
On the Relationship between the Duration of an Encounter and the Evolution of Cooperation in the Iterated Prisoner's Dilemma

David B. Fogel

Natural Selection, Inc.

3333 N. Torrey Pines Ct., Suite 200

La Jolla, CA 92037

e-mail: fogel@sunshine.ucsd.edu

Abstract

Evolutionary programming experiments are conducted to examine the relationship between the durations of encounters and the evolution of cooperative behavior in the iterated prisoner's dilemma. A population of behavioral strategies represented by finite-state machines is evolved over successive generations, with selection made on the basis of individual fitness. Each finite-state machine is given an additional evolvable parameter corresponding to the maximum number of moves it will execute in any encounter. A series of Monte Carlo trials indicates distinct relationships between encounter length and cooperation; however, no causal relationship can be positively identified.

Keywords

Iterated prisoner's dilemma, refusal to play, evolutionary programming.

1. Introduction

Simulated evolution provides a method for studying the interactions of coevolving individuals within a population (Hillis, 1992; Kauffman & Johnson, 1992; Ray, 1992; Angeline & Pollack, 1994; Rosen & Belew, 1995; and others). In such simulations, the fitness criterion for evaluating the behavior of an individual is defined as a function of the other extant individuals. Thus, the fitness of an individual can change as a temporal, stochastic process. Interest is often focused on identifying populational patterns of behavior over time and determining causal factors that induce such patterns. The broad hope is that explanations of individual and higher-level interactions in evolutionary systems might be forthcoming from idealized models that abstract the salient features of such systems. Recently, particular attention has been given to coevolutionary simulation in the study of social interactions.

Many of these simulations have centered on the iterated prisoner's dilemma, a two-person, nonzero-sum, noncooperative game. The term *nonzero-sum* indicates that benefits to one player do not imply similar losses to the other; the term *noncooperative* indicates that there is no preplay communication between the players. Two players meet for a series of moves, in which each may choose either to cooperate or to defect. *Cooperating* implies increasing the reward of both players; *defecting* implies increasing one's own reward at the expense of the other player. The typical payoff matrix for such interactions is shown in Figure 1. The Prisoner's Dilemma can be a useful tool for analyzing factors associated with

		Player B	
		C	D
Player A	C	(γ_1, γ_1)	(γ_2, γ_3)
	D	(γ_3, γ_2)	(γ_4, γ_4)
		C = Cooperate D = Defect	

Figure 1. The general payoff matrix of the Prisoner's Dilemma. Each entry (α, β) indicates the payoffs to player A and player B. γ_1 is the payoff for mutual cooperation, γ_2 is the payoff for cooperating when the other player defects, γ_3 is the payoff for defecting when the other player cooperates, and γ_4 is the payoff for mutual defection. The standard constraints on these payoffs are (1) $2\gamma_1 > \gamma_2 + \gamma_3$ and (2) $\gamma_3 > \gamma_1 > \gamma_4 > \gamma_2$.

the evolution of cooperation and aggression in social communities because it is conceptually simple, yet captures essential relationships between the relative payoffs for the alternative behaviors.

If the game is played for a single iteration, defection is the dominant move. Further, if the game is iterated for any fixed number of plays, with the duration known to each player *a priori*, the logical result is a continuous series of defections. Rational players will reason that they should defect on the final play. Knowing this, they will further reason that because the other player will defect on the final play, they should defect on the next-to-last play, and so forth. However, simulations with the iterated Prisoner's Dilemma involving long but fixed-length encounters with round-robin interaction across a population of players have not yielded endless defection. The results have been just the opposite: mutual cooperation has been the more prevalent result (Axelrod, 1987; Miller, 1989; Fogel, 1991, 1993; and others), although such cooperation has not always been stable (Lindgren, 1991). It appears reasonable to speculate that the logical structures involved in the previous simulations may not have been capable of learning the fixed duration of the game, or even learning that the game was of fixed duration, and therefore the evolutionary dynamics did not lead to endless mutual defection.

Three questions arise in light of previous efforts: (1) If the duration of an encounter is included as part of the behavioral strategy, will long encounters evolve at all? (2) Will cooperation be as prevalent when encounter length is adaptive as when long encounters

		Player B	
		C	D
Player A	C	(3,3)	(0,5)
	D	(5,0)	(1,1)

C = Cooperate
D = Defect

Figure 2. The traditional payoff matrix used in experiments with the iterated Prisoner's Dilemma. This matrix was also used for the current simulations.

are mandated? (3) What is the correlation between the mean duration of encounters in a population and the mean fitness of the population? These are the questions addressed by the current simulations.

2. Background

Previous simulations are rooted in the foundational tournaments described in Axelrod (1980a, 1980b). Axelrod conducted two tournaments involving the iterated Prisoner's Dilemma described by the payoff matrix in Figure 2. Strategies were sent to Axelrod and were played in a round-robin competition in which all played against all others, against themselves, and against a purely random strategy. It is now well known that the strategy Tit-for-Tat (i.e., cooperate on the first move, then imitate the opponent's preceding move) was the winner of both tournaments. Axelrod (1984) noted that 8 of 64 entries in the second tournament could be used to account reasonably for how well any specific strategy performed. Axelrod (1987) used a genetic algorithm to evolve strategies against these 8 representatives, with evolution proceeding on a static adaptive landscape. In subsequent experiments, the strategies were made to evolve against themselves as an instance of evolution on a dynamic adaptive landscape that changed at each generation in light of the current set of parents and offspring. Each strategy was a deterministic function of the three previous moves and all games iterated for 151 plays (the average length of the two tournaments from 1980). Lindgren (1991) investigated a similar approach with more limited memory of previous moves. Other related work has been offered by Miller (1989) and Fogel (1991, 1993), in which populations of finite-state automata (Figure 3) were evolved using genetic algorithms and evolutionary

