

The Impact of Payoff Function and Local Interaction on the N -Player Iterated Prisoner's Dilemma

Yeon-Gyu Seo¹, Sung-Bae Cho¹ and Xin Yao²

¹Department of Computer Science, Yonsei University, Seoul, S. Korea;

²School of Computer Science, University of Birmingham, Birmingham, UK

Abstract. The N -player iterated prisoner's dilemma (NIPD) game has been widely used to study the evolution of cooperation in social, economic and biological systems. This paper studies the impact of different payoff functions and local interactions on the NIPD game. The evolutionary approach is used to evolve game-playing strategies starting from a population of random strategies. The different payoff functions used in our study describe different behaviors of cooperation and defection among a group of players. Local interaction introduces neighborhoods into the NIPD game. A player does not play against every other player in a group any more. He only interacts with his neighbors. We investigate the impact of neighborhood size on the evolution of cooperation in the NIPD game and the generalization ability of evolved strategies.

Keywords: Co-evolutionary learning; Iterated prisoner's dilemma; Generalization; Local interaction

1. Introduction

One of the most well-known games for modeling complex social, economical, and biological systems is the iterated prisoner's dilemma (IPD) game (Colman, 1982). In the 2-player IPD game, each player can choose one of the two choices: defection (D) or cooperation (C). Table 1 shows the payoff matrix of the 2-player IPD game. The game is non-zero-sum and non-cooperative. One player's gain may not be the same as the other player's loss. There is no communications between the two players. The game is repeated infinitely. Neither of the players knows when the game is supposed to end.

Received 18 Aug 1999

Revised 27 Feb 2000

Accepted 15 May 2000

Table 1. Payoff matrix of the 2IPD game. $T > R > P > S$, $2R > T + P$.

	Cooperate	Defect
Cooperate	R	S
Defect	T	P

Table 2. The payoff matrix of the NIPD game, where the following conditions must be satisfied: (1) $D_x > C_x$ for $0 \leq x \leq N - 1$; (2) $D_{x+1} > D_x$ and $C_{x+1} > C_x$ for $0 \leq x < N - 1$; and (3) $C_x > (D_x + C_{x-1})/2$ for $0 < x \leq N - 1$. The payoff matrix is symmetric for each player. (Yao and Darwen, 1994).

No. of Cooperators	0	1	...	X	...	$N - 1$
Cooperate	C_0	C_1	...	C_x	...	C_{N-1}
Defect	D_0	D_1	...	D_x	...	D_{N-1}

According to the conditions given in Table 1, if the game is played for one round only, the dominant strategy is definitely defection (D). However, if the game is played for many rounds, mutual defection (D) may not be the dominant strategy. Mutual cooperation (C) certainly performs better than mutual defection. There has been a great deal of work in game theories investigating the dominant strategy for the 2IPD game under various conditions (Colman, 1982). This paper studies a different aspect of the IPD game, i.e., how simulated evolution can be used to learn to play the IPD game without being given any prior knowledge about how to play the game. The focus of this paper will be on evolutionary learning and those factors which influence such a learning process.

The evolutionary approach to strategy learning for the 2IPD game was popularized by Axelrod (1987). Much work following that line assumed that there were only two players involved. Yao and Darwen (Yao and Darwen, 1994; Yao and Darwen, 1995; Darwen and Yao, 1995) were among the first to study the N -player ($N > 2$) IPD game using the evolutionary approach. Up to 16IPD games were investigated (Yao and Darwen, 1994; Yao and Darwen, 1995; Darwen and Yao, 1995). A new encoding scheme for representing the NIPD game strategies was also proposed, which has a much shorter chromosome length than Axelrod's representation (Yao and Darwen, 1994; Yao and Darwen, 1995; Darwen and Yao, 1995).

Table 2 shows the payoff function for the NIPD game, which was defined by Yao and Darwen (Yao and Darwen, 1994; Yao and Darwen, 1995; Darwen and Yao, 1995). The basic principle of the NIPD game is the same as that of the 2IPD game: Defection is always better than cooperation for a single iteration, but mutual cooperation is better than mutual defection in the long run. The third condition in Table 2 was added to prevent a player from alternating between defection and cooperation. These three conditions are very similar to the three in Table 1. However, unlike the case in the 2IPD game where the (single) opponent can be retaliated if he defects, it is impossible to retaliate a defecting player in the NIPD game without affecting others because there is no way to distinguish different players in a group. If a player chooses to defect, he/she will defect against everyone in the group, including cooperators. This makes the NIPD game

		Number of cooperators among the remaining $n - 1$ players				
		0	1	2	$n - 1$	
Player A	C	0	2	4	\dots	$2(n - 1)$
	D	1	3	5	\dots	$2(n - 1) + 1$

Fig. 1. An example of the payoff matrix of the NIPD game (Yao and Darwen, 1994).

much more complex and interesting. The classical tit-for-tat type of strategies will no longer work well in the NIPD game since a defector cannot be identified and penalized effectively.

There are a number of issues in the NIPD game which can influence the evolution of game-playing strategies and the final evolved strategies, such as the payoff function (Banks, 1994), noise (Axelrod and Dion, 1988), population structure (Nowak and May, 1992), localization (Hoffman and Warning, 1996), the history length (Axelrod, 1984; Axelrod and Dion, 1988; Yao and Darwen, 1994), the number of players (Yao and Darwen, 1994), and so on. Most of the work did not use the evolutionary approach. The strategies used by players were largely fixed.

Yao and Darwen studied the NIPD game with N being up to 16 (Yao and Darwen, 1994; Yao and Darwen, 1995; Darwen and Yao, 1995). Using the payoff matrix given by Fig. 1 (Yao and Darwen, 1994), they showed that the evolution of cooperation became more and more difficult as the number of players in a group increased. While cooperation did emerge in the 4IPD game, it never did for the 16IPD game in multiple runs of simulation. Starting from a population of random strategies, the population actually converged to mutual defection (Yao and Darwen, 1994). One of the reasons for the failure to evolve cooperative strategies lies in the NIPD game's inability to identify and punish defectors effectively in a group. In the 2IPD game, a defector can be punished in the next round, which is one of the key points of the tit-for-tat strategy. In the NIPD game, a defector may also be punished in the next round. However, such punishment will hurt the cooperators in the population as they will be punished equally. Such unintended punishment hinders the evolution of cooperation. The larger the group is, the worse the situation will be for cooperators.

