Evolving Cooperation in the Non-Iterated Prisoner's Dilemma: The Importance of Spatial Organization

Michael Oliphant
Department of Cognitive Science
University of California, San Diego
e-mail: oliphant@cogsci.ucsd.edu

Published in: Brooks, R. and Maes, P, editors,
Proceedings of the Fourth Artificial Life Workshop,
Cambridge, MA, 1994. MIT Press.

Abstract

Most work on evolving cooperation in the Prisoner's Dilemma treats the non-iterated game as an undesirable simple case that should be risen above. It has been taken as a given that populations evolving to play the non-iterated game will always converge on defection. This paper questions this assumption, and demonstrates that organizing a population spatially dramatically changes the nature of the game and allows cooperation to emerge.

1 The Prisoner's Dilemma

A formalism known as the Prisoner's Dilemma has been adopted as the standard for studying the evolution of cooperative behavior. It describes a game played by two players, where each can decide whether or not they will cooperate with another (failure to cooperate is termed "defection"). A typical example of the payoff each player gets depending on their choice of action is summarized in Figure 1, with the payoff to the row player listed first.

	Coop	Defect
Coop	R/R	S/T
Defect	T/S	P/P

Figure 1: Payoff matrix for the Prisoner's Dilemma

The Prisoner's Dilemma occurs when $T > R > P > S$ and $2R > T + S$ (the most commonly used values are T=5, R=3, P=1, S=0). Although mutual cooperation is the best solution from the joint perspective of both players, defection is the best choice on a self-interested basis. If an individual can expect to encounter a percentage of cooperators Fc and a percentage of defectors Fd (where $Fc + Fd = 1$) over the course of an evaluation cycle, then its expected average payoff would be

$$E[Pc] = R * Fc + S * Fd \quad (1)$$

for cooperators, and

$$E[Pd] = T * Fc + P * Fd \quad (2)$$

for defectors. $T > R$ and $P > S$, so defectors always have an advantage, regardless of the makeup of the population. This mathematical fact is demonstrated if an evolutionary simulation of the Prisoner's Dilemma is done, with populations in such simulations converging to all defectors.

2 Previous Work on Evolving Cooperation

Dissatisfied with mathematical inevitability of defection, Axelrod and many others (Axelrod 1984; Axelrod 1987; Axelrod and Hamilton 1981; Lindgren 1991) have studied a more complex variant of the Prisoner's Dilemma. This variant, called the Iterated Prisoner's Dilemma, involves individuals playing each other more than once. In order to benefit from multiple games against the same opponent, individuals are given a three game history that documents both their actions and those of their opponent. They are also given a mechanism that allows them to modify their future behavior based their history of interaction with an opponent. This allows individuals to cooperate with those who cooperate with them and defect on others, and leads to a predominance of cooperation in the population.

Nowak and May (Nowak and May 1992) took a different approach – attempting to avoid convergence on defection while staying within the confines of the non-iterated game. They used the Prisoner's Dilemma as an update rule for two-dimensional cellular automata. In their simulations, cells could be in one of two states (cooperate or defect) and at each time step, every cell was replaced by the cell bordering it that has the highest summed payoff from playing the game with its eight neighbors. The result they found was not convergence on defection, but intricate patterns of cooperation and defection.

While this result seemed encouraging, it has been criticized as being an artifact of the simulation conditions. Huberman and Glance (Huberman and Glance 1993) argue that the results depend on the use of a synchronous update mechanism where all cells in the space are updated simultaneously. They demonstrate that when asynchronous updating (updating one cell at each time step) is used, the population quickly degrades into all defection.

3 Current Work

The present simulations are aimed at showing that, while the results of the simulations done by Nowak and May

may be artifactual, the idea of using a spatially organized population is a good one and can lead to results similar to what they thought they had found. There are good reasons to expect that organizing a population spatially might benefit cooperators.

3.1 Theoretical Motivation

There is an assumption present in the mathematical explanation (shown in equations (1)&(2)) of the superiority of defection as a choice in the non-iterated Prisoner's Dilemma. The assumption is that individuals will interact with the same percentage of cooperators and defectors regardless of their own game-playing strategy. If this assumption does not hold, neither does the mathematical result. A more generalized version of the payoff equations can be formulated as follows:

$$E[Pc] = R * Fcc + S * Fcd \tag{3}$$

$$E[Pd] = T * Fdc + P * Fdd \tag{4}$$

where Fcc is the percentage of the time a cooperator plays against a cooperator and Fcd is the percentage of the time a cooperator plays against a defector, etc. In the non-spatial, non-iterated case, $Fcc = Fdc = Fc$ and $Fcd = Fdd = Fd$, but this does not hold true in the spatial case.