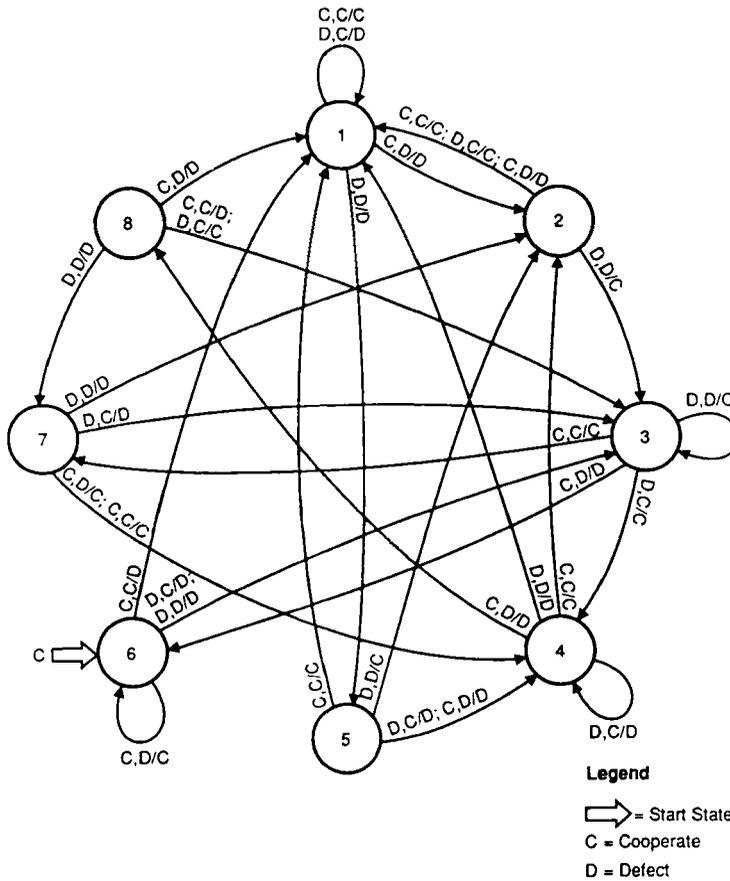


Figure 3. A finite-state machine used to represent a strategy in the Prisoner's Dilemma. The set of input symbols is $\{(C, C), (C, D), (D, C), (D, D)\}$, where the first entry indicates the machine's previous move and the second entry indicates the opponent's previous move. The set of output symbols is $\{C, D\}$, the two alternatives being *cooperate* or *defect*. Input symbols are shown to the left of the virgule. Output symbols are shown to the right of the virgule. The arrow shows the start state and starting play. This representation follows previous work in Fogel (1991, 1993).

programming, respectively (see Fogel, 1995, for a review of different methods in evolutionary computation); Fogel and Harrald (1994) evolved neural networks operating on the three previous moves to offer a continuous range of behavior between complete cooperation and defection. However, none of these simulations addressed the effect of an evolving encounter length between players. Axelrod (1984, 12–13) discussed the inclusion of a *discount parameter* to devalue future moves in relation to the current move, but did not include one in the experiments described in Axelrod (1987).

The previous investigation that appears most relevant to the current questions has been offered by Stanley, Ashlock, and Tesfatsion (1994).¹ Evolving strategies were allowed choice

¹ Dugatkin (1992) offered a version of the Prisoner's Dilemma involving spatial groups of interacting strategies and roving defectors,

and refusal of partners using expected outcomes based on previous encounters in light of the payoff matrix shown in Figure 2. Strategies in a population were given the potential to choose their opponents rather than to meet in round-robin or random encounters. Only those strategies that offered to play each other actually played out the sequence of moves in the game until either player's future payoff expectation fell below a given threshold, at which point the game was terminated. Strategies that were refused by others received a small *rejection* payoff; strategies that were completely inactive received a *wallflower* payoff. Stanley et al. (1994) indicated that the inclusion of this choice-and-refusal mechanism could accelerate the emergence of cooperative behavior, and often formed *payoff bands* (i.e., payoffs concentrated in a given range of values).

Rather than use the choice-and-refusal mechanism of Stanley et al. (1994), the current effort is an extension of previous research in Fogel (1991, 1993). Finite-state machines are again used to represent the behavioral logic of each player in the population (Figure 3), but an additional variable is also included that corresponds to the maximum length of any encounter for that machine. The duration of play between any two machines is determined to be the minimum of the two encounter lengths (i.e., one player breaks off the game). Fitness is considered as the average payoff of each move in all games across all players; selection is applied after each complete set of interactions to determine which machines survive to generate progeny in the next generation. Analysis of repeated patterns in several trials of such a simulation offers potential answers to the questions posed above.