This paper examines two issues in the evolution of NIPD game-playing strategies, which have not been studied in detail. The first is to explore different payoff functions from that used by Yao and Darwen (1994). It is found that different strategies emerged when different payoff functions were used. The evolved strategies were able to exploit the characteristics of a payoff function in order to maximize its return. Our experimental results also showed that co-evolutionary algorithms were able to learn different strategies for different payoff functions without human intervention.

The second issue studied in this paper is the impact of local interactions on the evolution of strategies. In the real world, a player may not interact with every other player in a group. He/she may interact only with his/her

immediate neighbors. How does such local interaction affect the evolution? Would a cooperative coalition (Seo and Cho, 1999) emerge in a group although the group as a whole is not cooperative? How robust would the strategies in the coalition be? Would they generalize well? These are among the questions we try to answer in this paper.

According to Schelling (1978), a coalition represents a group of players who choose the non-favorable choice, e.g., cooperation in the IPD game. The minimum coalition size is defined as the least number of players that can begin getting interested in choosing a specified action. We follow such a definition in our investigation of cooperative coalitions.

Generalization is a key issue in evolutionary learning. Little work has been reported on the generalization ability of evolved NIPD strategies except for some limited results (Darwen and Yao, 1995; Darwen and Yao, 1997). This paper will examine how localization (i.e., localized interaction) affects the generalization ability of evolved strategies.

The rest of this paper is organized as follows: Section 2 describes Yao and Darwen's representation scheme for encoding the NIPD game strategies. Section 3 discusses different payoff functions used in our study. Section 4 introduces localization into the NIPD game and explains how it was implemented in our program.

Section 5 presents the experimental results of our study. Several interesting observations can be made from our results. First, different strategies emerged when different payoff functions were used. The evolved strategies were able to exploit the characteristics of a payoff function in order to maximize its return. co-evolutionary algorithms were able to learn different strategies for different payoff functions without human intervention. Second, a small neighborhood size tended to encourage the evolution of cooperation although a neighbor was randomly selected from the population. Third, the history length of a strategy played a more important role in NIPD games with a large neighborhood size than in those with a small neighborhood size. Fourth, the group size (i.e., N) played a less important role in the NIPD games with localization than in those without localization. For example, cooperation did not emerge without localization in the 8IPD and 16IPD (Yao and Darwen, 1994), but it did with localization in this paper. Fifth, strategies evolved with localization tended to perform poorer than those evolved without localization against a group of unseen strong strategies.

Section 6 concludes with a brief summary of major findings of this paper.

2. Encoding Strategies for the NIPD Game (Yao and Darwen, 1994)

One of the most important issues in evolving game-playing strategies is their representation. While encoding strategies for the 2IPD game may be straightforward, representing strategies for the NIPD game can be difficult if we want to have a compact representation. There are two different possible representations, both of which are lookup tables that give an action for every possible contingency.

The first representation is a generalization of the representation scheme used by Axelrod (1987) for the 2IPD game. In this scheme, each genotype is a lookup table that covers every possible history of the last few steps. The Axelrod-style representation scheme for the NIPD game, however, suffers from two drawbacks. First, it does not scale well as the number of players increases. Second, it

provides more information than is necessary by telling which of the other players cooperated or defected, when the only information needed is how many of the other players cooperated or defected. Such redundant information has been shown to have reduced the efficiency of evolution greatly (Yao and Darwen, 1994).

A more compact representation scheme for the NIPD game strategies was proposed by Yao and Darwen (1994). In this scheme, each individual is regarded as a set of rules stored in a lookup table that covers every possible history. A history of l rounds is represented by:

1. l bits for the player's own previous l moves, where a '1' indicates defection, a '0' cooperation; and
2. another $l \log_2 N$ bits for the number of cooperators among the other $N - 1$ players, where N is the number of players in the game. This requires that N is a power of 2.

For example, if we are looking at the 8IPD game with history length 3, then one of the players would see the history as:

History for 8 players, 3 steps: 001 111 110 101 (12 bits).

Here, the 001 indicates the player's own actions: the most recent action (on the left) was a '0', indicating cooperation, and the action 3 steps ago (on the right), was an '1', i.e., defection. The 111 gives the number of cooperators among the other 7 players in the most recent round, i.e., there were $111_2 = 7$ cooperators. The 101 gives the number of cooperators among the other 7 players 3 steps ago, i.e., there were $101_2 = 5$ cooperators.

In the above example, there are $2^{12} = 2048$ possible histories. So 2048 bits are needed to represent all possible strategies. In the general case of the NIPD game with history length l , each history needs $l + l \times \log_2 N$ bits to represent and there are $2^{l+l \times \log_2 N}$ such histories. A strategy is represented by a binary string that gives an action for each of these possible histories. In the above example, the history 001 111 110 101 would cause the strategy to do whatever is listed in bit 1013, the decimal number for the binary 001111110101.

Since there are no previous rounds at the beginning of a game, we have to specify them with another $l(1 + \log_2 N)$ bits. Hence each strategy is finally represented by a binary string of length $2^{l+l(\log_2 N)} + l(1 + \log_2 N)$, which is exponentially shorter than Axelrod-style encoding scheme.

