

The Evolution of Agent Strategies and Sociability in a Commons Dilemma

Enda Howley
Department of Information Technology
National University of Ireland
Galway
enda.howley@nuigalway.ie

Jim Duggan
Department of Information Technology
National University of Ireland
Galway
jim.duggan@nuigalway.ie

ABSTRACT

This paper explores the evolution of strategies in a n-player dilemma game. These n-player dilemmas provide a formal representation of many real world social dilemmas. Those social dilemmas include littering, voting and sharing common resources such as sharing computer processing time. This paper explores the evolution of altruism using an n-player dilemma. Our results show the importance of sociability in these games. For the first time we will use a tag-mediated interaction model to examine the n-player dilemma and demonstrate the significance of sociability in these games.

Keywords

Evolution, Learning, Cooperation, Agent Interactions, Tragedy of the Commons, Tag-Mediated Interaction Models

1. INTRODUCTION

When a common resource is being shared among a number of individuals, each individual benefits most by using as much of the resource as possible. While this is the individually rational choice, it results in collective irrationality and a non Pareto-optimal result for all participants. These n-player dilemmas are common throughout many real world scenarios. For example, the computing community is particularly concerned with how finite resources can be used most efficiently where conflicting and potentially selfish demands on those resources are common. Those resources may range from access to processor time or bandwidth.

One example commonly used throughout existing research is the *Tragedy of the Commons* [5]. This outlines a scenario whereby villagers are allowed to graze their cows on the village green. This common resource will be over grazed and lost to everyone if the villagers allow all their cows to graze, yet if everyone limits their use of the village green, it will continue to be useful to all villagers. Another example is the *Diners Dilemma* where a group of people in a restaurant agree to equally split their bill. Each has the choice to exploit the situation and order the most expensive items on the menu. If all members of the group apply this strategy,

Cite as: The Evolution of Agent Strategies and Sociability in a Commons Dilemma, Enda Howley and Jim Duggan, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. XXX-XXX.
Copyright © 2009, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

then all participants will end up paying more [2].

These games are all classified as n-player dilemmas, as they involve multiple participants interacting as a group. These games involve only two players interacting through pairwise interactions. N-player dilemmas have been shown to result in widespread defection unless agent interactions are structured. This is most commonly achieved through using spatial constraints which limit agent interactions through specified neighbourhoods on a spatial grid. Limiting group size has been shown to benefit cooperation in these n-player dilemmas [14].

In this paper we will examine an n-player dilemma, and study the evolution of strategies when individuals can bias their interactions through a tag mediated environment. Furthermore, we will show how certain strategies evolve with respect to their sociability towards their peers. The simulations presented in this paper use the n-player Prisoner's Dilemma (NPD). The purpose of this paper is to examine the evolution of cooperation and sociability throughout the agent population in the NPD. The research presented in this paper will deal with a number of specific research questions:

1. Can a tag-mediated interaction model be used to determine group interactions in a game such as the NPD?
2. If agents have an evolvable trait which determines their sociability, then will this trait prove significant to the emergence of cooperation in the agent society?

The following section of his paper will provide an introduction to the NPD and a number of well known agent interaction models. In the Experimental Setup Section we will discuss our simulator design and our experimental parameters. Our Results Section will provide a series of game theoretic simulations. Finally we will outline our conclusions and future work.

2. BACKGROUND RESEARCH

In this section we will introduce the NPD game while also discussing some existing background research relevant to this paper.

2.1 The N-Player Prisoner's Dilemma

The n-player Prisoner's Dilemma is also known as the Tragedy of the Commons [5] and the payoff structure of this game is shown in Figure 1.

On the horizontal axis is the fraction of cooperators in the group of n players in a particular game. On the vertical axis is the payoff for an individual participating in a




Figure 1: The N-Player Prisoner's Dilemma

game. There is a linear relationship between the fraction of cooperators and the utility received by a game participant. Importantly, the payoff received for a defection is higher than for a cooperation. The utility for defection dominates the payoff for cooperation in all cases. Therefore, an individual that defects will always receive a higher payoff than if they had chosen to cooperate. The result of this payoff structure should result in an advantage to defectors in the agent population. Despite this, a cooperator in a group of cooperators will do much better than a defector in a group of defectors.