3. Method and Materials

The experiments described in this paper were conducted as follows: An initial population of 100 *parent* finite-state machines was created. The number of states for each machine was determined in accordance with a uniform random variable over the integers [1, 5]. All outputs and state transitions for each input symbol were chosen uniformly at random. The initial encounter-length parameter for each machine was set to two moves (i.e., the initial play and one subsequent play); this was also set as the minimum encounter length.²

Each parent created a single offspring machine through two forms of random mutation (both were applied to each parent). The first mode of mutation had six varieties: (1) add a state, (2) delete a state, (3) change the start state, (4) change an output symbol, (5) change a state transition, and (6) change the initial move. Each of these was given an equal probability of occurrence. Deleting a state and changing the start state were only allowed for machines of two or more states. Adding a state was allowed for machines having seven or fewer states (i.e., the maximum number of states was eight). For any added state, one state transition was chosen at random from the other states and redirected to the new state, thereby offering the potential for the new state to be expressed. The second mode of mutation varied the encounter length of the offspring by adding a Gaussian random variable with zero mean and variance of 5 (arbitrarily chosen to give a standard deviation slightly larger than the minimum encounter length) to the parent's encounter length. The resulting new encounter length was truncated to an integer. The encounter length was constrained to lie within [2,150] moves; if the mutated value fell outside these limits, it was reset to the appropriate limit it had exceeded.

in which the defectors were akin to con artists who move from location to location once their "cover is blown." The added complexity of spatial groupings makes it difficult to assess the relationship of that work to the present effort, although important similarities can be recognized.

² Two moves was chosen as a lower limit to require an iterated Prisoner's Dilemma. A lower bound of one move corresponds to the "one-shot" Prisoner's Dilemma.

All 200 finite-state machines competed in round-robin pairing (i.e., every machine met every other machine once). All encounters were iterated up to the minimum of both machines' encounter-length values. The payoff matrix in Figure 2 was used in all interactions. The fitness of each machine was assessed as the mean payoff per move averaged across every encounter. After all machines had competed, the 100 machines with the greatest fitness scores were selected to become parents of the next generation. This process was iterated for 4,000 generations to examine the long-term stability of any potential outcomes. Forty-nine trials were executed.

4. Results

The patterns of cooperation and defection generated in each of the 49 trials varied considerably. This precluded examining the average behavior across all trials, because the average was not representative of any particular trial. However, general patterns of behavior could be identified under some arbitrary definitions of what constituted cooperation or defection. For the purposes of this analysis, mutual cooperation was considered to have occurred whenever the mean fitness of all surviving parents at a given iteration exceeded 2.6 (complete mutual cooperation yields a mean fitness of 3.0), whereas mutual defection was taken as indicated by a mean fitness below 1.4 (complete mutual defection yields a mean fitness of 1.0).

Analysis was first given to the long-term stability of cooperation or defection. Each trial was assessed to determine whether or not it had generated 1,000 consecutive iterations of mutual cooperation or defection across the series of 4,000 generations. In all, 16 out of 49 trials generated long-term mutual cooperation (samples shown in Figure 4) and 15 out of 49 trials generated long-term defection (samples shown in Figure 5); some generated both conditions. As might be expected in light of these results, there were often rapid changes from cooperation to defection and vice versa. With *rapid change* defined as a mean fitness varying from under 1.4 to over 2.6, or the obverse, within 50 generations, 36 out of 49 trials generated rapid change (samples shown in Figure 6). Five trials did not generate any long-term stable cooperation or defection, nor did they generate any rapid changes between the two. Only three trials never generated any cooperation at all (i.e., never exceeded a mean fitness of 2.6).

Attention was subsequently focused on the mean encounter length of all surviving parents. Mean encounter lengths of over 140 moves were generated in 42 of the 49 trials. Since the maximum number of moves allowed was 150, it is reasonable to conclude that there was a strong tendency to generate long encounters within 4,000 generations, despite all initial encounter lengths' having been set to two.

Assessing the relationship between the mean encounter length and the mean fitness was more difficult. In many cases, there appeared to be a direct relationship, but a scatterplot of mean encounter length versus mean fitness for each iteration over the first three trials (i.e., 12,000 moves) suggested that such straightforward analysis would not be useful (Figure 7). Although a simple correlation test would yield statistically significant results, the regression line of best fit is clearly not a suitable explanation of the data, because there is almost no agreement between the line and the data for encounter lengths between 50 and 110.

A better understanding of the relationship between the two parameters can be obtained by examining the successive moves in which changes from mutual defection to mutual cooperation occurred, and also by assessing the relationship of the mean encounter length to those changes. In particular, because the populations tended toward mutual defection initially, attention was focused on the first rise from mutual defection to mutual cooperation.