3. Payoff Functions

Figure 1 gives the payoff function used by Yao and Darwen (1994). This section investigates other forms of payoff function and their impact on the evolution of strategies. Generally speaking, the payoff function in the NIPD game should satisfy at least the following conditions:

$$\begin{aligned} C_x &> C_{x-1}, \\ D_x &> D_{x-1}, \\ D_x &> C_x, \\ C_{N-1} &> D_0, \\ N-1 &\geq x \geq 1. \end{aligned}$$

The payoff function is a very important factor in determining the minimum

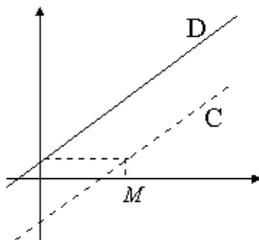


Fig. 2. Minimum coalition size, M , in the NIPD game. The horizontal axis indicates the number of cooperators in a group. The vertical axis indicates the payoff.

coalition size in the NIPD game (Banks, 1994; Schelling, 1978). In Fig. 2, a cooperative coalition makes sense only when the payoff is no worse than all defection. In other words, we need at least M cooperators to form a coalition. M is defined as the minimum coalition size. If M is larger than one half of the total number of players, a cooperative coalition cannot emerge in the game. If M is smaller than one half of the total number of players, a cooperative coalition may emerge depending on the number of players in the game, the payoff function used, etc. It is very useful to investigate different conditions under which a cooperative coalition may or may not be formed.

Although there are three conditions, as given in Table 2, which must be satisfied by the payoff function in order to define the NIPD game (Yao and Darwen, 1994), an occasional relaxation of these conditions enables us to examine other payoff functions which occur in the real world (Schelling, 1978). Figure 3 shows three examples of linear and quadratic payoff functions. Obviously, the quadratic payoff function shown by Figure 3(c) does not satisfy the conditions of the NIPD game, because the payoff of cooperation is higher than that of defection when the number of cooperators is greater than a certain threshold. There is no dilemma any more in this situation. However, when the number of cooperators is below this threshold, it is an NIPD game. It is interesting to see how such a payoff function affects the evolution of strategies.

The payoff function given by Fig. 3(a) is similar to Yao and Darwen's payoff function (Yao and Darwen, 1994). There is a constant difference, i.e., k , between C_x and D_x . In other words, the gradient of the C_x and D_x functions are the same. One variation of this type of payoff function is given by Fig. 3(b), where the gradient of the C_x and D_x functions are different. The D_x function increases more quickly than the C_x function as the number of cooperators (x) goes up. So the more cooperators in the group, the more attractive it is to defect.

4. Localization

The only interaction among players in the NIPD game is through playing games against each other in each generation. How well a player does depends on how strong/weak his opponents are. Localization restricts the opponents of a player to be selected only from a small neighborhood (i.e., among his neighbors). This restriction will understandably influence how the player is to interact with others and what payoff he is going to get. Such restriction (localization) also influences

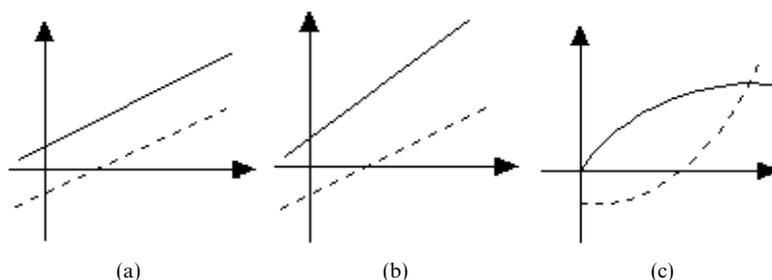


Fig. 3. Three different types of payoff function. Solid lines represent the payoff function for defection and dashed lines represent the payoff function for cooperation. The horizontal axis indicates the number of cooperators. The vertical axis indicates the payoff. (a) $C_x = 3x - k$ and $D_x = 3x$. (b) $C_x = 2x - k$ and $D_x = 3x + 1$. (c) $C_x = \frac{1}{2}x^2 - k$ and $D_x = \sqrt{2}x$. $0 \leq x \leq N$.

evolutionary learning of game-playing strategies because fit strategies (fit in the context of localized interaction) will reproduce and propagate in the population even though they are only fit locally. Very fit strategies without localization may not be fit anymore with localization. We would guess that strategies learned with localization would be less robust than those learned without it since players have less exposure to a wide range of different opponents. However, the NIPD game with localization is a more realistic model of many real world systems. It is interesting to find out the difference between NIPD games with and without localization.

Related work on the NIPD game with localization includes Nowak and May's work on spatial evolution (Nowak and May, 1992). In their model, a population of players are distributed on squares of a torus who are only capable of always defection (AD) and always cooperation (AC). Each player interacts only with his eight neighbors and imitates the strategy of any better performing one. Cooperative behavior can be sustained in clusters of players that insulate cooperators from hostile AD players under certain payoffs (Nowak and May, 1992). In our study, the evolutionary approach can explore a much larger set of different strategies (all possible strategies of a given history length), not just AD and AC. Strategies in a population are constantly changing from generation to generation under the influence of crossover, mutation and selection. It is a much more complex NIPD game that we are studying in this paper.

Hoffman and Warning (1996) have also experimented with localized interaction and learning between players on torus-employing Moore machines to represent strategies and reported that localized learning and interaction improve the cooperativity within a population.

There are two models, as shown by Fig. 4, which can be used in the study of localization in the NIPD game. We will use the one-dimensional model as shown by Fig. 4(a) in this paper. We will examine how the neighborhood size and the history length affect the emergence of cooperative coalition and the generalization ability of evolved strategies.

There are several methods of selecting $N - 1$ opponents among a player's neighbors in the NIPD game. Random selection will be used in our study. In other words, if one player for a game is determined, the other $N - 1$ players are selected at random among his neighbors. It is worth pointing out that such

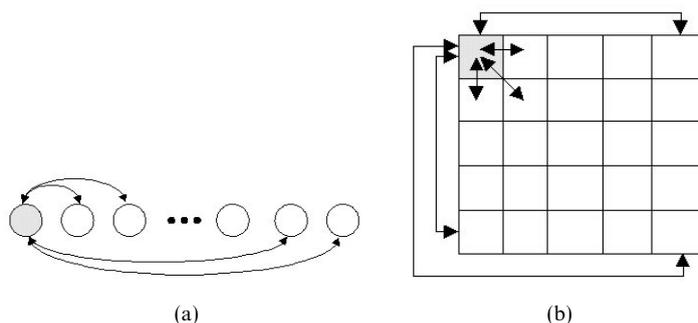


Fig. 4. Localization of interaction and learning. (a) One-dimension case with neighborhood size 2. (b) Two-dimension case with neighborhood size 1.

random selection leads to sparse evaluation of a player since he does not see all different groups of opponents. What he learns will be biased towards what he has seen. This is very similar to the situations in our social and biological environments.

