Achieving cooperation using Artificial Immune Systems

Tzai Der ang [1] and Colin Fyfe [2]

[1] Department of Industrial Engineering and Management
Cheng Shiu University, Taiwan, ROC
E-mail: douglaswang@giga.net.tw
[2] Applied Computational Intelligence Research Unit
The University of Paisley, Scotland, UK
E-mail: colin.fyfe@paisley.ac.uk

Abstract

In this paper, we discuss an artificial immune system algorithm which is capable of achieving cooperation in the context of the N-persons Iterated Prisoner's Dilemma. We use a representation for the antibodies which is identical to that used in our previous investigation of the same problem with evolutionary algorithms. We show that, while there is an interaction between the parameters, one of the most important factors is the number of clones of successful antibodies which are passed into the next generation.

Keywords: Artificial Immune Systems, N-player iterated Prisoner's Dilemma

1. Introduction

Game theory has grown significantly during the last 50 years so that it now has several distinct strands. One of the most actively researched is evolutionary game theory which has involved a mutually beneficial interaction between game theorists and evolutionary biologists: on the one hand, we have biologists using game theory to explain and interpret evolution [8]; on the other hand, we have game theorists investigating how evolutionary processes (which are, after all, the "blind watchmaker" [6]) can find optimal solutions to games. Perhaps the most famous outcomes of the second thrust is the seminal work of Axelrod [1] and Rapoport [9] who showed how cooperation could be evolved in the context of the Iterated Prisoner's Dilemma.

In this paper, we investigate the emergence of cooperation using Artificial Immune Systems.

2. The Prisoner's Dilemma

The Prisoner's Dilemma is a classic problem from Game Theory: it is a 2X2 non-zero sum non-cooperative game.

Axelrod [1] investigates the iterated prisoner's dilemma (IPD) to describe the appearance of cooperation. By iterating a potentially infinite number of games between the players, the "shadow of the future" within the current decisions is established. Because the players know that they might meet again in the future and also do not know when the last interaction between them will take place, this makes it possible for cooperation to emerge. This possibility means that the choices made today not only determine the outcome of this move, but also influence the later choices of the players.

Axelrod [1] used evolutionary algorithms to show how cooperation could be evolved using evolutionary algorithms.

The N-player Prisoner s Dilemma game (N-IPD), suggested by Colman [4], is qualitatively different from the two person game and that certain strategies that work well for individuals in the 2-Prisoner's Dilemma fail in large groups. Let C_i (D_i) refer to the payoff to the current strategy if it cooperates (defects) when there are i *t er* cooperators in the population and the payoff matrix of the N-IPD must satisfy,

1. It pays to defect: D_i C_i for all i in 0,…,N-1.
2. Payoffs increase when the number of cooperators in the population increases: $D_{i\ 1}$ D_i and $C_{i\ 1}$ C_i for all i in 0,…,N-1
3. The population as a whole gains more, the more cooperators there are in the population: $C_{i\ 1}$ (C_i $D_{i\ 1}$)/2 for all i in 0,…,N-2. Notice that this last gives a transitive relationship so that the global maximum is a population of cooperators.

Using a similar experimental design as Axelrod used for the 2-player IPD, Yao and Darwen [11] investigated the N-IPD via Genetic Algorithms. They proved that cooperation can still be evolved but it is

more difficult to evolve cooperation as the number of players increases. We have previously investigated a variety of evolutionary algorithms and problem representations [10] and shown that the most common outcome is a local fitness optimum in which players mutually defect at the first few rounds then mutually cooperate after that. These situations appear no matter which selection operators are used so that we consider the problem to be innate in genetic procedures. In this paper, we investigate whether algorithms derived from our immune systems can overcome this problem.

3. Artificial Immune Systems

The immune system is a complex system that enables a mechanism by which certain dangers to the organism can be identified. These dangers can be roughly classified as those which arise from dysfunction within the organism's own cells and those which are due to the action of exogenous pathogens. Most attention in the world of artificial immune systems has concentrated on the latter since this is the facet which gives rise to the pattern matching capabilities of AIS.