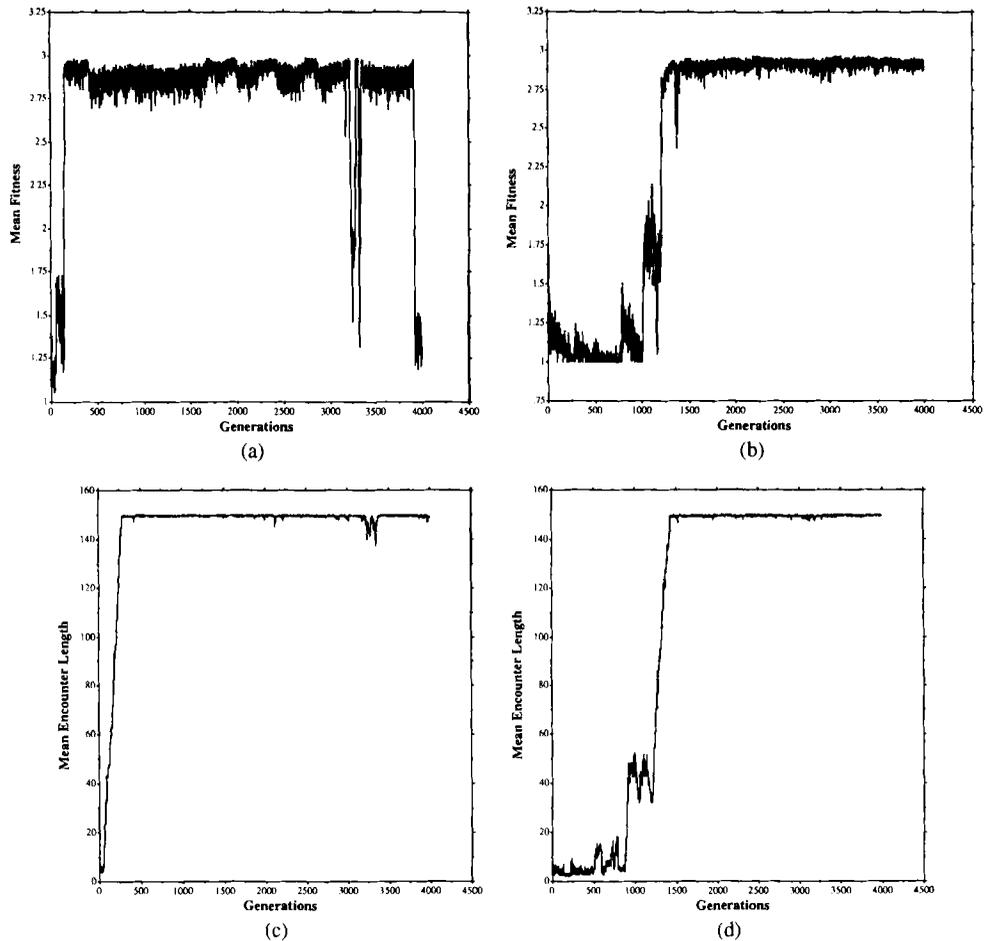


Figure 4. Two representative trials out of 16 that generated long-term mutual cooperation. Graphs (a) and (b) indicate the mean fitness of all surviving parents at each generation in trials 4 and 11, respectively. In each case, the mean fitness stayed above the threshold of 2.6 for more than 1,000 consecutive generations, although there were deviations from such mutual cooperation. For completeness, graphs (c) and (d) indicate the mean encounter length (the average of each surviving machine's maximum encounter length) for each of these trials, respectively.

Figure 8 shows a scatterplot of the mean encounter length for the last mean fitness value of less than 1.4 that was followed by the first rise in mean fitness to a value of greater than 2.6 versus the mean encounter length at the time when the mean fitness exceeded 2.6. It also shows a scatterplot of the mean encounter length for the last mean fitness value below 1.4 versus the size of the increase in the encounter length, for all trials that demonstrated such a rise (see Figure 9 for graphical explanation).³ In almost all cases, the mean encounter length

³ This description is somewhat complex, but the procedure is simple. Pairs of points are obtained for Figure 8a by finding the first point for which a mean fitness exceeds 2.6. The mean encounter length at this point is recorded. Then the evolution is retraced to the last time the mean fitness was below 1.4 (if it exists). The mean encounter length at this point is recorded; this gives the first pair of points. For Figure 8b, the mean encounter length at 2.6 is replaced by the size of the increase over the mean encounter length at 1.4.

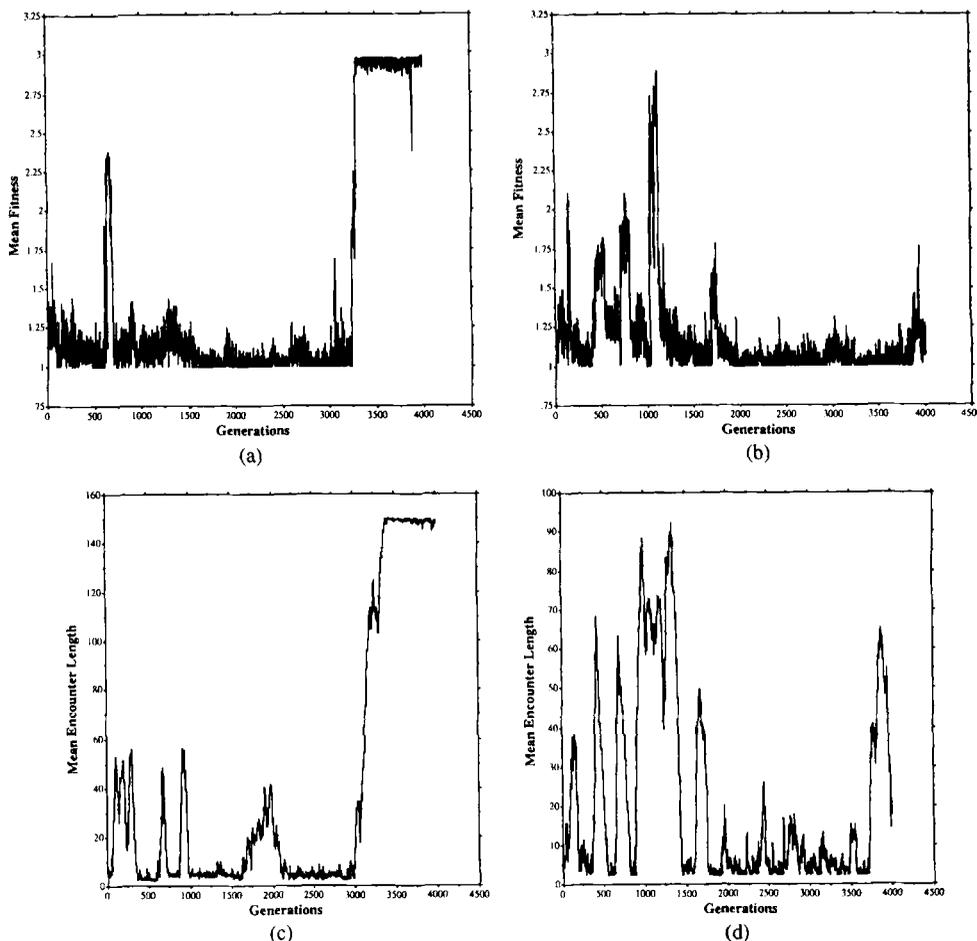


Figure 5. Two representative trials out of 15 that generated long-term mutual defection. Graphs (a) and (b) indicate the mean fitness of all surviving parents at each generation in trials 2 and 23, respectively. In each case, the mean fitness stayed below the threshold of 1.4 for more than 1,000 consecutive generations, although there were deviations from such mutual defection. For completeness, graphs (c) and (d) indicate the mean encounter length (the average of each surviving machine’s maximum encounter length) for each of these trials, respectively.

was shorter before the rise in mean fitness from mutual defection to mutual cooperation. That is, an increase in the mean encounter length was generally associated with an increase in mean fitness.

