Investigating the Mechanisms Underlying Cooperation in Viscous Population Multi-Agent Systems

James A. R. Marshall and Jonathan E. Rowe

Dept. of Computer and Information Sciences
De Montfort University
Milton Keynes MK7 6HP
UNITED KINGDOM
{jmarshall | jrowe}@dmu.ac.uk

Abstract

This paper presents a viscous population multi-agent system, which is claimed to provide scope for the emergence of cooperation both through iterated interaction and through kin selection. Theoretical examinations of iterated interaction and kin selection within the model are conducted and compared with empirical results. It is concluded that the model does allow for the operation both of iterated interaction and kin selection. The methods presented in the paper allow the operation of the two mechanisms to be distinguished in any instance of the model.

Introduction

The Prisoner's Dilemma (PD), first formalised by Tucker, is a well known metaphor for social interactions between individuals in which there is a dilemma over whether to act cooperatively or selfishly. In the single shot PD selfish behaviour is the only rational outcome (Nash 1950,1951). However (Axelrod 1984) proposed iterated interaction in the PD as a mechanism for promoting cooperation.

An alternative explanation for the evolution of cooperative behaviour in populations exists in (Hamilton 1964)'s theory of kin selection. Kin selection theory expects cooperative behaviour to be favoured in situations in which an individual's behaviour has an effect on the reproductive success of its relatives. In particular, so-called "viscous populations", in which individuals have limited dispersal and so tend to interact with relatives, provide such a situation. However, in these same populations, there is a potential force opposing the operation of kin selection, namely competition among relatives for finite resources.

Kin selection has previously been examined through computer simulation by (Oliphant 1994). Oliphant's model used genetic algorithms in a one-dimensional spatial environment to model individuals playing the non-iterated PD with their neighbours, and producing offspring within the same neighbourhoods. The results from the model showed that cooperation could emerge in such an environment through kin selection.

This paper presents a multi-agent evolutionary model with the PD as the model of social interaction. The model allows for the evolution of cooperation through kin selection as in (Oliphant 1994) as well as through iterated interaction. A mathematical investigation of the influence of repeated interaction and kin selection in the model is presented, and conclusions are drawn on empirical results in the light of these investigations.

The Model

The model described in this paper was implemented using the Swarm software from the Swarm Development Group (http://www.swarm.org). Source code for the model is available from one of the authors' websites (http://www.mk.dmu.ac.uk/~jmarshall).

Model Overview

The model used comprises a population of agents, and an environment in which they are situated. The environment is simply a grid of cells, with each edge of the grid wrapped around to meet its opposite edge, thus forming a torus. This type of toroidal environment has become widespread in artificial life models, such as that presented in (Nowak and May 1993). Each cell is capable of housing any number of agents from zero upwards. Cells are used as the local area of interaction, i.e. agents can only play the PD with, and mate with, agents in the cell they currently inhabit. At each time step, agents are able to move to any of the eight adjacent cells with a certain probability. In addition to this physical environment, there are eight environmental variables, namely the mutation rate, the parameterised crossover rate, the maximum population the environment can support, the initial population, the death probability, the movement probability, the initial agent energy, and the energy cost for living. These variables are described in the following sections.

Agent Description

Each agent is defined as having a chromosome, an energy level, and a memory of PD interactions with other agents. For every other "opponent" that an agent has interacted with during its lifetime the agent remembers both the last actions of itself and its "opponent" (cooperate (C) or defect (D) in each case). This memory is used to determine the action an agent will take next time it meets the same "opponent". The mapping of this interaction history to an action is achieved by the agent's strategy chromosome.

The agents' chromosomes specify characteristics of the agents, and are used during interaction and mating. These chromosomes are based on (Holland 1975)'s pioneering work using genetic algorithms in adaptive artificial systems.

Following (Mar and St. Denis 1994), a five loci chromosome is used, each locus having two alleles, a cooperation allele and a defection allele. This chromosome describes a "two-dimensional" strategy, which specifies a cooperate or defect action based on the previous action of each interacting individual. Four loci are used to specify actions based on this previous pair of interactions. The remaining locus is used to specify an initial action when the agent has never interacted with its opponent before.

Lastly, the agents also have an energy level, initialised at their "birth" and decreased by a certain amount every time step. Agents "die" when their energy level reaches zero, and the only way to replenish this energy and thus survive longer is by receiving payoffs from PD interactions with other agents. The PD payoffs used were temptation $T = 5$, reward $R = 3$, punishment $P = 1$ and sucker's payoff $S = 0$.