The root of the system we will be emulating is a set of cells known as the B-lymphocytes which are created in the bone marrow. They are responsible for producing specific antibodies in response to the identification of specific antigens: each antibody will respond optimally to a specific antigen rather like a key which fits into a keyhole. The key part of the antigen is known as the epitope and the keyhole part of the antibody is known as the paratope.

The immune system itself is partly innate – we are born with an active immune system –and partly adaptive – the body learns to respond appropriately to the attacks which the organism meets in life. There are limited resources available for the production of antibodies so that in responding to a particular antigen by creating specific antibodies, the system inevitably has fewer resources available to respond to other antigens.

The adaptive immune system is believed to be continually creating antibodies in a somewhat random fashion: it is rather as though it is exploring the space of antigens always on the lookout for new dangers to the organism. This can be used in creating probes on the environment and in readying the system for optimal response to new threats or opportunities; when an antibody is able to respond to an antigen, it is cloned so that more antibodies of a similar type are available to meet similar antigens. The closer the antibody is to matching the antigen, the more it is cloned. This cloning comes, though, with a high rate of mutation which allows the body to refine its response to any attack by antigens. The rate of mutation, too, is dependent on the closeness of the match which the antibody makes with the antigen: a close match will lead to a lower rate of mutation. One further feature of the immune system which is of interest is its memory: the adaptive immune system remembers which attacks the organism had to withstand and retains a response mechanism which is optimal for attacks of the same or a similar type in future.

A final feature of the immune system in which we will be interested is based on Jerne's immune network theory [7]. This proposes that the various parts of the immune system itself recognize other parts of the same system and indeed affect the production or suppression of other parts of the system. A positive reaction between components can lead to cell proliferation and activation while a negative response can lead to cell suppression. For a fuller discussion of AIS, see e.g. [2],[3],[5].

3.1. Representation

The first bit determines if the strategy will cooperate ("1") or defect ("0") on the first round. In round 2, the decision is taken dependent on whether the strategy cooperated or defected in round 1; this determines whether to use the second or third block of the chromosome. Then in each block, the bit that determines whether to cooperate or to defect in this round is identified dependent on the number of cooperators in round 1. This may vary from 0 to N, so N bits are necessary. In round 3 the decision depends on the strategy s previous decisions: DD, CD, DC, or CC. Within each block there are N N histories determined by the number of cooperators in each of the first two rounds, i.e., 00, 10, 20, …,N0, 01, 11, …,N1, and so on where AB means there were A cooperators in the first round and B in the second. The remainder of the antibody is determined in the same manner.

We chose this representation so that we could compare our results from the AIS with those achieved using genetic algorithms (GAs). Our main findings from the GA simulations were [10]:

1. GAs are generally good for developing cooperation in the N-IPD but often find local fitness optima.
2. GA with rank selection is better than roulette-wheel selection in terms of finding cooperation.
3. But rank selection suffers from premature convergence which leads to many local optima (approximately 1/3 of the optima were only local optima).

4. Simulations

Our GA investigation showed that the 7-players IPD is something of a watershed. With less than 7 players, cooperation is relatively easy to find but above 7, it is very difficult to evolve. A preliminary investigation with the AIS tended to confirm this and so we report on the 7-IPD. Each antibody plays a randomly chosen selection of 6 other antibodies; this competition is carried out for 20 rounds before the fitness of each antibody is calculated. At each iteration, we find the antibodies in the current population which have greatest fitness when they play a random sample of 6 other antibodies in the population over 20 rounds. This competition is performed 10 times for each antibody so that the fitness should be close to the fitness achieved by the antibody playing against any other strategy. We clone these best antibodies so that the best is cloned p times, the next best is cloned p times etc. Each of these clones is subject to mutation and the mutation rate is another parameter which we may vary. The mutation rate is greatest for the least fit clones. For the j^{th} individual of the its mutation rate = $(1 (f_1-f_j))*\mu$ where f_1 and f_j are the fitnesses of the most fit antibody and the j^{th} antibody, respectively.

