National University of Ireland, Galway
Ollscoil na hÉireann, Gaillimh

DEPARTMENT OF INFORMATION TECHNOLOGY

technical report NUIG-IT-141002

Forgiveness, noise and memory in the Iterated Prisoner's Dilemma

C. O' Riordan (NUI, Galway)

Forgiveness, noise and memory in the Iterated Prisoner's Dilemma

Colm O' Riordan,
colmor@geminga.nuigalway.ie
Dept. of Information Technology,
National University of Ireland,
Galway, Ireland.

October 14, 2002

Abstract

This paper focuses on the iterated prisoner's dilemma(IPD). We outline results obtained in simulations of the IPD in both clean and noisy environments. We discuss attributes we believe to be of importance with respect to the fitness of strategies in both environments. Specifically we discuss notions of *forgiveness* (where strategies attempt to forgive strategies that defect in the game) and *memory* (whereby strategies are less reactive to immediate events but maintain a longer memory of past interactions).

1 Introduction

Multi-agent systems have been adopted and researched as a means to develop robust, distributed reactive and intelligent systems. Within the field, much research has been informed by previous and ongoing research in the fields of Computer Science, namely those of distributed artificial intelligence (DAI) and distributed problem solving (DPS), and also by research within the social sciences and economics.

We typically view a Multi-Agent System as comprising a set of intelligent agents[26][12] each capable of behaving in an autonomous manner to satisfy their own goals (which may be individual or joint goals involving other agents).

Much research has focused on the development of theories and languages to aid the development of agent-based systems. Work has also progressed in the development of tools and ideas to analyse the behaviour of agents and multi-agent systems.

An important strand of this research has been the field of game-theoretic and decision-theoretic analysis of agent-based systems[28][3][14]. The game-theoretic approach is typified by the abstraction of agent behaviour to a series of games, in which one attempts to capture the salient notions and features of agent behaviour. Analysis and exploration of ideas regarding agent behaviour can be effected by studying these games (through formal analysis of these games and through the simulation of strategies playing these games).

This paper concentrates on one such game, namely the well-known iterated prisoner's dilemma. The paper focuses on two aspects of strategies for playing the game—the willingness of strategies to *forgive* and the use of a *memory* of past interactions to guide future behaviour. We explore these features for strategies within both noise-free environments and noisy environments.

The motivation for further research in the iterated prisoner's dilemma, given that it has been well-studied in a variety of domains[2] are two-fold:

- the assumption of cooperative behaviour in multi-agent systems is not necessarily valid in domains where agents are truly autonomous—agents may be selfish and/or competitive. More research is necessary to allow the creation and deployment of multi-agents systems which permit agents to act in a non-cooperative manner and yet have means to reason or possibly constrain the emergent global system behaviour.
- the assumption of a noise-free environment is also not valid in certain domains. An *intended* cooperative gesture by one agent may be interpreted as a non-cooperative gesture by the receiving agent due to a number of potential reasons (for example, ambiguity in the message; conflicting goals of agents; differences in ontologies maintained and used by the agents; and 'cultural differences' of agents). A transmitted gesture may be lost or damaged in transmission which may result in the receiver mis-interpreting the message. These potential problems may result in a cooperative gesture being recognised by the receiver as a defective gesture and vice versa.

These issues have not been widely studied, particularly the latter problem with only a handful of researchers considering noise and errors in games such as the iterated prisoner's dilemma [6][17][23][24].

In this paper, a review of research in the iterated prisoner's dilemma is given together with a review of work in the IPD in noisy environments. Subsequent sections outline some recent results obtained in experiments dealing with the iterated prisoner's dilemma in both noise-free and noisy environments. The approach taken in this research has been to design strategies based on heuristics, to validate these strategies in fixed environments and to search the range of the features of these strategies using evolutionary computational search strategies. The paper concludes with some observations regarding the evolved strategies and their features which may be of use in the development and deployment of multi-agent systems.

2 Prisoner's Dilemma

In the prisoner's dilemma game, two players are both faced with a decision—to either cooperate(C) or defect(D). The decision is made by a player with no knowledge of the other player's choice. If both cooperate, they receive a specific punishment. If both defect they receive a larger punishment. However, if one defects, and one cooperates, the defecting strategy receives no punishment and the cooperator a punishment (the sucker's payoff). The game is often expressed in the canonical form in terms of pay-offs:

		`Player 1`	
		C	D
`Player 2`	C	(λ_1, λ_1)	(λ_2, λ_3)
	D	(λ_3, λ_2)	(λ_4, λ_4)

where the pairs of values represent the pay-offs (rewards) for players `Player 1` and `Player 2` respectively. The prisoner's dilemma is a much studied problem due to it's far-reaching applicability in many domains. In game theory, the prisoner's dilemma can be viewed as a two-person, non-zero-sum, non-cooperative and simultaneous game. In order to have a dilemma the following must hold: $\lambda_3 < \lambda_1 < \lambda_4 < \lambda_2$, where λ_2 is the sucker's payoff, λ_4 is the punishment for mutual defection, λ_1 is the reward for mutual cooperation and λ_3 is the temptation to defect. The constraint $2\lambda_1 > \lambda_2 + \lambda_3$ also holds.

The prisoner's dilemma and applications has been described in many domains including biology[10][11][20], economics[29] and politics[7].