Model Operation

The operation of the model is now described. In the first stage an initial population is created, of a size specified by an environmental variable, and randomly distributed over the environment. Each agent's initial energy level is set to a specified level. In the second stage, each agent's energy level is decreased by the "living cost" specified. Any agent whose energy level reaches zero is removed from the population. The next stage terminates agents randomly according to the probability of death in the environment. These two methods of termination correspond to "death by starvation" and by "natural causes". In the fourth stage all the agents move to an adjacent cell with a certain probability. This probability can be changed to vary the population viscosity, the implications of which are discussed in the next section. Next, agents are randomly paired up within each cell, and each pair plays one round of the Prisoner's Dilemma. The action each agent chooses will be determined by their interaction history together, and their individual strategy chromosomes. As a result of this, agents' energy levels are increased in the next stage by the payoff received from the PD interaction. Finally the agents in a cell are again randomly paired, and produce offspring with a probability $1 - P$, where P is the proportion of the maximum possible population that the current population represents, and $0 <= P <= 1$. On completion of this stage the model re-enters the second stage and executes as before.

A pair of agents reproduce to combine their genetic material into one offspring agent, whose energy level is initialised appropriately. The genetic operators used in the reproduction process are (Spears and DeJong 1991)'s parameterised crossover operator, and a standard mutation operator.

Kin Selection

A Necessary Condition for Kin Selection

There is one fundamental condition for the operation of kin selection, that the behaviour of an individual has an effect on the reproductive success of its relatives (Hamilton 1964). Our model addresses this condition in two ways. First, the local random movement of agents leads to PD interactions among relatives, as a result of the agents' mating and interaction neighbourhoods being identical. Second, agents' behaviour has an effect on relatives' reproductive success, because payoffs from PD interactions are used as "energy" which is necessary for an agent's continued survival. Therefore exploitation of a relative (or indeed any other agent) will directly harm them by depriving them of vital "energy". This can directly contribute to their "death". The method of local interaction and mating described above creates what is termed a viscous population, which provides a suitable environment for kin selection to operate. As (Hamilton 1964) said, "we would expect to find giving traits commonest and most highly developed in the species with the most viscous populations whereas uninhibited competition should characterise species with the most freely mixing populations."

The Opposing Forces of Kin Selection

When considering kin selection as a potential explanation for cooperative behaviour, it is important to note the possible existence of mechanisms counter-acting its operation. Such a counteracting mechanism could exist in the form of competition between relatives for finite resources. (Wilson, Pollock and Dugatkin 1992) studied a viscous population model and concluded that the effects of these opposing mechanisms exactly cancelled each other, reducing evolutionary fitness to the simplest form of individual fitness. (Taylor 1992a) concluded the same for homogenous environments, as well as extending their results to patch-structured populations (Taylor 1992b). However others, such as (Kelly 1992,1994) and (Queller 1994), argued against these claims. (Kelly 1994) in particular suggested that regulation of a population at the global rather than local level would still allow kin selection to be effective. The model studied in this paper implements such global population regulation.

Examination of the Implicit Shadow of the Future

Based on the model described above, it is possible to make some simple calculations about the probability of repeated interaction between two agents. Due to the undecidability inherent in such models (Grim 1994), even those without stochastic elements, these are limited to calculations from one time step to the next. However, we can still use these calculations to get an impression of the implicit “shadow of the future” (Axelrod 1984). The shadow of the future is defined as the probability that two agents will interact again at some point in the future. If the shadow of the future is large then iterated interaction is likely, and so cooperation is favoured. If, however, the shadow of the future is small, then single iteration encounters are common, and so defection is favoured.

The following equation calculates the probability of two consecutive interactions between the same agents, where *m* represents the agents’ movement probability, *d* the environmental death probability, and *p* the size of the local population in an agent’s cell. It should be noted that it is difficult to take into account the probability of an agent’s death through running out of energy.

$$\frac{(1-d)^2}{(p-1)}\left((1-m)^2+\frac{m^2}{8}\right)$$

Repeated interaction in the model is a geometric process, ignoring death by “starvation”. Given one interaction, the probability there will be another is $1/(1-q)$, where q is the repeated interaction probability from the previous formula.

