Using Tags to Evolve Trust and Cooperation Between Groups

Aviv Zohar Jeffrey S. Rosenschein
School of Engineering and Computer Science
Hebrew University, Jerusalem, Israel
{avivz, jeff}@cs.huji.ac.il

ABSTRACT

Tags [2] are externally visible markers that are determined by an individual's genes. We study a simple model of interactions between large groups of boundedly rational players playing the prisoner's dilemma [1], who are allowed to see one another's tags prior to choosing an action, but may not choose their opponent. Since tags are genetically determined, they are correlated with the behavior of the tagged individual, and possessing a similar tag implies similar behavior. Our tag model exploits this correlation so as to enable beneficial interactions between groups of players. Computer simulations show that with the tag mechanism in place, cooperation between different groups of players can become common.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed AI—*Multiagent Systems*

General Terms

Experimentation

Keywords

Multi-agent simulation & modeling, evolution, artificial social systems, tags, trust and reputation, prisoner's dilemma

1. THE MODEL

Our model contains a population of players that is allowed to interact by playing one-shot prisoner's dilemma games, with actions **C**ooperate and **D**efect. Each player also possesses a tag from a discrete tag space T and memory cells, one for each tag it encountered in the past. We shall denote the memory cell of player i regarding tag t as mem_i^t. Each memory cell may be in one of the states S. The behavior of the player is then determined by two functions: $MemUpdate : S \times A \rightarrow S$; $Action : S \rightarrow A$.
The tag and the functions $MemUpdate$, $Action$, as well as the initial state of memory cells, are all encoded in the player's genes.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.

When player i encounters player j that has tag $t \in T$, player i acts as follows: (1) observe j's tag t; (2) play $Action(mem_i^t)$; (3) observe j's action $a \in A$ and receive a reward; (4) set $mem_i^t \leftarrow MemUpdate(mem_i^t, a)$. The behavior of players with regard to a fixed tag can be modeled by an automaton. At every generation, players get to play several rounds of games, after which their fitness is determined by the average reward they obtained. Next, a reproductive stage occurs where 5% of the population's fitter members are asexually reproduced. A small probability of mutation is introduced during reproduction.

2. SIMULATION RESULTS

After a transient phase at the beginning of the run, the system usually maintains a high degree of cooperation. Reciprocating players that exhibit forgiving variations of Tit-For-Tat are commonly found in the population. Oscillations in cooperation during the run emanate from the invasion of defectors into groups of cooperating agents. These defectors initially do better than other players, but as the number of cooperators in their group shrinks, these defectors, having killed off exploitable victims and having established their "reputation" as exploiters, earn less and are displaced by other tag groups. Simulations show that only about 4.5% of the games are played with members of the same group. Thus, the main factor in determining the fitness of players is how well they do when playing with dissimilar players.

Figure 1: Tournament rounds each generation vs. social utility

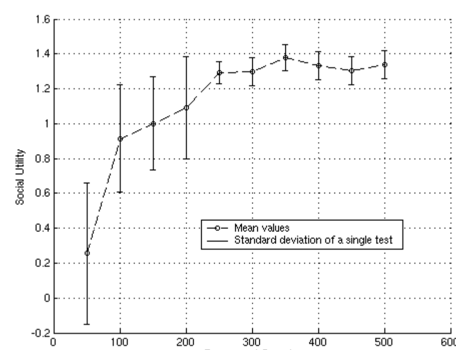

2.1 Social Utility & Number of Game Rounds

It is to be expected that if more game rounds are played every generation, individuals that make use of their memory of past interactions will do better (for example, reciprocating players). As seen in Figure 1, this was found to be true. Social utility is usually higher if more game rounds are allowed in each generation. It is remarkable to observe that the system does indeed reach very high levels of cooperation. A social utility score of 1.4 is reached (from the possible range of [-1,2]). This means that in 75% of all games, both players chose to cooperate. Adding more game rounds also makes the emergence of cooperation from an entirely defecting population more likely. This can be seen by the decrease in the standard deviation of the tests as the number of rounds is increased.

2.2 Increasing the Number of Memory States

Tests were run for populations of players with different numbers of memory states. The results appear in Figure 2. Giving players more memory states promotes cooperation in the system, but only up to a certain point. It appears that most of the advantages of added memory states are realized with 3 memory states. Further increases in memory affect the social utility very little. Automatons with more reachable states also require more tuning by evolution, and may not realize their full potential in the time before they are infiltrated by defecting individuals and become extinct.

