Emergence of Cooperative Societies in Structured Multi-Agent Systems

Marguerite Moran and Colm O'Riordan

National University Of Ireland, Galway.
m.moran4@nuigalway.ie
colmor@it.nuigalway.ie

Abstract. In this paper we investigate the emergence of cooperation in spatially organised games. We extend traditional spatial models and use a graph to model the environment. In the graph representation of the environment, each node represents an individual, and an edge between two individuals represents a neighbourhood relationship.
In our model, players interact in a Prisoner's Dilemma. We examine various learning mechanisms where the agent's strategies are selected and propagated. We investigate the effect of allowing agents learn from their neighbours to improve their individual performance.
We also explore the evolution of neighbourhoods by enabling them to grow or shrink depending on their relative fitness to other neighbourhoods.

1 Introduction

Many approaches have been investigated in an attempt to understand how cooperation may emerge in societies of autonomous, rational agents. The Prisoner's Dilemma has been adopted as the standard for studying cooperative behaviour [1], [2], [12], [11]; and has also been used in work focussing on spatially organised games [6], [9], [3], [13].

Classical game theory [8] doesn't include the effect of spatial structures on a population. In many populations, both real and artificial, individuals are more likely to interact with their neighbours than interact with a player chosen at random from the population. In order to model such scenarios more realistically, it is necessary to spatially organise agents in a multi-agent system, and for interactions to take place according to these spatial constraints. Furthermore, it is unlikely that neighbourhoods will be of uniform size throughout the population. We adopt a graph model to represent the neighbourhoods.

While our model encompasses many features and extensions of traditional models, our primary focus is on the effect of defining different social structures and neighbourhoods based on our graph representation of the environment.

In the experiments discussed in this paper, we investigate the emergence of cooperation or defection in a spatially structured society. We compare the spread of cooperation in a graph based model to traditional grid representations. We investigate a range of neighbourhood sizes and learning mechanisms. Furthermore, we investigate the spread of cooperation in a spatially structured society where the neighbourhood structure itself is allowed to evolve and change, i.e., fit neighbourhoods are allowed to increase in size. Much research in traditional models deals with patterns of cooperators and defectors [9] — clusters of cooperators and defectors situated on a grid with cooperator-defector interactions along the edges of the clusters. They don't assess the overall fitness of each neighbourhood. Our model allows fit neighbourhoods to propagate and conversely, if a neighbourhood has a low fitness value it is reduced in size.

2 Related Research

2.1 The Prisoner's Dilemma

The Prisoner's Dilemma is described as a non-cooperative, non-zero-sum, two person simultaneous game. In the prisoner's dilemma, two players are separated and faced with a decision. Each has two alternatives — to cooperate or defect. Neither has knowledge of the other player's choice. If they both cooperate, they receive a payoff, R. If both defect, they receive a smaller payoff, P. If one defects and the other cooperates, the defecting strategy receives the largest possible payoff, T, and the cooperator the smallest possible payoff, S.

For a dilemma to exist, the following must hold: $T > R > P > S$. The constraint $T + S \leq 2R$ must also hold. The constraint $T + S \leq 2R$ prevents a form of cooperation where two players obtain an average payoff greater than cooperation by alternating between cooperation and defection.

2.2 Spatial Prisoner's Dilemma

Nowak and May [9] model a simple, deterministic, spatial version of the Prisoner's Dilemma, with no memories among players and no strategic elaboration; showing that cooperators and defectors and co-exist indefinitely for a subset of the parameter space.

Many researchers have investigated the effect of varying payoffs [13]. Some have suggested that if agents are indistinguishable from each other, genuine cooperation cannot emerge (Frank [4], Maynard-Smith [8], Hofbauer [7], Weibull [18], Samuelson [15]). Evidence in social science shows that cooperation can emerge even if creatures cannot recognise individual players [17], [11]. Epstein[3] shows that cooperation can emerge and endure if negative payoffs are introduced. He requires that $T > R > 0 > P > S$, and his results depend largely on the values of T, R, P and S.

Cooperators in the population will benefit those close to them and flourish if the bonus for defection is not too large [5], [11]. Hence, groups of cooperators situated adjacent to each other can allow cooperation to flourish through mutual cooperation. Conversely, for defectors, they hurt their own kind and hence groups of defectors will not do well.

Nowak et al. [10] also experimented with the grid size and topology and also the neighbourhood type. They extended their analysis to explore simulations where players are randomly distributed on a two-dimensional plane. Players are said to be neighbours of each other if their distance is less than a certain radius of interaction r; so the number of neighbours each player has can vary. The simulations were deterministic and in discrete time. Results showed that populations on random grid were more static than on rectangular lattices and that cooperation can emerge and endure.