Clearly, cloning is the method whereby exploitation is carried out and the accompanying mutation introduces exploration. Note, however, that as the population converges, the difference between the fittest antibody and the th (the last to be cloned) decreases and so the exploration decreases and the exploitation increases. Initially, the exploration rate around the lower fitness clones is rather high; therefore we have experimented with different numbers of winners, , different number of clones, p, and different mutation rates, μ. These parameters interact e.g. adopting more winners (increasing) leads to a greater difference between f_1 and f_m and so gives greater mutation around the less fit clones.
This still leaves a substantial parameter space to investigate. In order to keep this investigation manageable, we have used a linear relationship between the number of clones of the jth most fit antibody and the number of clones we make of it: w_j j . For most of the reported experiments, . For the initial ones, *i e* we make clones of the best antibody, ( ) of the second best antibody and so on down to of the th best antibody.

Each statistic we report is based on 25 simulations using that particular parameter set and with the particular methods being discussed. To give all our results would require a whole book and so we give a general flavour of our findings and try to illustrate specific points.

4.1. The Mutation Rate μ

We did not find a parameter setting which guaranteed total mutual cooperation on the 7-IPD. Figure 1 shows a simulation in which total mutual cooperation was found. This was with μ p and

Figure 2 shows a second simulation with the same parameter set in which a local optimum is found. After one defection, mutual cooperation is created i.e. we have a local optimum. However only 2 (out of 25) simulations at this parameter set achieved total cooperation (such as Figure 1) and only another 3 achieved partial cooperation such as shown in Figure 2. The remainder did not converge. However it should be noted that, if convergence to cooperation is achieved, it is very stable. On the other hand, if we increase the mutation rate to 0.03, we find 6 simulations with convergence to cooperation but also simulations such as that shown in Figure 3: the global optimum is achieved but is subsequently broken because of the high rate of mutation.

4.2. The umber of inners *m*

If we choose the above regime, we are only putting 15 clones (5 4 3 2 1) into the population of 100 at each stage; however, this leaves at least 85 antibodies which are simply passing into the next population without ever having had to undergo a selection operation. We can treat this as noise which has no beneficial effect on the emergence of cooperation.

We consider two methods of varying the number of winners. We first consider the effect of increasing . It could be argued that, in the experiments described in the previous section, we do not have enough winners: since we are investigating the 7-IPD, each antibody's fitness is determined by its interaction with 6 other antibodies in the current game. Therefore we should have ≥ 7. Thus we select 10 but keep μ=0.001. With this parameter set, we found that, in 25 simulations, 16 showed convergence to cooperation with 3 of these finding the global fitness optimum – much more cooperation is in evidence than when 5. We find a gradual monotonic decline in the level of cooperation when we decrease from 10 to 9 to 8 to 7 to 6 to 5, with somewhat of a knee at the transition between 7 and 6. However, it could be argued that we are merely replacing a greater proportion of the population than before since we now are putting 55 clones into the population compared with the 15 when =5. If we repeat the simulation, with 9 but =2, (i.e. we clone the best 10 times, the next best 9 times and the 9th twice) there are still 54 clones put into the population but we only get 9 trials in which there is a convergence to cooperation. Further, in these trials, cooperation tends to emerge late and is sometimes unstable. Thus, we find that the number we choose to

clone is important in the emergence of cooperation as well as the actual number of clones themselves. It is also worth noting that with the 10 winners discussed above the globally optimal convergence is more stable than the locally optima.

However, this comparison does not give the whole story; again there is an interaction with the other parameters. Figure 4 shows the number of times cooperation was achieved (out of 25 simulations at each parameter setting) with some example sample sizes and mutation rates.

5. Discussion

One of the advantages that artificial immune systems are said to have over evolutionary algorithms is that there are fewer parameters to tune. We have seen that, while this may be so, there are certainly enough parameters to make a search of the parameter space a time-consuming and complex process. In any case, it may be that AIS is simply a newer technology than GAs and, as they are investigated, more "bells and whistles" will be added leading to an increase in the number of parameters.