This game is considered a valid dilemma due to the fact that individual rationality favors defection despite this resulting in state which is less beneficial to all participants. In our case where all individuals defect they all receive 0.5. This state is a non-pareto, sub-optimal, and collectively irrational outcome for the agent population. For all values of x this can be expressed as follows: $U_d(x) > U_c(x)$. x is the fraction of cooperators while U_d and U_c are utility functions based on the fraction of cooperators in the group.

2.2 Agent Interaction Models

A number of alternative agent interaction models have been proposed and examined, such as spatial constraints [11, 10] and tag mediated interactions, [13]. The importance of group size has been demonstrated explicitly through tags in the PD by [7]. Similarly in the NIPD [14]. Yao and Darwin demonstrated the effects of limiting group size, which was shown to benefit cooperation. Increasingly complex aspects of agent interactions have been examined by a number of authors, these include the effects of community structure on the evolution of cooperation [12, 1]. These have shown that neighbourhood structures benefit cooperation.

In this paper we are most concerned with tag-mediated interactions. Tags are visual markings or social cues which can help bias social interactions [6]. They are a commonly used agent interaction model and can be considered akin to football supporters identifying each other through wearing their preferred team colours. Similarly individuals can identify each other in conversations through a common language, dialect, or regional accent. Tag-mediated interaction models are often considered as more abstract interaction models, and thereby useful to represent agent interactions more abstractly without the specific characteristics of a specific topology or implementation. The research presented by Riolo demonstrated how tags can lead to the emergence of cooperation in the Prisoner's Dilemma [13]. Riolo investigated both a fixed and an evolved tag bias. More recently tags have been successfully applied to multi-agent problems [3,

4]. Tags have been shown to promote mimicking and thereby have major limitations where complimentary actions are required by agents. Cooperation that can be achieved through identical actions is quite easily achieved using tags, yet behaviours that require divergent actions are problematic [9, 8].

In this paper we will augment existing research to show the effects of using a tag-mediated interaction model to determine group interactions in the NPD. The following section will provide a detailed specification of our simulator and the overall design of our experiments.

3. EXPERIMENTAL SETUP

In this section we will outline our agent structure, our agent interaction model and our evolutionary algorithm.

3.1 Agent Genome

In our model each agent is represented through an agent genome. This genome holds a number of genes which represents how that particular agent behaves.

$$Genome = G_C, G_T, G_S, \quad (1)$$

The G_C gene represents the probability of an agent cooperating in a particular move. Each agent has G_C gene which never changes throughout their lifetime. The G_T gene represents the agent tag. This is represented in the range $[0..1]$ and is used to determine which games each agent participates. Finally, the G_S gene represents the sociability of each agent. This gene is also a number in the range $[0..1]$ which acts as a degree of sociability for that individual agent. Initially these agent genes are generated using a uniform distribution for the first generation. Over subsequent generations new agent genomes are generated using our genetic algorithm.

3.1.1 Tag Mediated Interactions

In our simulations each agent interacts through a simple tag mediated interaction model. We adopt a similar tag implementation as that outlined by Riolo [13]. In our model each agent has a G_T gene which is used as their tag value. Each agent A is given the opportunity to make game offers to all other agents in the population. The intention is that this agent A will host a game and the probability other agents will participate is determined as follows.

$$d_{A,C} = 1 - |A_{GT} - C_{GT}| \quad (2)$$

This equation is based on the absolute value between the tag values of two agents A and C . This value is used to generate two roulette wheels R_a and R_c for A and C . These two roulette wheels will then be used to determine agent A 's attitude to C and agent C 's attitude to A . An agent C will only participate in the game when both roulette wheels have indicated acceptance. The distribution of these roulette wheels are also influenced by each agents sociability gene. This gene acts like a scalar value which is used to reflect that some agents are more sociable than others and will therefore be more willing to play with their peers. This is shown in the following equation, where R_a represents the roulette wheel probability of entering a game.

$$R_a = d_{A,C} \times A_{GS} \quad (3)$$

Each agent in the population makes a game offer to all other agents, and the set of agreed players then participate in the NPD game.