Ashlock [16] focuses on a static population and the effectiveness of a choice / refusal strategy undertaken by some agents in the population. He implements a graph based model, where each node represent an individual, and an edge connecting two vertices represent some relationship between the two individuals. Initially, the graph is fully connected and each edge is assigned a weight of zero. When two individuals play, the weight is incremented by 1. These weights are plotted in order to show population characteristics. Analysis is performed on a significant play graph. Edges are considered significant if they are greater than a given threshold value; and these edges determine the significant play graph. The significant play graph adds valuable information about the social behaviours of a system.

3 Simulator Design and Description

Our simulator allows us to examine the emergence (or not) of cooperation for a wide range of parameters. We compare different topologies, neighbourhood types and update strategies for varying radius values.

For a finite number of iterations (unknown to the player), each agent plays a PD game against each of its neighbours. Each player is either a pure cooperator or a pure defector. At the end of each game, depending on the payoffs received by each agent and the update strategy in place, the agent adopts the strategy of one of its neighbours. Following the completion of each game, and the application of the update rules, the number of cooperators is recorded.

3.1 Parameter Range Modelled

The parameter space within which agents play the PD game is very important. In the following section we discuss the parameters that are modelled and studied in our system. We discuss topology, initial configuration of the population, neighbourhood type and learning mechanisms.

Topology. A common example of a topology used in simulations is that of a $N \times M$ rectangular lattice. An $N \times M$ grid is a planar graph with $N \times M$ vertices arranged on a rectangular grid, and with edges connecting horizontally and vertically adjacent vertices.

We compare two types of topology. In our model, the population is placed randomly on either an $N \times M$ lattice or on a graph type structure. The $N \times M$

lattice is standard for traditional spatial models. This planar grid type structure is not an accurate or flexible representation of real social systems so we have also simulated a more general graph based environment. Our graph is made up of nodes — one node denotes one member of the population — and depending on the connectivity, edges which connect nodes at random.

A general unidirectional graph is a better representation of real systems. Individuals in our simulated environment are connected to other individuals via an edge. Rather than representing distance or adjacency, this edge can be a representation of a different type of relationship. An individual's neighbourhood can be built depending on some social relationship taken into account. Kinship, similarity, mutual interest are typical criteria used when establishing the social components of a community [14].

Initial configuration of cooperators and defectors. The configuration of the population can also be defined. This includes the ratio of cooperators to defectors and where they are located in their environment. They can be placed either randomly, or a specific configuration can be imposed.

Neighbourhood type. The neighbourhood refers to the set of agents with which a player interacts. There exist two traditional neighbourhood types: Von Neumann and Moore. In the Von Neumann neighbourhood, each player interacts with its four nearest neighbours to the north, south, east and west. In the Moore neighbourhood each individual interacts with the eight neighbours reachable by a chess-king's move - i.e. one square in any direction.

Our model simulates these neighbourhoods types as well as a graph type network. In the graph network, individuals are connected to a set of individuals, chosen at random. The size of this neighbourhood depends on the connectivity imposed at the outset of the experiment. Our simulator also allows us to experiment with different radius values in order to determine what effect this has on a system. As the radius, and thus the neighbourhood size increases, we can determine the resulting impact on cooperation.

Player Update. The model enables us to compare many different update rules, some of which are deterministic, some stochastic. Depending on its own score and the scores of its neighbours, each player may update its strategy. Different update rules exist: e.g. Best Takes Over, Imitate the Better and Proportional. Our model is capable of simulating each of these. If strategies are updated using Best Takes Over, a player adopts the strategy of its most successful neighbour. Imitate the Better involves a player comparing its own score with that of it neighbours. If the difference is positive, the player imitates the neighbour with a probability proportional to the difference. The proportional update is where a player adopts the strategy of one of their neighbours with a probability proportional to their scores. Other learning models can also be employed. Update rules can be either deterministic or stochastic. In the stochastic game, update rules will be executed with some degree of probability.

We implement stochastic update rules to discourage swift update convergence which can sometimes lead to a sub-optimal outcome. Deterministic local player update means that after the first round, players are more likely to face opponents playing the same strategy as themselves. This can tend to a bias in favour of cooperators if defectors end up playing defectors. This can quickly drive to extinction strategies that would normally survive and eventually dominate under more stochastic update rules.