5. Experimental Results

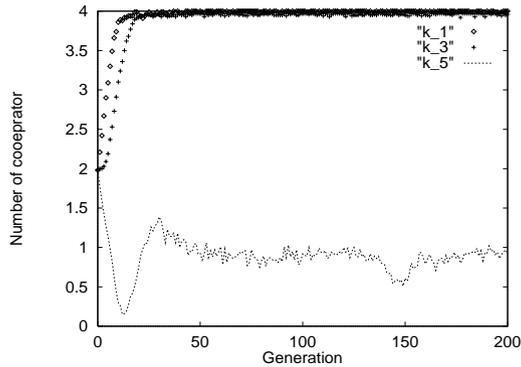
For all our experiments, we used an evolutionary algorithm with population size 100, crossover rate 0.6, and mutation rate 0.001. Two-point crossover with elite preserving strategy was adopted. The selection scheme used was ranking.

5.1. NIPD Games with Different Payoff Functions

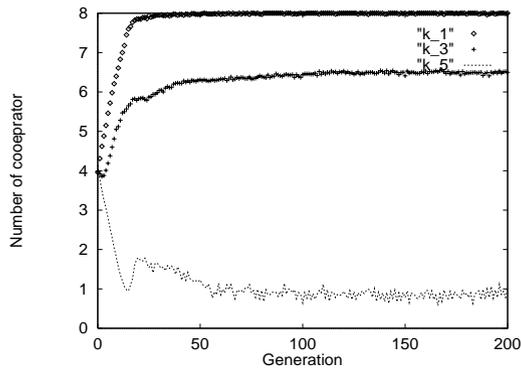
We have applied our evolutionary algorithm to a number of NIPD games using three different types of payoff function as described in Section 3. We used $N = 4$ and 8. Three different k values (i.e., $k = 1, 3, 5$) were used for each type of payoff function. In total, 18 different NIPD games were simulated. Each game was repeated for 10 independent runs. All the results reported in this section have been averaged over 10 runs.

Figure 5 shows the experimental results using the payoff function given by Fig. 3(a), i.e., $C_x = 3x - k$ and $D_x = 3x$, where $0 \leq x \leq N - 1$ and $k = 1, 3, 5$. This type of payoff function is the more typical one, similar to the one used by Yao and Darwen (1994). Two observations can be made from Fig. 5. First, a small k encourages the evolution of cooperation. As k becomes larger, the payoff of defection becomes so high in comparison with that of cooperation that most players are willing to take the risk (of being punished/retaliated) and defect. That is what happened when $k = 5$.

Second, a small group encourages the evolution of cooperation. Cooperation is easier to emerge in a small group of four players than in a large group of eight players. This confirms Yao and Darwen's earlier finding about the group size (Yao and Darwen, 1994). However, what Yao and Darwen did not explain clearly is the cooperative coalition. For example, the population was not fully cooperative in the 8IPD game with $k = 3$, as shown by Fig. 5(b). However, there was a sustained coalition of six cooperators in the 8IPD game.



(a)

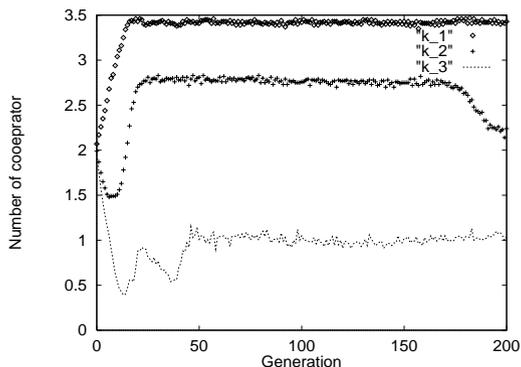


(b)

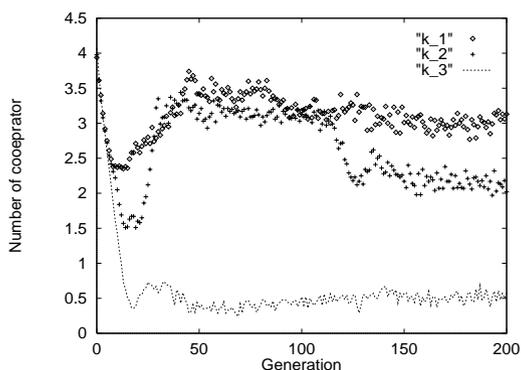
Fig. 5. The evolution of strategies for the NIPD game using the payoff function $C_x = 3x - k$ and $D_x = 3x$, where $0 \leq x \leq N - 1$. $N = 4, 8$. $k = 1, 3, 5$. (a) Four players; (b) eight players.

Figure 6 shows the experimental results of the NIPD games using the payoff function given by Fig. 3(b), i.e., $C_x = 2x - k$ and $D_x = 3x$, where $0 \leq x \leq N - 1$ and $k = 1, 3, 5$. An immediate observation from Fig. 6 is that the average number of cooperators dropped in comparison with the results in Fig. 5. Moreover, the larger the group size, the more decrease in the number of cooperators. For example, the average number of cooperators in the 8IPD with $k = 1$ dropped substantially from eight shown by Fig. 5(b) to less than four shown by Fig. 6(b). Such a decrease in the number of cooperators was less dramatic for the 4IPD game. This is because the more cooperators we have in the NIPD game, the larger the payoff difference will be between defection and cooperation. Hence it is more attractive to defect. It is hard to maintain a high level of cooperation within a group in this case.