The fact that a new population is not created at each generation as it is with most GAs means that there is a great deal of information passed onto the next round of the game which comes from antibodies which were not selected because of their fitness. We therefore found that increasing the number of winning antibodies (antibodies to be cloned) at each time instant had a beneficial effect on the emergence of cooperation. However, this should be treated with some caution: in some circumstances the more winners there are, the greater the mutation rate for the poorer winners and the more the search resembles a random search.

Overall our results are in line with those achieved using GAs and other evolutionary algorithms but there is one area which is specific to the AIS: only in the AIS is the globally optimal solution of total cooperation broken; with all the other algorithms, if the global optimum is found, it is retained. This aspect and the means to either prevent or circumvent it, is the subject of future research.

6. References

[1] Axelrod, R. "The Evolution of Co-operation", Penguin books, 1984.

[2] De Castro, L. N. and von Zuben, F. J., "Artificial Immune Systems: Part 1 – Basic Theory and Applications", Technical report TR DCA 01/99, from ftp://ftp.dca.fee.unicamp.br/pub/docs/vonzuben/Inunes. 1999.

[3] De Castro, L. N. and von Zuben, F. J., "Artificial Immune Systems: Part 2 – A Survey of Applications", Technical report TR DCA 02/00, (same site) , 2000.

[4] Colman, A. M., "Game Theory and Experimental Games", Pergammon Press, Oxford, England, 1982.

[5] Dasgupta, D., "Artificial Immune Systems and their Applications", Springer-Verlag, 2000.

[6] Dawkins, R." The Blind Watchmaker", Harlow: Longman, 1986.

[7] Jerne, N. K., "Towards a network theory of the immune system". *Ann l n l* y, 125C, pp. 373-389, 1974.

[8] Maynard Smith, J. "Evolution and the Theory of Games", Cambridge University Press, 1982.

[9] Rapoport, A., "N-Person Game Theory, Concepts and Applications", Dove edition. 2001.

[10] Wang, T.-D., "The Evolution of Cooperation in Artificial Communities", PhD thesis, the University of Paisley, Scotland, 2002.

[11] Yao X., & Darwen, P. J "An experimental study of N-person iterated prisoners dilemma game", *n r ti* , pp. 435-450, 1994.

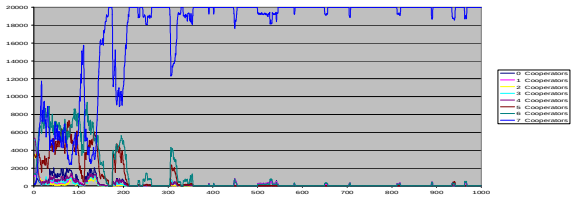

Fig. 1: The global optimum: total mutual cooperation.

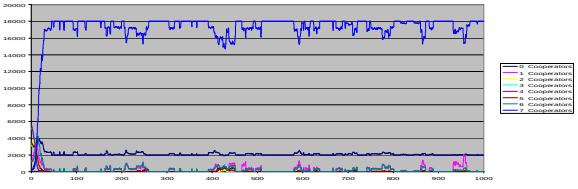

Fig. 2: A local optimum: mostly coop. but some def..

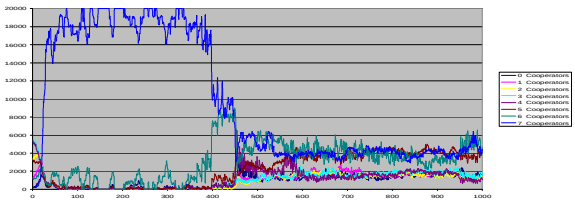

Fig. 3: Global cooperation is achieved but lost.

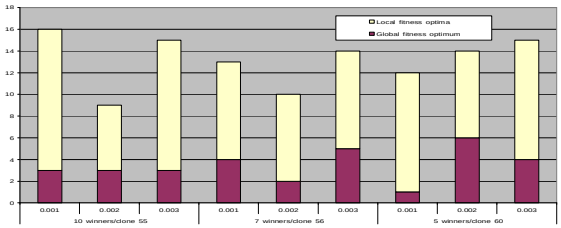

Fig. 4: The relationship between , μ and the emergence of cooperation is rather complex.