Neighbourhood Update. As well as individual learning mechanisms described in the previous section, our model allows us to update and evolve neighbourhoods. The evolution is executed at discrete time steps and in a stocahstic nature to prevent swift update convergence which could lead to a sub-optimal outcome. Depending on a neighbourhood's fitness relative to the average neighbourhood fitness, a new agent can be propagated or indeed an agent can die off. Fitter neighbourhoods are allowed to grow while less fit neighbourhoods decrease in size.

4 Results

In our experiments, we aim to compare agent behaviour in grid and graph topologies. We then extend our graph experiments and examine behaviour as a result of a neighbourhood's ability to grow or shrink relative to neighbourhood fitness.

Initially, we varied the environment size and ran simulations with 100, 400 and 900 agents. Results for the varying sizes showed similar graph patterns emerging for mean values. The results in the rest of this section are based on simulations run for 900 agents.

4.1 Grid topologies

In our experiments, we vary the radius size (R = 1, 2, 3). We bias the initial configuration such that 90% of agents are cooperators. We explore the behaviour of the agents given different update strategies (Best Takes Over and Imitate the Better), across different neighbourhood types (Von Neumann and Moore). The results depicted in Figures 1, 2, and 3 show us that for the Von Neumann neighbourhood with $R = 1$, cooperation can coexist with defectors for all update strategies. This is due to the small neighbourhood size defined and results are consistent with traditional models.

For the Moore neighbourhood type with $R = 1$, levels of cooperation can be maintained for deterministic update rules. However, the stochastic update mechanism results in the emergence of defection. For both deterministic and stochastic updates, the *imitate the better* update strategy showed slightly better results in terms of cooperation for Moore and Von Neumann neighborhoods with a radius of 1.

Due to the increase in neighbourhood size, for all neighbourhood types with a radius value of $R > 1$, defection becomes the norm, with one exception —