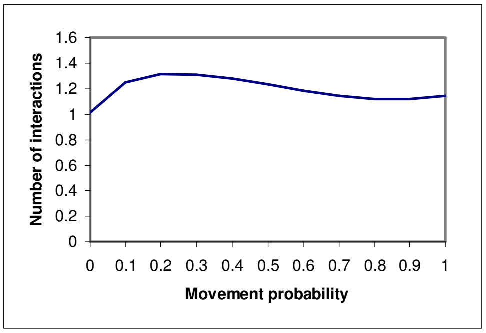

Figure 1. Expected number of interactions between agent pairs, under varying movement probability

The results of this calculation are presented in figure 1 above. The parameters used to draw this graph are: d = 0.05. As local population level p varies according to movement probability m in the model, the value for p at each value of m was taken from empirical data gathered from the model.

Examination of Inclusive Fitness of Cooperation and Defection Alleles

Not all the loci on an agent’s chromosome are subject to kin selection. Some loci are used only in interactions where both agents will see the same interaction history and therefore use the same locus to determine their actions. Hence kin selection cannot operate on these loci. However interactions using the remaining loci will always result in each interacting agent using a different locus to determine their action. At these loci alleles can benefit copies of themselves in an “opponent” agent through increasing the fitness of the “opponent” agent, possibly at the expense of agent in which they are carried. For these loci it is interesting to calculate the inclusive fitness of cooperation and defection alleles, to gain a greater understanding of the potential for the operation of kin selection within the model. This calculation is presented below. In the following equations, T, R, S and P represent the payoffs in the Prisoner’s Dilemma, c represents the frequency of cooperators in the environment, r represents the relatedness of the agents at the locus of the allele whose fitness we are considering.

Inclusive fitness of cooperation allele:

$$(Rc + S(1 - c)) + ((Rc + T(1 - c))r)$$

Inclusive fitness of defection allele:

$$(Tc + P(1 - c)) + ((Sc + P(1 - c))r)$$

In each of these equations, the first term represents the fitness of the allele in one interacting agent, the second term represents the fitness of a copy of that allele in its “opponent” agent, multiplied by a coefficient of relatedness. This is the standard inclusive fitness calculation (Hamilton 1964), interpreting payoff from a Prisoner’s Dilemma interaction as reproductive potential. Each fitness term is calculated as the average payoff received by the agent weighted by the frequency of cooperators in the environment (i.e. the probability that an agent’s “opponent” will cooperate in an interaction).

These two equations were used to draw the graph shown below in figure 2, given the following parameters: $T = 5$, $R = 3$, $S = 0$, $P = 1$.

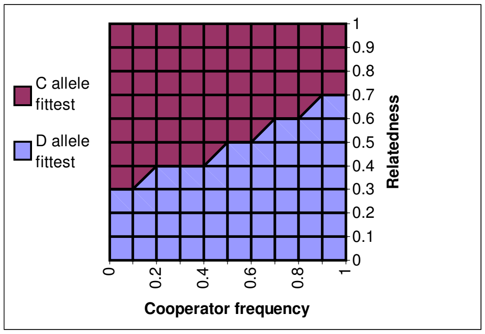

Figure 2. Relative fitness of cooperation and defection alleles at loci subject to kin selection

The inclusive fitness results in figure 2 above clearly show that the model supports kin selection. This can be seen by observing that, for the loci we are considering, the C allele is fittest at high relatedness levels.

Presentation and Discussion of Results

Experiments were run to observe the effect of varying population viscosity in the model, through changing the movement probability of agents, as follows: the experiments were conducted in 11 parameter configurations, sweeping the movement probability from

0 to 1 inclusive in increments of 0.1. For each parameter configuration 20 runs of 1000 time steps each were conducted. For each run, the average cooperation over the entire run was calculated, and this was used to calculate the average and standard deviation of the cooperation level for each value of the movement probability. The average number of interactions between any two interacting agents was also calculated, as well as the average relatedness of interacting agent pairs. Relatedness was calculated using the following formula from (Collins and Jefferson 1991), where f_i = frequency of allele 1 (for example) at locus *i*, *l* = number of loci being compared and $0 \le D \le 1$. The relatedness of two interacting agents was therefore calculated as 1-*D*, with *l* = 5.

$$D = \frac{\sum_{i=1}^{l}(1-4(0.5-f_i)^2)}{l}$$

The other model parameters were as follows: environment width and height = 10, maximum agent population = 200, initial agent population = 75, mutation rate = 0.01, parameterised crossover rate = 0.1, death probability = 0.05, initial agent energy = 15, living cost = 3.