3.1.2 Genetic Algorithm

In our simulator we have implemented a simple genetic algorithm. In each generation individuals participate in varying numbers of games. Therefore, fitness is determined by summing all their payoffs received and getting an average payoff per game. In each generation, the top 10% of agents are copied directly into the following generation. The other 90% of the agent population in generation $G + 1$ are generated through evolving new strategies based on agent fitness in G . Individuals are selected through roulette wheel selection based on their fitness from generation G . Parent pairs are selected and then these are used to generate a single new agent offspring for generation $G + 1$. Crossover occurs through averaging the genes between the two parent strategy genomes G_C, G_T, G_S . These averaged strategy genes are then used for the new agent. A 5% chance of mutation on each of these strategy genes is also used, and once this occurs a gaussian distribution is used to determine the degree of change.

4. EXPERIMENTAL RESULTS

In this section we will present a series of simulations showing the results of our experiments. Firstly, we will examine a set of graphs depicting the results from a single run over 1000 Generations. The aim of this single run is to show the inherent links between certain agent gene values and the overall cooperation throughout the agent population. Later in this section we will present simulations showing results from a number of experimental runs. These will demonstrate the overall stability of our results over multiple runs. All our simulations were conducted using an agent population of 100 agents.




Figure 2: Average Gene Value (1 Run)

Figure 2 shows the rapid emergence of cooperation throughout the agent population. This graph depicts the average G_C and G_S genes throughout the agent population in each generation. The results show the emergence of cooperation as the average G_S gene falls throughout the population. These results show a rapid drop in the average G_S gene which reflects the tendency of the agent population to interact with fewer peers. The increased levels of cooperation throughout the population are closely linked with the tendency of individuals to act less sociably. It is clear from the results that the heightened cooperative gene is linked directly with the lower sociability gene.




Figure 3: Average Number of Games (1 Run)

The results in Figure 3, depict the average number of games each agent participates in throughout successive generations. These results show the underlying dynamics that resulted in the heightened average cooperation shown in Figure 2. Once agents begin to participate in multiple n-player dilemmas they are exposed to exploitation and they are then heavily penalised. It is clear that cooperation is achieved through agents participating in as few games as possible. This serves to limit their exposure to potential exploitative peers.

The simulations shown are from a single run over 1000 generations. These simulations show the close relationship between the various agent gene values, and the collective behaviour of the agent population. For example around Generation 440 we can identify a period of increased sociability and a corresponding drop in cooperativeness throughout the population. This feature is clearly identifiable through examining the average gene values in Figure 2 and also the average game participation results in Figure 3.

These results are confirmed when examined over multiple experimental runs. The following graphs are averaged over 25 experimental runs. The purpose of these experiments is to confirm that the overall trends identified previously are repeated over many runs.




Figure 4: Average Gene Value (25 Runs)

The data shown in Figure 4 show the average strategy genes averaged over many experiments. The results show that the agent population consistently converges on cooperation throughout multiple experiments. We also notice the low G_S genes recorded throughout the simulations. Through limiting game participation to a tiny number of games, each agent minimises the opportunity of less cooperative individuals to exploit them. Once cooperative strategies benefit heavily by limiting their interactions they receive heightened payoffs and then this feature is propagated throughout new agents in the population.




Figure 5: Average Number of Games (25 Runs)

The tendency to interact in a small number of games is confirmed in Figure 5, which depicts the average number of games each agent participates in. Our results indicate a clear benefit to individuals who are less sociable and thereby choose to be far more discerning regarding game participation. This facilitates the emergence of cooperation and helps to maintain cooperation it over successive generations.