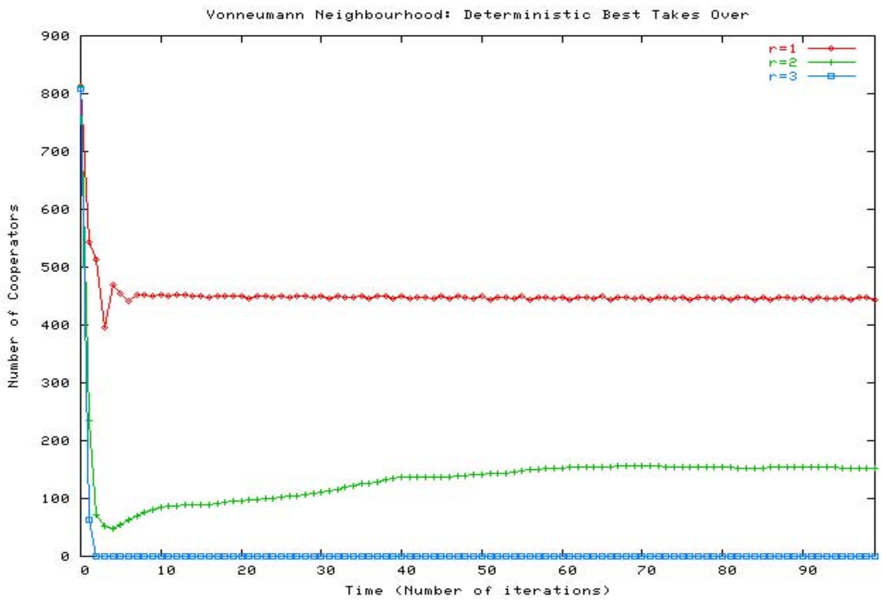

Fig. 1. Von Neumann: Deterministic Best Takes Over

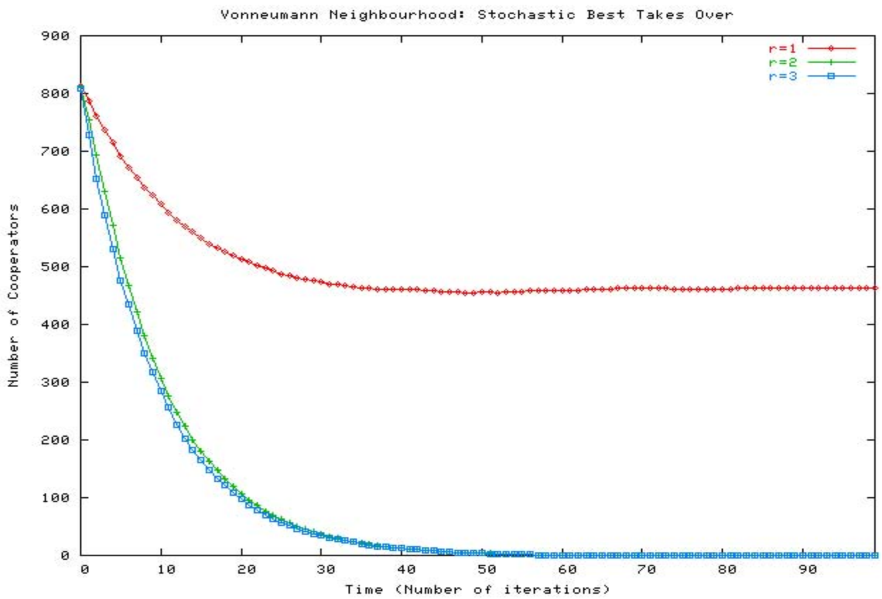

Fig. 2. Von Neumann: Stochastic Best Takes Over

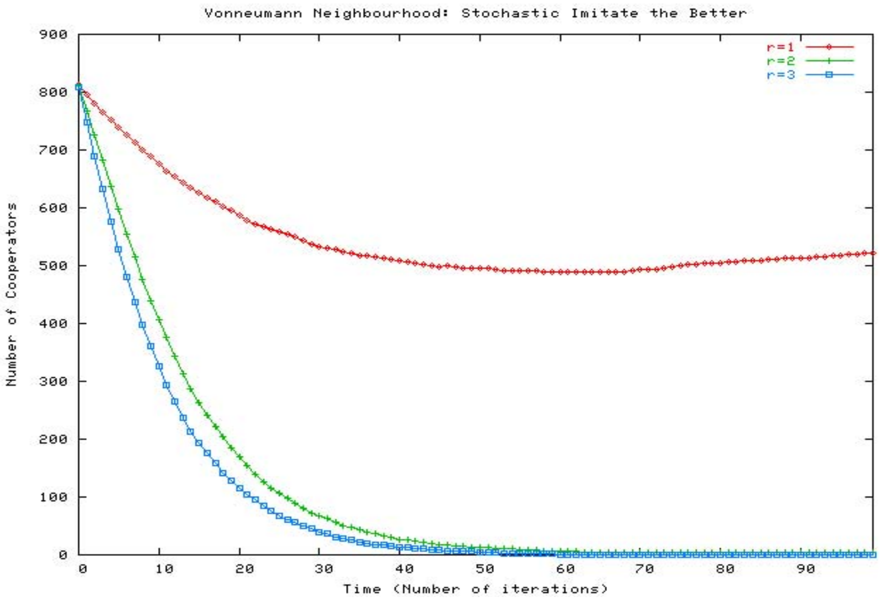

Fig. 3. Von Neumann: Stochastic Imitate the Better

at $R = 2$, the Von Neumann neighbourhood, with a deterministic Best Takes Over update mechanism, a small degree (10% - 15%) of cooperators survive. As expected, the rate at which cooperation decreases is proportional to the increase in R.

4.2 Graph topology

We also explore the behaviour of agents given different update strategies, using the graph topology and neighbourhoods. For deterministic updates widespread defection is almost immediate. For some runs of the experiment, a small number of cooperators coexist with defectors. However, this could be a result of a small neighbourhood of cooperators being disconnected from the rest of the graph.

Using stochastic update rules has little effect on our simulated graph environment (Figures 4 and 5). Convergence is similar but occurs over a longer period of time.

Defection spreads more quickly through our graph topology for a number of reasons. Firstly, depending on the seed, the graph neighbourhood can be much larger than that of the Moore or Von Neumann, and since defection is proportional to neighbourhood size this seems reasonable. Also, our graph topology represents a social system where behaviour can spread more quickly than on traditional lattice structures.