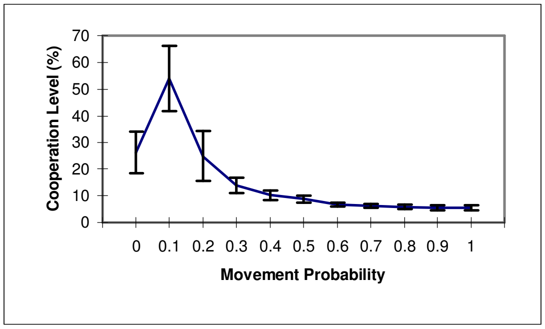

Figure 3. Global cooperation level among all agents in model

In figure 3 the cooperation level, averaged globally over all types of interaction, behaves as predicted both by iterated interaction and kin selection theory, increasing as population viscosity increases (agent movement probability decreases). It is interesting to note, however, that cooperation levels remain low across all movement probabilities, only approaching the mean at the lowest movement probabilities.

Figure 4 below presents the average number of interactions within agent pairs in the experiment, while figure 5 presents the average relatedness of pairs of interacting agents.

In figure 4, average number of interactions could be considered to be well correlated with average cooperation level in the model. Both values exhibit the same increase under decreasing movement probability, peaking at movement probability 0.1, subsequently falling at movement probability 0. In fact the only obvious difference between the curves of average cooperation level and average interaction length is the magnitude of fall at movement probability 0.

Comparing the theoretical predictions shown in figure 1 with the empirical results presented in figure 4, it can be seen that there is a close match, apart from some discrepancies at movement probabilities 0.1 and 0.2. It is suggested that the reason for these discrepancies is that the theoretical calculations cannot take account of death by ‘starvation”, which may become increasingly important at the lower movement probabilities.

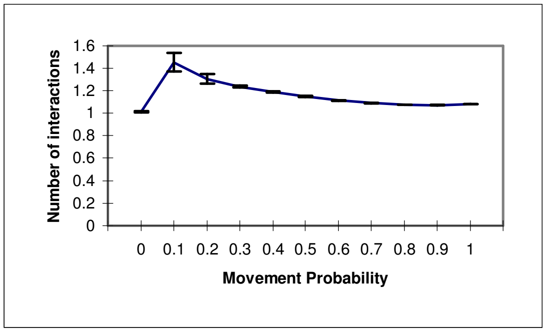

Figure 4. Average number of interactions within agent pairs

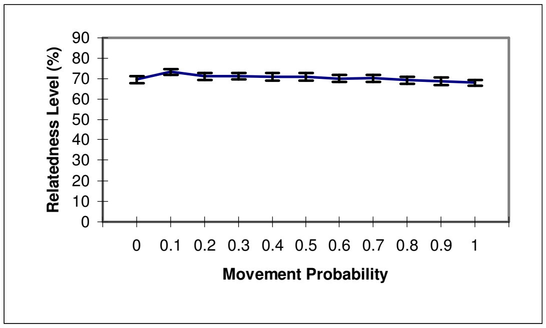

Figure 5. Average relatedness of interacting agent pairs

Figure 5 shows a high degree of relatedness within interacting agent pairs. However this relatedness level is comparatively constant over movement probability.

In addition to these results, further cooperation and relatedness data was collected from the experiments, focussing on a subset of the initial data. Figure 6 and figure 7 below present the average cooperation level and relatedness of agent pairs respectively, for interactions determined by the two loci subject to kin selection. The motivation for studying this subset of the data was presented in the section entitled ‘Examination of Inclusive Fitness of Cooperation and Defection Alleles”.

Figure 6 shows a noisy cooperation level that does not vary with movement probability, but remains around the mean cooperation level found in the absence of selection.

Figure 7 shows the same high relatedness level observed in figure 5. It is interesting to note that while relatedness does increase as population viscosity increases (movement probability decreases), this increase is not as pronounced as (Hamilton 1964) suggested it might be. Possible explanations for this may include the existence of forces counteracting the operation of kin selection, as outlined under the heading ‘The Opposing Forces of Kin Selection”. Alternatively, the particular parameters used for the experiments, more particularly the PD parameters, may not allow kin selection to operate effectively. This is suggested by the cooperation levels presented in figure 6, which are based around the mean cooperation level found in the absence of selection, and do not show any trend over varying movement probability.