Another interesting observation from Fig. 6 is that the number of cooperators in a group was much less stable in comparison with the results in Fig. 5. The players appeared to be more cautious and ready to defect whenever opportunities arose. For example, the number of cooperators varied greatly between two and four for the 8IPD game even with $k = 1$, as shown in Fig. 6(b). The seemingly



(a)



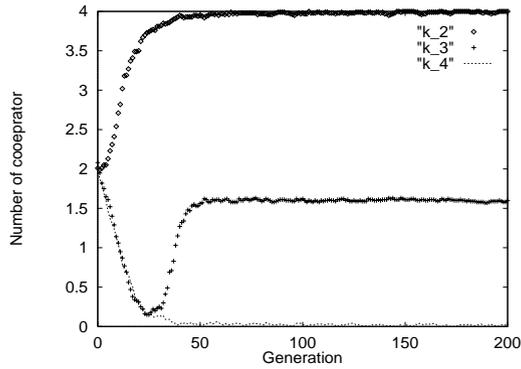
(b)

Fig. 6. The evolution of strategies for the NIPD game using the payoff function $C_x = 2x - k$ and $D_x = 3x$, where $0 \leq x \leq N - 1$. $N = 4, 8$. $k = 1, 3, 5$.

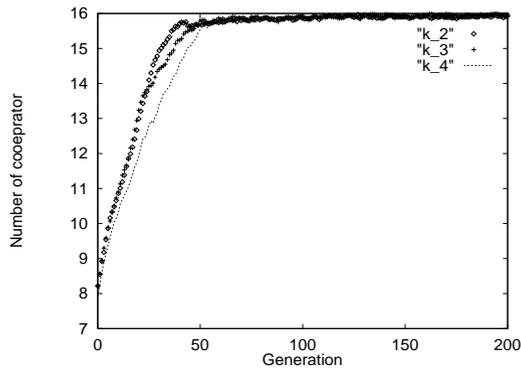
stable cooperative coalition of (average) 2.7 players suddenly collapsed around 170 generations for the 4IPD game with $k = 3$, as shown in Fig. 6(a). The typical size of a cooperative coalition was between 1.5 and 3.5 on average, which was substantially smaller than the coalition sizes observed in Fig. 6.

Figure 7 shows the experimental results of the NIPD games¹ using the payoff function given by Fig. 3(c), i.e., $C_x = \frac{1}{2}x^2 - k$ and $D_x = \sqrt{2x}$, where $0 \leq x \leq N - 1$ and $k = 2, 3, 4$. According to the payoff function, C_x grows much faster than D_x as x , i.e., the number of cooperators goes up. Although the payoff for cooperation is usually smaller than that for defection when x is small due to $-k$, it catches up quickly. The time it takes to go past the payoff for defection depends on the value of k . The smaller the k , the quicker it is. This is shown clearly in Fig. 7(a) for the 4IPD games. In the case of 8IPD games, the payoff for having more cooperators was so high that every player decided to cooperate. We would

¹ Strictly speaking, these payoff functions do not define an NIPD game because they do not satisfy all three conditions given in Table 2.



(a)



(b)

Fig. 7. The evolution of strategies for the NIPD game using the payoff function $C_x = \frac{1}{2}x^2 - k$ and $D_x = \sqrt{2x}$, where $0 \leq x \leq N - 1$. $N = 4, 8$. $k = 2, 3, 4$. (a) Four players; (b) eight players.

have to use a fairly large k (probably around 27) if we wanted to observe some non-cooperative behaviors in the 8IPD game.

In general, there may be oscillation in co-evolutionary process. The fitness of strategy is evaluated by other strategies in the population. At the beginning of the game, there may be defective strategies in the population and many cooperative strategies do not know how to deal with the strategies, so many cooperative strategies are hurt by defective strategies. However, as the time goes by, there come out conditional cooperative strategies that are adaptive to defective strategies. The more these strategies come out, the less the defective strategies are. Conditional strategies can get more payoff better than defective strategies, because they cooperate with cooperative strategies, while they defect against defective strategies. Once an appropriate number of conditional cooperators remain and a large number of cooperators is established in the population, though there are defective strategies, cooperators can survive in the next generation by interaction among cooperators. This is because mutual cooperation gets a better payoff than mutual defection.

5.2. NIPD Games with Localization

Localization limits a player to interact only with his neighbors. This section presents experimental results of the NIPD games with different neighborhood sizes, history lengths and group sizes (i.e., N). The one-dimensional model as shown by Fig. 4(a) will be used here. The main purpose is to investigate how the neighborhood size influences the evolution of cooperative strategies. The issue of history length and group size has been studied by Yao and Darwen (Yao and Darwen, 1994; Seo et al, 1999) for the NIPD games without any localization. We will use the same payoff function as that used by Yao and Darwen (1994), i.e., $C_x = 2x$ and $D_x = 2x + 1$, where $0 \leq x \leq N - 1$ and $N = 4, 8, 16$. The history length is varied from one to three to investigate the effect. Here, the player's partner is randomly selected among his adjacent neighbors.

Figure 8 shows the experimental results of the 4IPD games with different neighborhood sizes and history lengths. It is obvious that the neighborhood size played a critical role in the evolution of cooperation. When the neighborhood size was small (i.e., two), all runs converged to mutual cooperation regardless of the history length, although it appeared to take longer to cooperate when the history length was large (i.e., three), as shown by Fig. 8(a), (c) and (e). This is rather interesting as the previous study showed that a longer history length encouraged cooperation (Yao and Darwen, 1994). When the neighborhood size was increased to 30, the history length became much more important. A short history length (i.e., one) led to many non-cooperative behaviors, as observed from Fig. 8(b). Some runs did converge to mutual cooperation, but others did not. The number of cooperators appeared to fluctuate greatly in the evolution. In some runs, the number of cooperators climbed to a fairly high level (around 3.5) after the initial dip, and then gradually decreased again. It seemed to be very difficult to maintain a stable cooperative coalition no matter what the coalition size was. As the history length was increased to two (Fig. 8(d)) and then to three (Fig. 8(f)), most players tended to benefit from the extra memory and seemed to realize that mutual cooperation was better than mutual defection. Hence we can observe the increase in cooperation among players. For example, most runs in Fig. 8(f) converged to mutual cooperation among all four players. For the two runs where mutual cooperation among all four players was not achieved, there was a stable cooperative coalition of (average) 3.3 players.