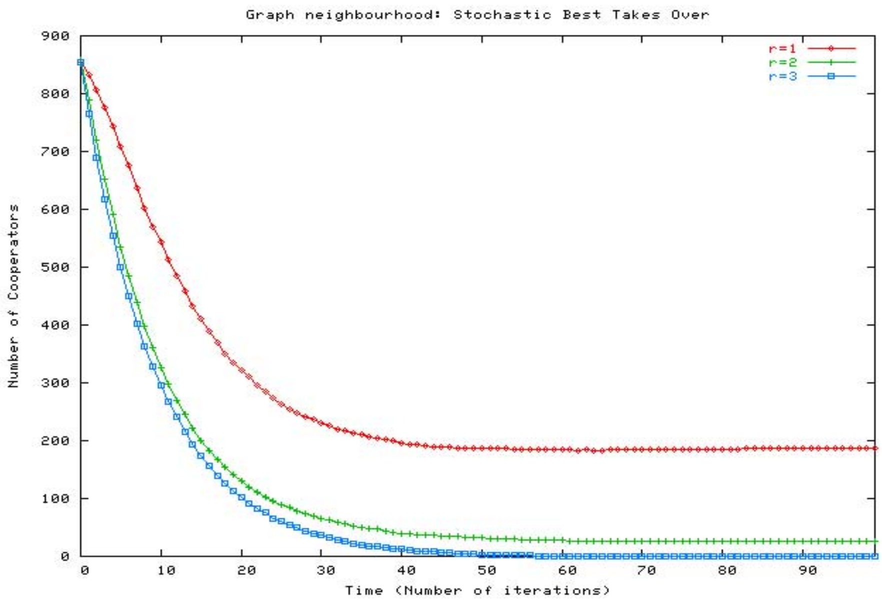

Fig. 4. Graph: Stochastic Best Takes Over

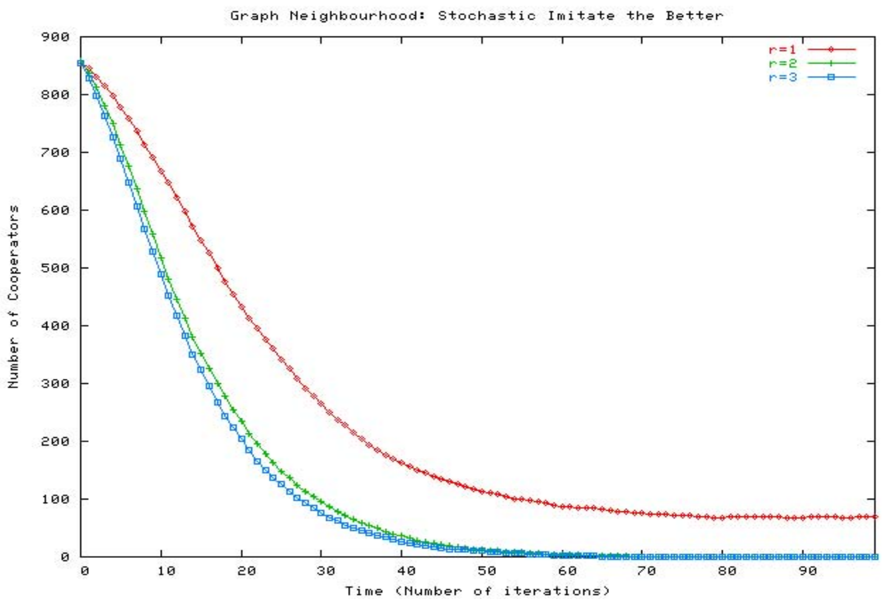

Fig. 5. Graph: Stochastic Imitate the Better

4.3 Graph neighbourhood propagation

In the previous experiments, we updated strategies by allowing agents learn from their neighbours, i.e. strategies which are fit in comparison to the neighbours are propagated. We paid no attention to the fitness of neighbourhoods relative to other neighbouhoods in the population. In this experiment, we allow fit neighbourhoods to grow. A neighbourhood which is fitter than the average neighbourhood is allowed to grow in proportional to its fitness above the average. Similarly, less fit neighbourhoods decrease in size. We investigated the effect of neighbourhood propagation on our graph topology. Figure 6 shows the outcome for one of our experiments. Here we initiated our population to 900 with 95% cooperators and experiments were run for 1000 generations. The resulting graph shows that, initially, defecting strategies quickly flourish and propagate. However, as the simulation continues and neighbourhoods evolve by propagating or shrinking depending on their fitness, we can see the decline of defecting strategies and emergence of a cooperative environment.