Figure 2: Number of memory states vs. social utility

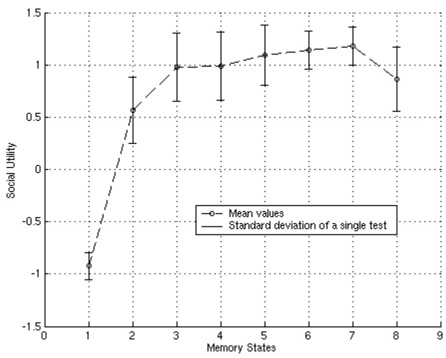

However, since players can evolve automatons with unreachable states, thereby decreasing the effective size of the automaton, there does not appear to be any obvious disadvantage to having more memory states. Note that if the players are only allowed 1 memory state, they are playing a fixed strategy. In this case, the population nearly always converged to a population of defecting players. Two memory states are already enough to encode the Tit-For-Tat strategy, but here as well, the population did not always converge to high values (the large variance indicates that there were very large differences between various tests). When given three memory states, players could encode forgiving versions of Tit-For-Tat, which have better noise tolerance, and help avoid escalating defections.

2.3 The Effects of Varying the Mutation Rate

It is important to note that even though the same mutation rate was used for all genes, automatons with unreachable states are not affected by mutations pertaining to these states. This way, the effective mutation rate of the strategy-related genes and their proportion to the tag mutation rate are modified by evolution itself. Figure 3 shows the results of runs of the simulation with various mutation rates. At a mutation rate of zero, the population converges to complete cooperation, since there is a very good chance that there is a player within the randomly initialized population that plays a variation of Tit-For-Tat and has a unique tag. The lack of mutation in tags and strategies means there are no possible invasions by exploiting individuals that may threaten such players. Adding a small mutation rate allows for some small probability of invasion into each tag group, and so the social utility declines abruptly.

Increasing the mutation rate further brings about a rise in the systems' social utility; the mutation rate of tags becomes large enough to allow groups of cooperators to expand quickly to unoccupied regions of the tag space. As the mutation rate is increased further, the system degrades slowly, possibly because of the loss of correlation between tag and behavior that higher mutation rates cause.

Figure 3: Effects of mutation rate on social utility

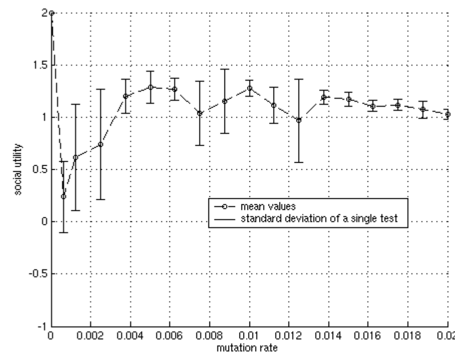

2.4 The Size of the Tag Space

Tag-based systems contain a never-ending race between cooperative players and exploiters. Exploiters behave much like viruses and epidemics — they spread from one cooperative group to another via mutation of their tags. A larger tag space decreases the chance that such a mutation will randomly find a cooperative group. Exploiters are then formed mostly from within the group — by mutation of one of the behavior encoding genes. Another benefit of a large tag space is that it allows "nice" groups a chance to form even in an entirely defecting population. Mutations are more likely to fall on an empty tag and may start to evolve and interact with some separation from the rest of the population. The results displayed in Figure 4 support these hypotheses and demonstrate that a larger tag space allowed for more cooperation on average in the system.

Figure 4: Size of tag space vs. social utility

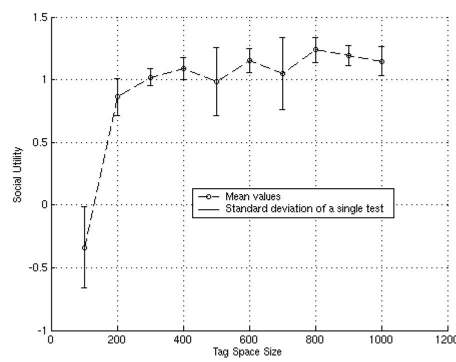

Players need to maintain a memory cell for every tag they encounter but the tag space can still be made very large with very little cost since only a small portion of it is actually used. Simulations show that for a tag space of size 4000 only around 100 tags are actually occupied.

3. REFERENCES

[1] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

[2] D. Hales and B. Edmonds. Evolving social rationality for MAS using "tags". In *AAMAS'03*, pages 497–503, Melbourne, Australia, July 2003.