Figures 10 and 11 show histograms compiled from Figure 8. Figure 10 demonstrates that long mean encounters were not sufficient to generate mutual cooperation, because in several trials the mean encounter length exceeded 60 (or even 140) without generating mutual cooperation. The histogram indicates a potential exponential decline of the number of observations as the encounter length is increased, reminiscent of a random walk away from a barrier (possibly biased toward the barrier). In contrast, Figure 11 (mean encounter length when first achieving a mean fitness greater than 2.6) shows no such decline. The

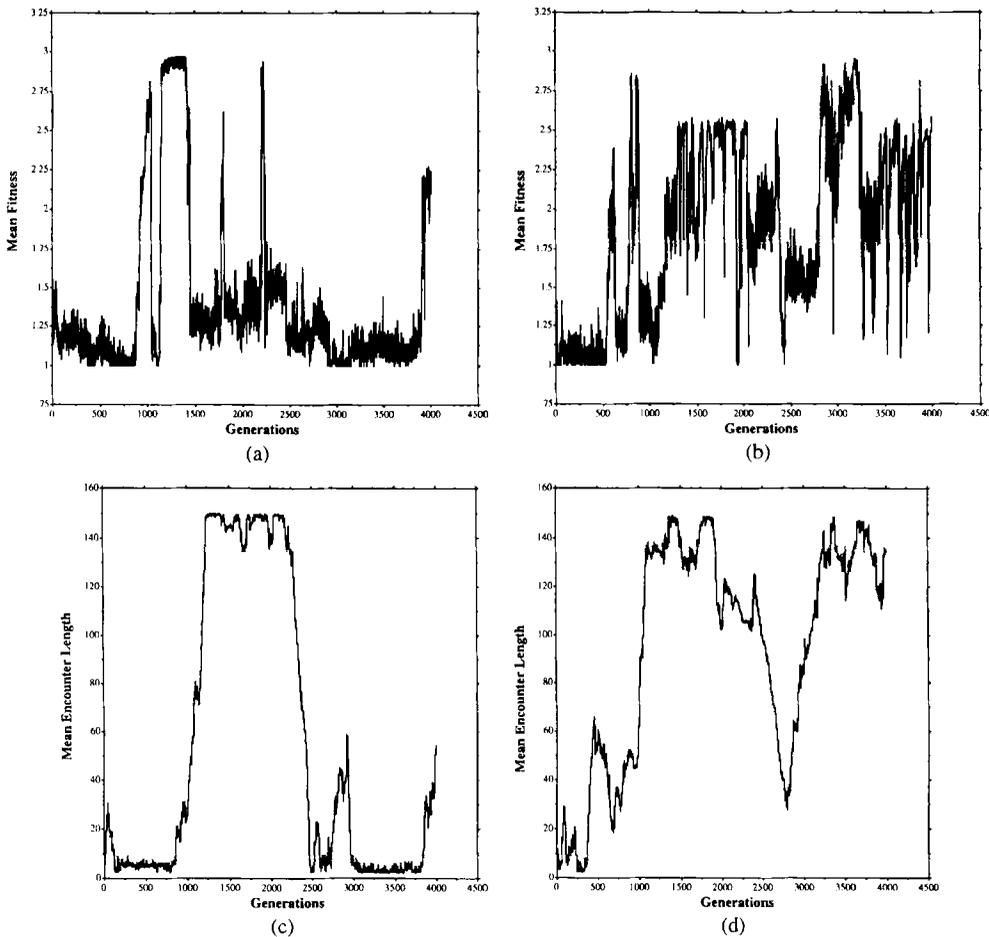


Figure 6. Two representative trials out of 36 in which the mean fitness changed at least once over 4,000 generations from below 1.4 (mutual defection) to above 2.6 (mutual cooperation), or vice versa, within 50 generations. Graphs (a) and (b) indicate the mean fitness of all surviving parents at each generation in trials 5 and 16, respectively. For completeness, graphs (c) and (d) indicate the mean encounter length (the average of each surviving machine's maximum encounter length) for each of these trials, respectively.

graph provides no support for the idea that mean encounter lengths greater than 140 are more effective in generating mutual cooperation than are mean encounter lengths of 40. Although there did appear to be some minimum duration for encounters before mutual cooperation would be generated (the shortest observed mean encounter length for mutual cooperation was 18.33), there was no requirement that excessively long encounters precede or accompany mutual cooperation. These conditions are better illustrated by Figure 12, which shows the relationship between the first rise to mutual cooperation and the mean encounter length in the first three trials.