In the bulk of previous simulations involving the Prisoner's Dilemma, the populations were not spatially organized. This meant that when picking an opponent to communicate with, or when selecting a mate, all other individuals were equally likely to be chosen. In a population where spatial organization is being imposed, this is not the case. Individuals are more likely to play the game with those close to them than they are with those farther away. Also, when an individual has offspring, they are placed in the area of the space where the parent was. What these factors result in is a space where individuals are more related (genetically closer) to those nearer to them.

This means that the expected number of cooperators and defectors an individual will play against depends on their own status, providing a mechanism for cooperation to spread through the population. Cooperators in the population will benefit those close to them. Since those close to them will be more likely to be cooperators as well, they are, in a sense, benefitting their own genes. The opposite holds true for defectors, who will have a negative impact on those around them. Since their neighbors are more likely to be defectors, they hurt their own kind. Getting back to equations (3)&(4), what spatial organization does is to increase the probabilities Fcc and Fdd, thereby increasing the terms $R * Fcc$ and $P * Fdd$. Since $R > P$, this benefits the cooperators.

3.2 Simulations

Instead of using cellular automata, the present simulations use the genetic algorithm (a function optimization technique based on biological evolution developed by John Holland (Holland 1975)). As in Nowak and May's simulations, individuals are fixed as being cooperators or defectors (coded on a one-bit genome). All

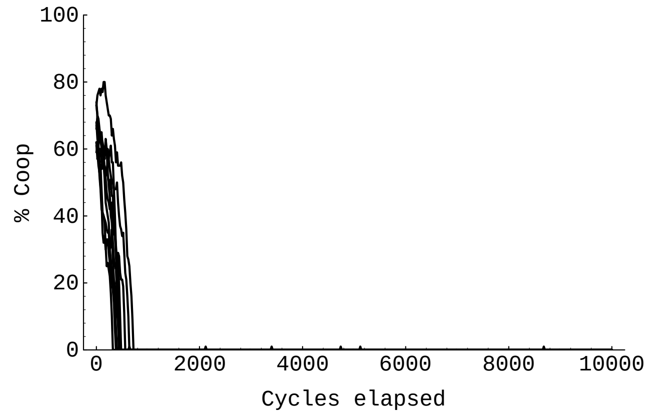

Figure 2: Non-spatial runs with T=5, R=4.5, P=0.5, S=0

simulations use a fixed population size of 100 individuals, a low mutation rate (1 in 1000), and no crossover (as it isn't very useful with when the genome is comprised only of a single bit). The fitness metric used is the average payoff an individual receives when playing the the Prisoner's Dilemma against approximately 32 [1] in the current population. Initial simulations were done using syncronous updating, but to avoid the criticism given by Huberman and Glance of the Nowak and May work, the simulations described here used asynchronous updating. In each cycle of the simulation a single individual was chosen randomly (but biased so that high fitness increased the chance of selection) to reproduce, and a single individual was chosen randomly (but biased so that low fitness increased the chance of selection) to "die" and make room for the new individual.

To replicate the standard all-defect result in the non-spatial, non-iterated game, a series of runs were done without imposing any kind of spatial organization on the structure of the population. The result (10 runs of which are shown superimposed in Figure 2) is as expected. In every run, the population quickly converges to defection.

A very different picture is seen when the same simulations are done with the only modification being to impose spatial organization. Figure 3 shows that each run quickly converges to cooperation, and is maintaining this equilibrium after 10000 cycles have elapsed. In these simulations, the population was organized as a one-dimensional wrap-around space (a ring). Spatial organization was imposed by selecting opponents for an individual to play based on a gaussian distribution around that individual. In this particular set of runs, the distance to an opponent was selected from a distribution with a standard deviation of 1 and incremented by one

[1] An average of 32 games are played by each individual because an individual selects 16 opponents, and is expected to be selected approximately 16 times by other individuals.