Yao and Darwen have shown that it becomes more difficult to evolve cooperative strategies as the group size (i.e., N) increases in the NIPD game (Yao and Darwen, 1994). Consensus is more difficult to achieve in a large group. Figure 9 shows some interesting results on what happened when each player was restricted to interact only with his neighbors. With such localization, the group size (i.e., N) did not seem to have much impact on the evolution of cooperation anymore, as indicated by Fig. 9(b) and 9(d). Instead, the neighborhood size had a critical impact on the evolution. For example, Yao and Darwen (1994) could not evolve any cooperative strategies for the 16IPD games; we could evolve cooperative strategies in a coalition even with a rather large neighborhood size of 30, as shown by Fig. 9(c). When the neighborhood size was reduced to eight, the coalition size was very close to 16, as shown by Fig. 9(d). In other words, almost all players were cooperating with each other. This was impossible without localization (Yao and Darwen, 1994).

One of the key issues in evolutionary learning is generalization (Darwen and Yao, 1995). It is important to examine how well evolved strategies perform against

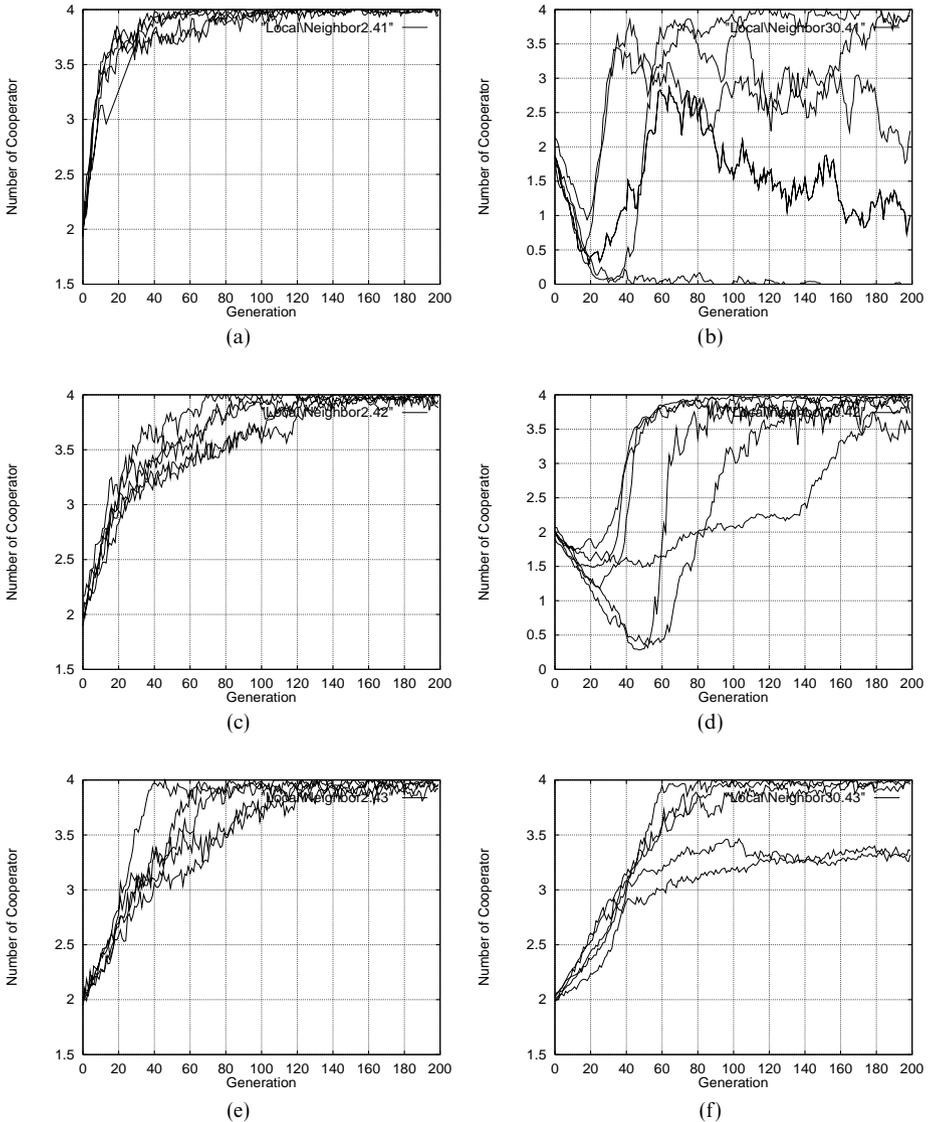


Fig. 8. The 4IPD games with different neighborhood sizes (2 and 30) and different history lengths (1, 2, and 3). (a) Neighborhood size 2 and history length 1. (b) Neighborhood size 30 and history length 1. (c) Neighborhood size 2 and history length 2. (d) Neighborhood size 30 and history length 2. (e) Neighborhood size 2 and history length 3. (f) Neighborhood size 30 and history length 3.

different opponents. We select two sets of strategies as potential opponents of evolved strategies in this paper. The first set consists of the 15 best strategies from a randomly generated population of 100 strategies. The second set consists of evolved strategies themselves. This approach of testing the generalization ability of evolved strategies follows that used by Yao and Darwen (Yao and Darwen, 1994; Darwen and Yao, 1997). The 8IPD game with history length 1 and neighborhood size 4 will be used.

