pp. 17-30.

R. Conte, C. Castelfranchi: Understanding the functions of norms in social groups through simulation, in: N. Gilbert, R. Conte (eds.): Artificial Societies. The computer simulation of social life, London, 1995, pp. 252-267.

R. Dawkins: The Selfish Gene, Oxford University Press, 1989.

J. M. Epstein, R. Axtell, Growing Artificial Societies, Cambridge, MA, 1996.

N. J. Saam, A. Harrer: Simulating Norms, Social Inequality, and Functional Change in Artificial Societies, in: Journal of Artificial Societies and Social Simulation vol. 2, no. 1, www.soc.surrey.ac.uk/JASSS/2/1/2.html, 1999.