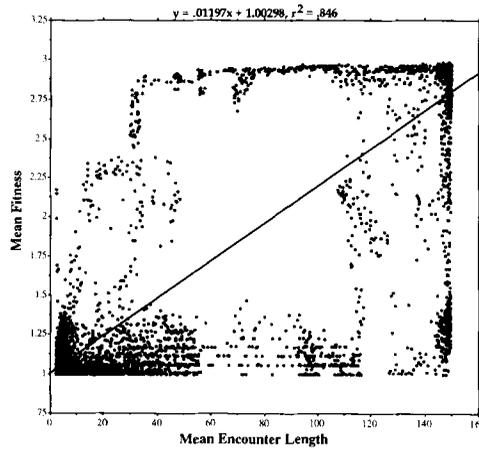


Figure 7. Scatterplot of the mean encounter length and the corresponding mean fitness of the population at each generation of the first three trials (12,000 total data). Although the best least square regression line (the equation appears above the scatterplot) is statistically significant ($P < 0.0001$, $R^2 = 0.846$), it is not appropriate for describing these data, because it does not model the data pairs well when the mean encounter length varies between 50 and 110.

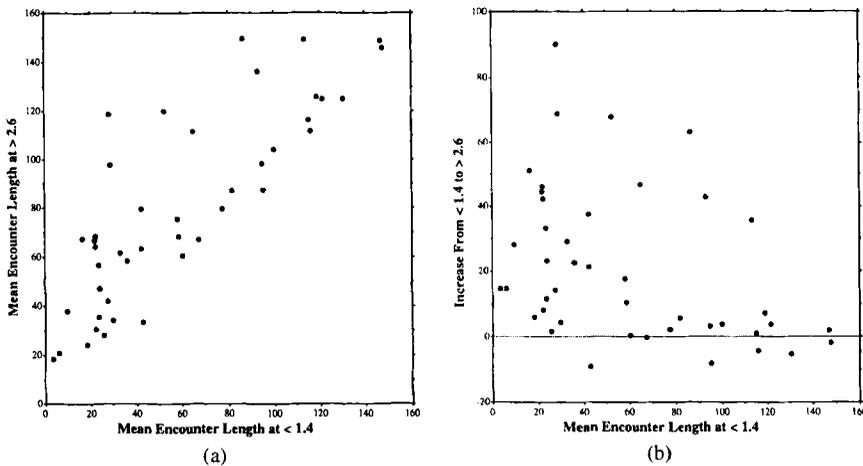


Figure 8. (a) Scatterplot of the mean encounter length at the last generation with a mean fitness of less than 1.4 paired with the mean encounter length at the first generation with a mean fitness of greater than 2.6, for the first rise from mutual defection to mutual cooperation. (b) Scatterplot of the difference between the mean encounter length at mean fitness greater than 2.6 and the mean encounter length at the last generation with a fitness of less than 1.4. The smaller the mean encounter length is at the time of the rise from defection to cooperation, the greater the increase.

5. Conclusions

Both mutual cooperation and long encounters can be generated under the conditions of the above-described experiment, despite the potential for any of the evolving strategies to adapt

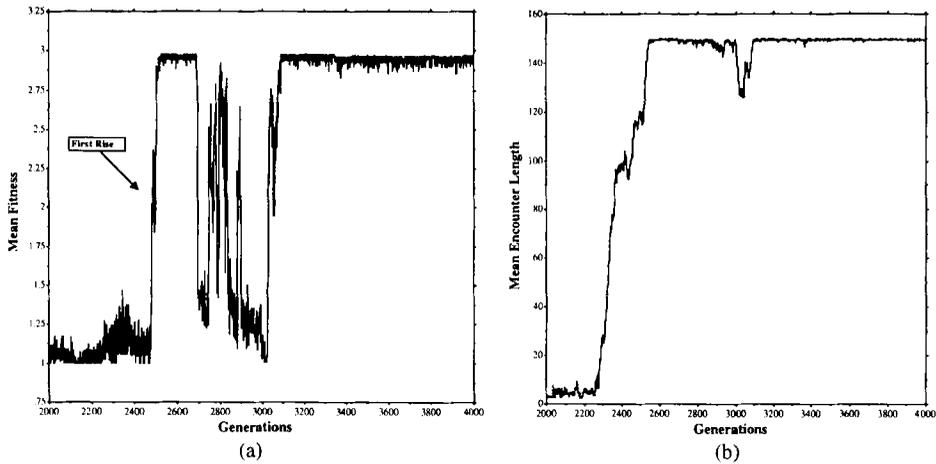


Figure 9. The first rise from mutual defection to mutual cooperation. (a) Plot of mean fitness versus generations for generations 2000–4000 from trial 3 (also see Figure 6). The rise in mean fitness that occurs around the 2500th generation is the first rise to mutual cooperation (mean fitness > 2.6). (b) Plot of the mean encounter length versus generations over the same sample for trial 3. The rapid increase in the mean encounter length precedes the rise in mean fitness.

their encounter length and break off encounters very early in the engagements. However, the tendency to generate mutual defection was comparable to the tendency to generate cooperation; this finding appears to differ from the general cooperation observed when long encounters were mandated (Fogel, 1993). Moreover, neither cooperation nor defection was a particularly stable outcome.

The relationship between mean encounter length and mean fitness appears more complex. It may best be summarized as follows:

- The mean encounter length can increase up to the maximum bound without being accompanied by mutual cooperation. Encounters of long duration do not guarantee cooperative behaviors.
- With the current representation, mutual cooperation is unlikely to arise under very short encounters; however, a trend toward increasing fitness and away from mutual defection can arise under very short encounters.
- The amount of increase in the mean encounter length during the first rise from mutual defection to mutual cooperation is dependent on the mean encounter length that exists at the time of the rise. The longer the mean encounter length before the rise, the less likely it will be to increase during the rise.
- Departures from mutual cooperation can occur without any significant decrease in the mean encounter length. Long departures from mutual cooperation tend to drive the mean encounter length down.