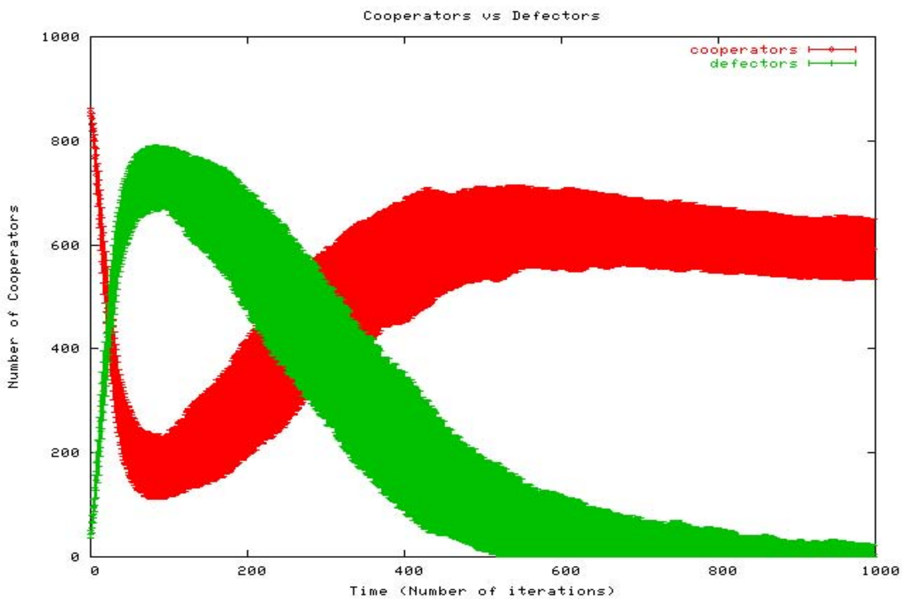

Fig. 6. Graph: Stochastic Imitate the Better including Propagation

5 Conclusion & Future Work

In this paper we explored the parameter space, with particular emphasis on the environment topology. This paper describes a graph topology within which agents can participate in a Prisoner's Dilemma game. A comparison of the graph topology and traditional topologies for varying parameter ranges shows that behaviour spreads more quickly through the graph type environment. We see that in most cases, defection spreads. However, by modifying the graph structure and allowing neighbourhoods to grow or shrink depending on relative fitness, cooperation dominates.

Future work will involve exploring a larger strategy set and fuller investigation into the evolution of the social structure (i.e. the graph topology).

References

1. Hamilton W.D. Axelrod, R. The evolution of cooperation. *Science*, 211(4489):1390–1396, Mar. 1981.
2. R Axelrod. *The Evolution of Cooperation*. Basic Books, New York, USA, 1984.
3. J. M. Epstein. Zones of cooperation in demographic prisoner's dilemma. Technical Report 97-12-094e, Santa Fe Institute, December 1997. available at http://ideas.repec.org/p/wop/safire/97-12-094e.html.
4. R. H. Frank et. al. *Principles of Micro Economics*. McGraw-Hill, 1st edition, 1997.
5. Ch. Hauert. Fundamental clusters in 2 x 2 games. In *Proceedings of the Royal Society B: Biological Sciences*, volume 268, pages 761–769, 2001.
6. Ch. Hauert. Effects of space in 2 x 2 games. *International Journal of Bifurcation and Chaos*, 12:1531–1548, 2002.
7. J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
8. J Maynard Smith. *Evolution and the Theory of Games*. Camridge University Press, Cambridge, England, 1982.
9. M. A. Nowak and R. M. May. Evolutionary games and spatial chaos. *Nature*, 359:826–829, 1992.
10. M. A. Nowak, R. M. May, and S. Bonhoeffer. More spatial games. *International Journal of Bifurcation and Chaos*, 4(1):33–56, 1994.
11. May R.M. Sigmund K. Nowak, M.A. The arithmetics of mutual help. *Acientific American*, 272, 1995.
12. Sigmund K. Nowak, M. Chaos and the evolution of cooperation. In *Proc. Natl. Acad. Sci. USE*, volume 90, pages 5091–5094, June 1993.
13. M. Oliphant. Evolving cooperation in the non-iterated prisoner's dilemma: The importance of spatial organisation. In Maes P. Brooks, R., editor, *Proceedings of the fourth artificial life workshop*, pages 349–352, 1994.
14. J. M. Pujol, R. Sang, and J. Delgado. Extracting reputation in multi agent systems by means of social network topology. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 467–474, New York, NY, USA, 2002. ACM Press.
15. L. Samuelson. *Evolutionary Games and Equilibrium Selection*. MIT Press, Cambridge, MA, March 1997.
16. M. D. Smucker, Stanley, E. Ann, and D. Ashlock. Analyzing social network structures in the iterated prisoner's dilemma with choice and refusal. Technical Report CS-TR-94-1259, University of Wisconsin - Madison, Dept. of Computer Science, December 1994.
17. R. L. Trivers. The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46(1):35–57, 1972.
18. J. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.