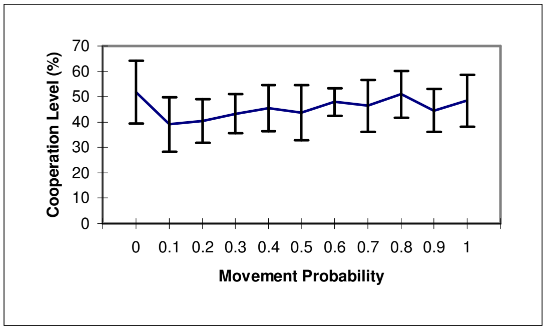

Figure 6. Average cooperation level for interactions using loci subject to kin selection

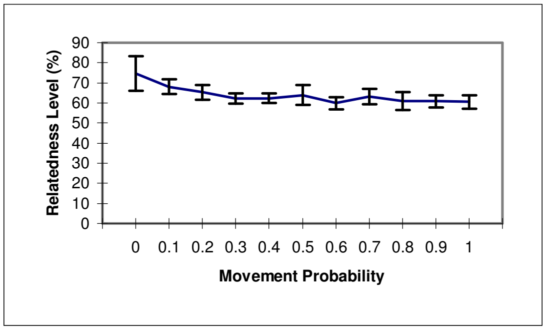

Figure 7. Average relatedness of interacting agent pairs for interactions using loci subject to kin selection

Returning to the theoretical result presented in figure 2, and taking the empirical relatedness level from figure 7, we can suggest an explanation for the observed mean cooperation level by examining the parameter space of relative C and D allele fitness. Looking at figure 2 with a relatedness level of between 0.6 and 0.7 (taken from figure 7), we find at cooperator frequencies 0.7 to 1 the threshold of relative C allele/D allele fitness. There is little difference between the fitness of the C and D alleles at this point in the parameter space, thus there is an absence of strong selection on these loci, resulting in the mean cooperation level observed.

Conclusions

This paper has presented a theoretical examination of the possible mechanisms underlying emergent cooperation in a viscous population multi-agent system. These examinations, in conjunction with empirical data from the model, conclude that the model presented allows for cooperation to emerge both through iterated interaction and through kin selection. The methods presented allow for the differentiation of the two mechanisms underlying cooperation in different cases of the presented model.

Acknowledgements

This research was funded by Sony Computer Entertainment Europe and De Montfort University.

References

Axelrod, R. 1984. *The Evolution of Cooperation*. Penguin Books, London.

Collins, R. J. and Jefferson, D. R. 1991. Selection in Massively Parallel Genetic Algorithms. In Proceedings of the Fourth International Conference on Genetic Algorithms , 249-256.

Grim, P. 1997. The Undecidability of the Spatialized Prisoner's Dilemma. *Theory and Decision* 42, 1:53-80.

Hamilton, W. D. 1964. The Genetical Evolution of Social Behaviour (I and II). *Journal of Theoretical Biology* 78:1-16,17-52.

Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press.

Kelly, J. K. 1992. Restricted Migration and the Evolution of Altruism. *Evolution* 46, 5:1492-1495.

Kelly, J. K. 1994. The Effect of Scale Dependent Processes on Kin Selection: Mating and Density Regulation. *Theoretical Population Biology* 46:32-57

Mar, G. and St. Denis, P. 1994. Chaos in Cooperation: Continuous-Valued Prisoner's Dilemmas in Infinite-Valued Logic. *International Journal of Bifurcation and Chaos* 4,4 :943-958.

Nash, J. F. 1950. Equilibrium Points in N-Person Games. *Proceedings of the National Academy of Sciences of the United States of America* 36:48-49.

Nash, J. F. 1951. Non-Cooperative Games. *Annals of Mathematics* 54:286-295.

Nowak, M.A. and May, R.M. 1993. The Spatial Dilemmas of Evolution. *International Journal of Bifurcation and Chaos* 3, 1:35-78.

Oliphant, M. 1994. Evolving Cooperation in the Non-Iterated Prisoner's Dilemma: The Importance of Spatial Organisation. In Proceedings of the Fourth Artificial Life Workshop, 349-352.

Queller, D.C. 1994. Genetic Relatedness in Viscous Populations. *Evolutionary Ecology* 8:70-73.

Spears, W. M., DeJong, K. A. 1991. On the Virtues of Parameterized Uniform Crossover. In Proceedings of the Fourth International Conference on Genetic Algorithms, 230-236.

Taylor, P. D. 1992a. Inclusive Fitness in a Homogenous Environment. *Proceedings of the Royal Society of London B* 249:299-302.

Taylor, P. D. 1992b. Altruism in Viscous Populations – An Inclusive Fitness Model. *Evolutionary Ecology* 6:352-356.

Wilson, D. S., Pollock, G. B. and Dugatkin, L. A. 1992. Can Altruism Evolve in Purely Viscous Populations? *Evolutionary Ecology* 6:331-341.