5. CONCLUSIONS

This paper has examined the NPD game with respect to group participation. For the first time this game has been investigated using a tag-mediated interaction model. Our results demonstrate that despite there being a clear incentive to defect, cooperation can still emerge. This stems from the ability of individuals in our agent population to determine their degree of sociability towards their peers. This reinforces much of the existing literature involving the traditional Prisoner’s Dilemma [7] and also the NIPD [14]. Our models reinforces these observations through an alternative approach. In our case we have not explicitly determined the sociability of our agent population. Instead we have allowed the agent population to evolve with respect to their cooperative and sociability genes. Our results have demonstrated the significance of sociability in games such as the NPD. Furthermore, we have also demonstrated the advantage to cooperative individuals who act less sociably towards their peers. Limiting game participation provides a very effective defence against exploiters. Earlier in our introduction we posed two specific research questions.

1. Our results show that tags can successfully bias interactions in the the NPD. We believe this is the first time a tag model has been applied to the NPD. Our results show the resulting levels of cooperation that emerged.
2. The significance of the sociability gene in our simulations is clear from the obvious link between cooperation and sociability in our simulations.

This paper has presented an evolutionary model capable of modeling sociability within the agent strategy genome. We have also shown how tags can be used to determine n-player games. Finally, our results have shown through an evolutionary model that there is a clear benefit to agent strategies who are cooperative in tandem with being less sociable through limiting their exposure to exploitation.

In summary this paper has shown that tags can be successfully adapted to bias agent interactions in a n-player game such as the NPD. Furthermore, we have demonstrated how an agent population can engender and maintain cooperation

through an evolvable sociability trait. In future work we hope to examine how cooperation can be engendered without limiting game participation so dramatically.

6. ACKNOWLEDGMENTS

The authors would like to acknowledge the continued support of Science Foundation Ireland (SFI).

7. REFERENCES

- [1] R. Chiong, S. Dhakal, and L. Jankovic. Effects of neighbourhood structure on evolution of cooperation in n-player iterated prisoner’s dilemma. In *IDEAL*, pages 950–959, 2007.
- [2] N. S. Glance and B. A. Huberman. The dynamics of social dilemmas. *Scientific American*, 270(3):76–81, 1994.
- [3] D. Hales and B. Edmonds. Evolving social rationality for mas using “tags”. In *AAMAS ’03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 497–503, New York, NY, USA, 2003. ACM.
- [4] D. Hales and B. Edmonds. Applying a socially inspired technique (tags) to improve cooperation in p2p networks. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 35(3):385–395, 2005.
- [5] G. Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, December 1968.
- [6] J. Holland. The effects of labels (tags) on social interactions. *Working Paper Santa Fe Institute 93-10-064*, 1993.
- [7] E. Howley and C. O’Riordan. The emergence of cooperation among agents using simple fixed bias tagging. In *Proceedings of the 2005 Congress on Evolutionary Computation (IEEE CEC’05)*, volume 2, pages 1011–1016. IEEE Press, 2005.
- [8] M. Matlock and S. Sen. Effective tag mechanisms for evolving coordination. In *AAMAS ’07: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pages 1–8, New York, NY, USA, 2007. ACM.
- [9] A. McDonald and S. Sen. The success and failure of tag-mediated evolution of cooperation. In *LAMAS*, pages 155–164, 2005.
- [10] M. Nowak and R. May. The spatial dilemmas of evolution. *Int Journal of Bifurcation and Chaos*, 3:35–78, 1993.
- [11] M. Oliphant. Evolving cooperation in the non-iterated prisoner’s dilemma: the importance of spatial organisation. In *Proceedings of Artificial Life IV*, 1994.
- [12] C. O’Riordan and H. Sorensen. *Stable Cooperation in the N-Player Prisoner’s Dilemma: The Importance of Community Structure*, volume 4865 of *Lecture Notes in Computer Science (lncs) 4865*, pages 157–168. Springer-Verlag Berlin, 2008.
- [13] R. Riolo. The effects and evolution of tag-mediated selection of partners in populations playing the iterated prisoner’s dilemma. In *ICGA*, pages 378–385, 1997.
- [14] X. Yao and P. J. Darwen. An experimental study of n-person iterated prisoner’s dilemma games. *Informatica*, 18:435–450, 1994.