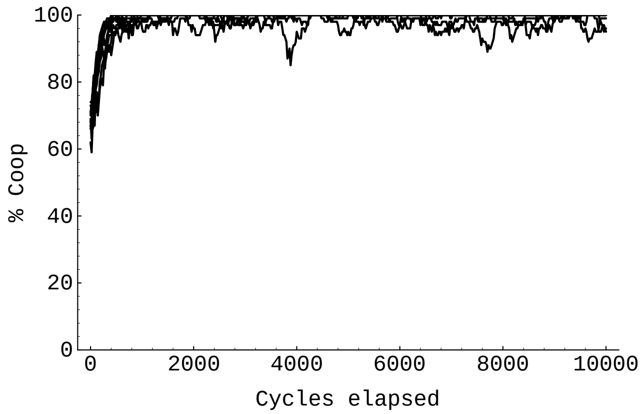

Figure 3: Spatial runs with T=5, R=4.5, P=0.5, S=0

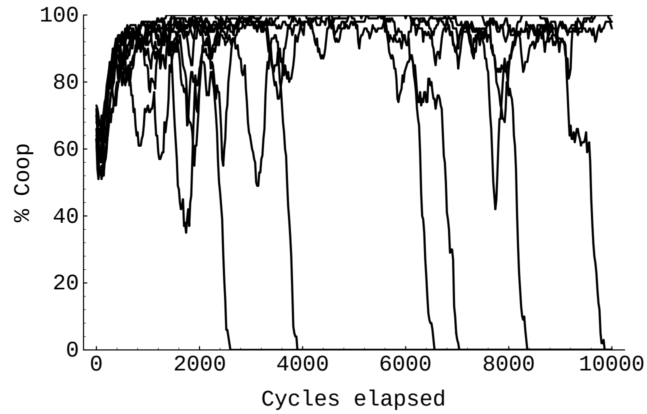

Figure 4: Spatial runs with T=5, R=3, P=1, S=0

(to prevent individuals from playing themselves). Also, offspring were placed in the population at a position near their parents (using the same deviation formula that was used for opponent selection). Using larger standard deviations reduces the spatial organization of the population (and hence the advantage to cooperators), and a sufficiently large deviation makes the population effectively non-spatial.

The results of the previous simulations are somewhat misleading, however, as they used a payoff matrix (T=5, R=4.5, P=0.5, S=0) that gives the defectors only a slight advantage (although enough of an advantage for defectors to easily dominate in the non-spatial case). If this advantage is increased, the results are less clear. Figure 4 shows 10 runs done with a payoff matrix where T=5, R=3, P=1, and S=0. The runs show initial convergence on cooperation, with sharp transitions to defection later on in 6 of the 10 runs. The reason for this is that while a population converged on cooperation is in an equilibrium state, this state is not stable. Enough mutations to defection within a short period of time can put the system in a position to be pulled into the other attractor (convergence on defection), which is a stable state. The probability that a population will be jostled out of the cooperation equilibrium state and then transition to the defection equilibrium state depends on the payoff matrix (which determines the strength of the two attractors) and the mutation rate (where a higher mutation rate is more likely to mutate enough cooperators into defectors in a short period of time to move the population into the attraction region of the all-defect state).

4 Conclusions

The simulations presented in this paper show that when a spatial structure is imposed on populations evolving to play the Non-Iterated Prisoner's Dilemma, a much different pattern emerges. While non-spatial populations quickly fall into defection, spatial populations are able to evolve and maintain cooperative behavior. Because of this, it is a mistake to treat the non-iterated game as a simple case where defectors always have the advantage, and researchers should be careful in motivating their move to more complex forms of the game, such as the iterated version. The iterated game requires a sophisticated game player that can recognize the individuals it has played with previously and remember the other's behavior in those prior interactions – requirements one might not want to assume when trying to apply simulation results to real-world phenomena.

An individual playing the non-iterated game is very simple – all that is required is a hardwired strategy that is either "always cooperate" or "always defect". In demonstrating that cooperation can emerge in the non-iterated game, a lower bound has been placed on the complexity required by an individual to have a population of such individuals evolve cooperative behavior.

References

Axelrod, R. (1984). *The evolution of cooperation.* New York: Basic Books.

Axelrod, R. (1987). The evolution of strategies in the iterated prisoner's dilemma. In L. Davis (Ed.), *Genetic algorithms and simulated annealing,* Chapter 3, pp. 32–41. Los Altos, CA: Morgan Kaufmann Publishers, Inc.

Axelrod, R. and W. Hamilton (1981). The evolution of cooperation. *Science 211*, 1390–1396.

Holland, J. (1975). *Adaptation in natural and artificial systems.* Ann Arbor, MI: The Univ. of Michigan Press.

Huberman, B. and N. Glance (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences (USA) 90*(16), 7715–7718.

Lindgren, K. (1991). Evolutionary phenomena in simple dynamics. In C. Langton, C. Taylor, J. Farmer, and S. Rasmussen (Eds.), *Artificial life II*, pp. 295–311. Redwood City, CA: Addison-Wesley.

Nowak, M. and R. May (1992). Evolutionary games and spatial chaos. *Nature 359*, 826–829.