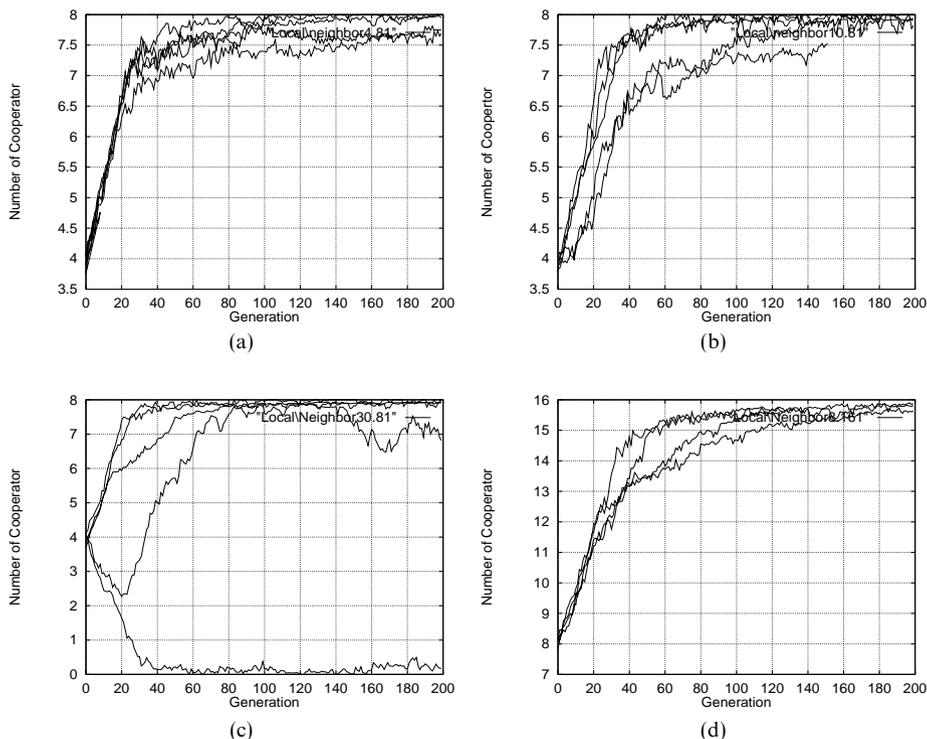


Fig. 9. The NIPD games with $N = 8, 16$ and neighborhood size 4, 8, 10 and 16.

Table 3 gives the 15 best strategies which were selected from a random population of 100 strategies for the 8IPD game. Table 4 shows the game-playing results between the evolved strategies and the 15 best strategies given by Table 3. Table 5 shows the game-playing results among the evolved strategies themselves. It is obvious from Tables 4 and 5 that the evolved strategies generalized poorly against best random strategies although they did very well against themselves. In comparison with the results presented by Yao and Darwen where no localization was used (Yao and Darwen, 1994; Darwen and Yao, 1997), it is very clear that localization had created a small and cosy local environment in which all strategies tended to cooperate with other in order to gain high payoff. The strategies were seldom exposed to a more diverse and hostile environment. That is why they were so vulnerable and performed so poorly against best random strategies. In fact, they performed much worse than the strategies evolved without any localization (Yao and Darwen, 1994) when playing against the best random strategies. However, they did spectacularly well among themselves.

6. Concluding Remarks

Localization and payoff function can have a dramatic impact on evolved strategies and their generalization ability in the NIPD game. In general, a larger difference between the payoff for defection and that for cooperation will make non-cooperative moves more attractive and hence make a cooperative coalition

Table 3. The 15 best random strategies selected from a population of 100 random strategies.

Opponent strategy				
1110	1100	1111	1001	1001
1010	1001	0111	1101	0110
1100	1111	1001	1101	0100
0010	1111	1001	1101	1000
1101	1010	1111	1111	1010
0000	1111	1111	1101	1110
0001	1011	1101	1100	1100
0111	0111	0111	1111	0010
0110	1110	0101	1011	1010
0011	1011	1111	1111	1011
1101	1101	1101	1111	0100
1010	0011	1111	1100	1011
1101	0100	0011	1100	0100
1110	1100	1101	1101	1000
0011	1110	1011	1110	0111

Table 4. The game-playing result between evolved strategies and the best random strategies given by Table 3.

Evolved strategy	Self		Opponents	
	Mean	S.D.	Mean	S.D.
00011100101001000001	4.364	0.268	5.336	0.183
00001001000010001011	4.137	0.194	5.575	0.160
01001000011111010011	4.366	0.201	5.446	0.197
00011000010100111110	4.268	0.241	5.486	0.211
00100100111000110000	4.064	0.180	5.452	0.146
00000000100101000000	4.219	0.200	6.315	0.157
00111011011110101000	4.795	0.191	5.025	0.124
00100011010010100100	4.769	0.158	5.497	0.136
01101110100011011110	4.856	0.201	4.929	0.209
00000100010011000110	4.286	0.229	5.965	0.196

smaller, if it exists at all. Localization can encourage the evolution of cooperation. The smaller the neighborhood size is, the easier it is for cooperation to emerge, even in the NIPD game with a large N (i.e., $N = 16$). The neighborhood size appears to be a much more important factor than the group size (i.e., N) in influencing the evolution of cooperation. The history length plays a role only when the neighborhood size is sufficiently large.

Although localization can encourage the evolution of cooperation and achieve very high payoff for cooperative strategies, they are very brittle and generalize poorly against best random (non-cooperative) strategies. The cooperative coalition tend to create a small and cosy environment for strategies to pat each other's back. As a result, strategies lose their competitiveness against a wide range of different opponent strategies.

The NIPD game is an extremely rich and complex model for many real world applications. What we have studied do far are all based on a relatively simple model where each player can only have two choices: cooperation or defection. In reality, few people cooperate entirely (i.e., 100%) with others or defect entirely

Table 5. The game-playing result among the evolved strategies themselves.

Evolved strategy	Self		Opponents	
	Mean	S.D.	Mean	S.D.
00011100101001000001	14.000	0.000	14.000	0.000
00001001000010001011	14.000	0.000	14.000	0.000
01001000011111010011	14.000	0.000	14.000	0.000
00011000010100111100	14.000	0.000	14.000	0.000
00100100111000110000	14.000	0.000	14.000	0.000
00000000100101000000	14.000	0.000	14.000	0.000
00111011011110101000	14.000	0.000	14.000	0.000
00100011010010100100	14.000	0.000	14.000	0.000
01101110100011011110	14.000	0.000	14.000	0.000
00000100010011000110	14.000	0.000	14.000	0.000

against others. There exist many different levels of cooperation. One person may cooperate with another 80% or 40%. The latest study along this line (Yao and Darwen, 1999) has shown some very interesting results which are quite different from what we found in the simple binary case. This is one of the promising directions for future study.

Another promising direction of future study is the investigation of a player's reputation in the NIPD game. We do not interact with everyone in a human society, yet we tend to trust and cooperate with some people more often than with others even though we have never interacted with them. Reputation is one of the factors which affect our choice of cooperation or defection in this situation. The NIPD game with player's reputation added (Yao and Darwen, 1999) enables us to study a number of more realistic and interesting behaviors which one might observe in real life.