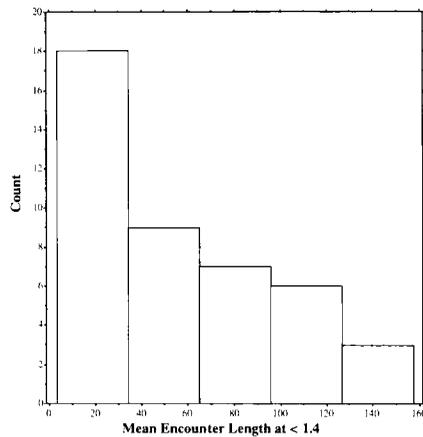


Figure 10. Histogram of the mean encounter length at the last generation in which the mean fitness of the population was below 1.4 before the initial rise to mutual cooperation. The histogram shows an apparently exponential rate of reduction as the mean encounter length increases. Such a result is similar to what would be expected if the mean encounter length were determined by a random walk away from a barrier (i.e., the minimum length). Note, though, that the functional fit to these data with an exponential model yields an R^2 value of 0.942, whereas a logarithmic fit has a corresponding R^2 value of 0.968. The exponential form could be favored because of its asymptotic properties, but in this case the random walk would have an upper barrier at 150, so the asymptotic properties may not be particularly important. Although an exponential model is reasonable, the logarithmic model cannot be discounted. Note also that here the random walk would be concluded when the mean fitness started to increase toward mutual cooperation, and this may be more likely at any generation when the mean encounter length is longer. Thus, there could be an additional mechanism leading to the decline evidenced in the histogram.

These experiments demonstrated more instability than was reported in Stanley et al. (1994). This is likely due to their sophisticated mechanism of estimating future expected payoffs, used to decide when to break off encounters. It is not clear that either model is more reasonable than the other, although explicit estimation of future payoffs requires an assumption of greater information processing capability for both strategies.

A question remains regarding identification of the precursors of mutual cooperation and defection. An increase in the mean encounter length preceded an increase in mean fitness in many cases, but not in all. In some trials, a rise in mean fitness occurred before an increase in the mean encounter length, but in others, an increase in the mean encounter length did not correspond to any increased fitness at all. Further, rapid departures from mutual cooperation and defection occurred often—even under long-duration encounters, in which greater support for stability in mutual cooperation would be expected. It seems reasonable to expect that the answers to these questions will require careful examination of the parent and offspring finite-state machines over successive generations just before the first rise to mutual cooperation and any departure from it.

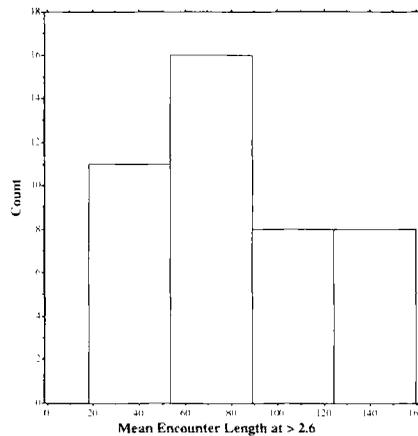


Figure 11. Histogram of the mean encounter length at the first generation in which the mean fitness of the population was above 2.6 after the initial rise to mutual cooperation. The histogram does not display the nearly geometric rate of reduction as the mean encounter length increases. Mean encounter lengths of about 20 would appear to be sufficient for mutual cooperation. Moreover, mean encounter lengths of substantially less than 20 can give rise to increasing mean fitness and subsequently develop into mutual cooperation (i.e., mean population fitness of > 2.6).

The results of these experiments highlight a point that is important to the study of complex adaptive systems. Such systems are often parametrized by a very large number of interactive variables. For example, in the iterated Prisoner's Dilemma, these include (1) the length of the game, (2) the form of representation for a strategy, (3) the maximum memory length of a strategy, (4) the specific payoff function, including possible random effects, (5) the form of competition (e.g., round-robin vs. spatial), (6) the severity of selection pressure, and many more. Each of these variables offers a dimension in the state space of possible configurations for the simulation. It must be recognized that holding many of these variables constant (e.g., encounter length, payoff function, and so on) and varying only a few at a time may render important interactions invisible to the investigator. A stationary point generated by variations within some subset of variables may be overcome by including variation on another dimension (Conrad, 1983, 1993). Although adding degrees of behavioral freedom to a model makes it more difficult to analyze, if the study of models of complex adaptive systems is to be of general benefit in advancing an understanding of real-world phenomena, the interactions among system components cannot be ignored.

Acknowledgments

The author would like to thank P.J. Angeline, D. Ashlock, W. Atmar, K. A. DeJong, G. B. Fogel, L. J. Fogel, P. G. Harrald, and the anonymous referees for their helpful comments and suggestions. Thanks also to J. J. Grefenstette for his time and effort.