Acknowledgements. This work has been supported partially by a grant (975-0900-004-2) from the Korea Science and Engineering Foundation and a grant from the Australian Research Council (to the third author).

References

- Axelrod R (1984) The evolution of cooperation. Basic Books, New York
- Axelrod R (1987) The evolution of strategies in the iterated prisoner's dilemma. In Davis L (ed). Genetic algorithms and simulated annealing. Morgan Kaufmann, San Mateo, CA, pp 32–41
- Axelrod R Dion D (1988) The further evolution of cooperation. *Science* 242: 1385–1390
- Banks S (1994) Exploring the foundation of artificial societies: experiments in evolving solutions to iterated N-player prisoner's dilemma. In Brooks R, Maes P (eds). *Artificial life IV: Proceedings of the fourth international workshop on the synthesis and simulation of living systems*. MIT Press, Cambridge, MA, pp 337–342
- Colman AM (1982) Game theory and experimental games. Pergamon Press, Oxford
- Darwen PJ Yao X (1995) On evolving robust strategies for iterated prisoner's dilemma. In Yao X (ed) *Progress in evolutionary computation. Lecture notes in artificial intelligence*, vol 956. Springer, Berlin, pp 276–292
- Darwen PJ, Yao X (1997) Speciation as automatic categorical modularization. *IEEE Transactions on Evolutionary Computation* 1(2): 101–108
- Hoffmann R, Warning N (1996) The localization of interaction and learning in the repeated prisoner's dilemma. Technical report 96-08-064, Santa Fe Institute working paper
- Nowak MA, May RM (1992) Evolutionary games and spatial chaos *Nature* 359: 826–829
- Schelling TC (1978) *Micromotives and macrobehaviors*. Norton, New York

- Seo YG, Cho SB (1999) An evolutionary study on cooperation in n-person iterated prisoner's dilemma game. In *Simulated Evolution and Learning*. Lecture notes in artificial intelligence, vol 1585. Springer, Berlin, pp 301–308
- Seo YG, Cho SB, Yao X (1999) Emergence of cooperative coalition in NIPD game with localization of interaction and learning. In *Proceedings of congress on evolutionary computation (CEC'99)*, Piscataway, NJ. IEEE Press, pp 877–884
- Yao X, Darwen P (1994) An experimental study of N-person iterated prisoner's dilemma games. *Informatica* 18: 435–450
- Yao X, Darwen PJ (1995) An experimental study of N-person iterated prisoner's dilemma games. In Yao X (ed). *Progress in Evolutionary Computation*. Lecture notes in artificial intelligence, vol 956. Springer, Berlin, pp 90–108
- Yao X, Darwen P (1999) How important is your reputation in a multi-agent environment? In *Proceedings of the 1999 IEEE international conference on systems, man, and cybernetics (SMC'99)*, Piscataway, NJ, October 1999. IEEE Press, pp II-575–II-580

Author Biographies



Yeon-Gyu Seo received a B.Sc. degree in computer science from Suncheon National University, Suncheon, Korea, in 1997 and an M.S degree in computer science from Yonsei University in 1999. Since then, he has been studying for his Ph.D. in computer science at Yonsei University. He is interested in genetic algorithms, neural networks, and agent systems. He is a member of the Korea Information Science Society, Korea Cognitive Science Society and Korea Information Processing Society.



Sung-Bae Cho received a B.Sc. degree in computer science from Yonsei University, Seoul, Korea, in 1988 and M.S. and Ph.D. degrees in computer science from KAIST (Korea Advanced Institute of Science and Technology), Taejeon, Korea, in 1990 and 1993, respectively. He worked as a Member of the Research Staff at the Center for Artificial Intelligence Research at KAIST from 1991 to 1993. He was an Invited Researcher of Human Information Processing Research Laboratories at ATR (Advanced Telecommunications Research) Institute, Kyoto, Japan from 1993 to 1995, and a Visiting Scholar at the University of New South Wales, Canberra, Australia in 1998. Since 1995, he has been an Associate Professor in the Department of Computer Science, Yonsei University. His research interests include neural networks, pattern recognition, intelligent

man-machine interfaces, evolutionary computation, and artificial life. Dr Cho was awarded outstanding paper prizes from the IEEE Korea Section in 1989 and 1992, and another one from the Korea Information Science Society in 1990. He was also the recipient of the Richard E. Merwin prize from the IEEE Computer Society in 1993. He was listed in *Who's Who in Pattern Recognition* from the International Association for Pattern Recognition in 1994, and received the best paper awards at International Conference on Soft Computing in 1996 and 1998. Also, he received the best paper award at the World Automation Congress in 1998, and was listed in *Marquis Who's Who in Science and Engineering* in 2000. He is a member of the Korea Information Science Society, INNS, the IEEE Computer Society, and the IEEE Systems, Man, and Cybernetics Society.



Xin Yao received his B.Sc. from the University of Science and Technology of China (USTC) in 1982, his M.Sc. from the North China Institute of Computing Technologies (NCI) in 1985, and his Ph.D. from USTC in 1990. He is currently a professor of computer science at the University of Birmingham, England, and a guest professor of USTC. He was a lecturer, senior lecturer and associate professor at University College, the University of New South Wales, and the Australian Defence Force Academy (ADFA), between 1992 and 1999. He held post-doctoral fellowships from the Australian National University (ANU) and the Commonwealth Scientific and Industrial Research Organization (CSIRO) between 1990 and 1992. He is an associate editor or a member of the editorial board of six international

journals, including *IEEE Transactions on Evolutionary Computation* and *Knowledge and Information Systems: An International Journal*, and an editor/co-editor of seven journal special issues. He has chaired/co-chaired a number of international conferences in evolutionary computation and artificial intelligence in recent years. His major interests include evolutionary learning, co-evolution, evolutionary design and evolvable hardware, neural network ensembles, global optimization, simulated annealing, computational complexity and data mining.

Correspondence and offprint requests to: X. Yao, School of Computer Science, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK. Email: x.yao@cs.bham.ac.uk