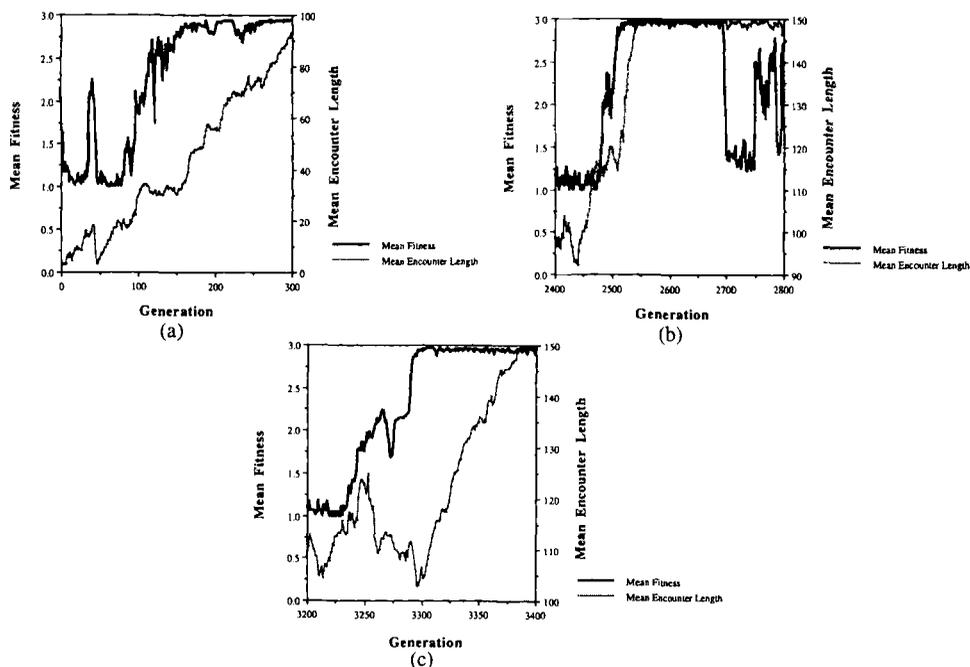


Figure 12. Three graphs indicating different relationships between the mean fitness of the population (read on the left-hand vertical axis) and the mean encounter length (read on the right-hand vertical axis) as a function of the number of generations. (a) Trial 1. The first rise from mutual defection to mutual cooperation begins around the 100th generation. The nearly linear increase in the mean encounter length begins before the increase in mean fitness. (b) Trial 2. The first rise from mutual defection to mutual cooperation begins around the 3250th generation. Although the mean encounter length again appears to increase just before the mean fitness starts to increase, it quickly declines, even while mean fitness continues to increase toward complete mutual cooperation. Further, the mean encounter length had reached values over 100 without generating any mutual cooperation within the first 3250 generations. (c) Trial 3. The first rise from mutual defection to mutual cooperation begins around the 2450th generation. Again, the mean encounter length increases sharply just before the increase in mean fitness and continues to increase along with it, as in Trial 1. However, the mean encounter length remains near the maximum bound of 150 despite the fact that the mean fitness rapidly declines around the 2675th generation.

References

- Angeline, P. J. & Pollack, J. B. (1994). Coevolving high-level representations. In C. G. Langton (Ed.), *Artificial Life III* (pp. 55–71). Reading, MA: Addison-Wesley.
- Axelrod, R. (1980a). Effective choice in the Prisoner's Dilemma. *Journal of Conflict Resolution*, 24, 3–25.
- Axelrod, R. (1980b). More effective choice in the Prisoner's Dilemma. *Journal of Conflict Resolution*, 24, 379–403.
- 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* (pp. 32–41). London: Pitman.
- Conrad, M. (1983). *Adaptability*. New York: Plenum Press.
- Conrad, M. (1993). Structuring adaptive surfaces for effective evolution. In D. B. Fogel & W. Atmar (Eds.), *Proceedings of the Second Annual Conference on Evolutionary Programming* (pp. 1–10). La Jolla, CA: Evolutionary Programming Society.
- Dugatkin, L. A. (1992). The evolution of the “con artist.” *Ethology and Sociobiology*, 13, 3–18.
- Fogel, D. B. (1991). The evolution of intelligent decision making in gaming. *Cybernetics and Systems*, 22, 223–236.
- Fogel, D. B. (1993). Evolving behaviors in the iterated Prisoner's Dilemma. *Evolutionary Computation*, 1, 77–97.
- Fogel, D. B. (1995). *Evolutionary computation: Toward a new philosophy of machine intelligence*. Piscataway, NJ: IEEE Press.
- Fogel, D. B., & Harrald, P. G. (1994). Evolving continuous behaviors in the iterated Prisoner's Dilemma. In A. V. Sebald & L. J. Fogel (Eds.), *Proceedings of the Third Annual Conference on Evolutionary Programming* (pp. 119–130). River Edge, NJ: World Scientific.
- Hillis, W. D. (1992). Co-evolving parasites improve simulated evolution as an optimization procedure. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II* (pp. 313–324). Reading, MA: Addison-Wesley.
- Kauffman, S. A., & Johnson, S. (1992). Co-evolution to the edge of chaos: coupled fitness landscapes, poised states, and co-evolutionary avalanches. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II* (pp. 325–370). Reading, MA: Addison-Wesley.
- Lindgren, K. (1992). Evolutionary phenomena in simple dynamics. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.) *Artificial Life II* (pp. 295–312). Reading, MA: Addison-Wesley.
- Miller, J. H. (1989). The coevolution of automata in the repeated Prisoner's Dilemma. Working Paper No. 89-003, Santa Fe Institute, July.
- Ray, T. S. (1992). An approach to the synthesis of life. In C. G. Langton, C. Taylor, J. D. Farmer & S. Rasmussen (Eds.) *Artificial Life II* (pp. 371–408). Reading, MA: Addison-Wesley.
- Rosen, C. D., & Belew, R. K. (1995). Methods for competitive co-evolution: Finding opponents worth beating. In L. J. Eshelman (Ed.), *Proceedings of the Sixth International Conference on Genetic Algorithms* (pp. 373–380). San Francisco, CA: Morgan Kaufmann.
- Stanley, E. A., Ashlock, D., & Tesfatsion, L. (1994). Iterated Prisoner's Dilemma with choice and refusal of partners. In C. G. Langton (Ed.), *Artificial Life III* (pp. 131–175). Reading, MA: Addison-Wesley.