3 Iterated Prisoner's Dilemma

In the iterated version 2 players play numerous games (the exact number unknown to either player). Each player adopts a strategy to determine whether to cooperate or defect at each of the moves in the iterated game.

3.1 Strategies

Strategies for the IPD may be:

periodic: strategies play C or D in a periodic manner. Common strategies: *ALL-C*, *ALL-D*, *(CD)**, *(DC)**, *(CCD)**, etc.

random: strategies that have some random behaviour. Totally random, or one of the other types (e.g. periodic) with a degree of randomness.

based on some history of moves: *tit-for-tat* (C initially, then D if opponent defects, C if opponent cooperates), *spiteful* (C initially, C as long as opponent cooperates, then D forever), probers (play some fixed string, example (DDC) and then decides to play *tit-for-tat* or *ALL-D* (to exploit non-retaliatory)), *soft-majo* (C initially, then cooperate if opponent is not defecting more than cooperating).

There are many variations on each of the above type of strategies.

3.2 Results

A computer tournament[1] was organised to pit strategies against each other in a round-robin manner in an attempt to identify successful strategies and their properties. The winning strategy was *tit-for-tat* (TFT); this strategy involved cooperating on first move and then mirroring opponents move on all subsequent moves.

The initial results and analysis showed that the following properties seemed necessary for success—niceness (cooperate first), retaliation, forgiveness and clarity.

In a second tournament[1], 15 of the top 16 strategies, were found to be nice. These results seem to indicate that cooperative strategies are useful if there is a high chance that the strategies will meet again.

Further analysis involved the development of a genetic algorithm to evolve successful strategies. The more successful strategies tended to be more complex than the traditional TFT and violated the fourth heuristic (that of clarity) proposed by Axelrod :"Don't be too clever"; these strategies are quite complex.

Beaufils et al[4] question the fourth property and develop a strategy *gradual*[1] which is far more complex than *tit-for-tat* and outperforms *tit-for-tat* in experiments. The *forgiving*[27] strategy also challenges the final property; *forgiving* is not clear or simple and has proven strong in environments similar to those used by Beaufils[4].

No best strategy exists; the success of a strategy depends on the other strategies present. For example, in a collection of strategies that defect continually (*ALL-D*) the best strategy to adopt is *ALL-D.* In a collection of strategies adopting a *tit-for-tat* strategy, an *ALL-D* strategy would not perform well.

Some ideas to promote cooperation in environments have been posited by Axelrod; these include genetic kinship, clustering of like strategies, recognition, maintaining closeness when recognition capabilities are limited or absent (e.g limpets in nature), increasing the chance of future interactions

---

[1] gradual performs like *tit-for-tat*, in that it cooperates on the first move. It retaliates upon defection. On the first defection it responds with a defection(D), followed by 2 cooperations(CC). Following the second defection, it responds with 2Ds, followed by 2Cs and so forth.

(certain social organisations, hierarchies in companies etc.), changing the pay-offs, and creating social norms where one is encouraged to cooperate.

4 Noisy environments

The majority of work in the iterated prisoner's dilemma has focused on the games in a noise-free environment, i.e. there is no danger of a signal being misinterpreted by the opponent or the signal being damaged in transit.

However, this assumption of a noise-free environment is not necessarily valid if one is trying to model real-world scenarios.

There are different means that can be used to introduce noise to the simulation:

- mis-implementation (when the player makes a mistake implementing its choice)
- mis-perception (when one player perceives incorrectly the other player's signal or choice)

Bendor[6] effects noise by introducing payoffs that are subject to error. Upon cooperation in the face of defection by an opponent, a person receives the payoff $\lambda_2 + e$, where e is random with expected value 0.

In [13], it is argued that "if mistakes are possible evolution may tend to weed out strategies that impose drastic penalties for deviations".

Kahn and Murnighan [17] find that in experiments dealing with the iterated prisoner's dilemma in noisy environments, cooperation is more likely when players are sure of each other's payoffs. Miller's experiments in genetic algorithms applied to the prisoner's dilemma result in the conclusion that cooperation is at its greatest when there is no noise in the system and that this cooperation decreases as the noise increases[23].

Hoffman[16] reports that results are sensitive to the extent to which players make mistakes (mis-implementation of mis-perception). In particular, cooperation is vulnerable to noise as it is supported by conditional strategies. For example, in a game between two TFTs, a single error would trigger a series of alternating defection. Axelrod [1] repeated his initial round-robin tournament with added 1% chance of players misunderstanding their opponent's move in any round. He found that TFT still came first despite some echoes of retaliation between cooperative strategies.

It can be shown that higher degrees of noise can be detrimental to TFTs performance. Given noise of p percent, we can show that TFT strategies playing against each other will spend 25% of their time in mutual cooperation, 25% in mutual defection and the remainder of time with one strategy cooperating and the other defecting.

A number of of authors confirm the negative effect of noise on TFT and find that more forgiveness promotes cooperation in noisy environments[6][24].

Other interesting results are also reported. These include 'pavlovian' strategies which are more likely to avoid spirals of defection that `tit-for-tat`[18] (also shown to perform well in[22]), the lowering of levels of cooperation in a society without the introduction of defecting strategies[23][21], the effect of highlighting differences between strategies that would coexist in noise-free environments[8].

5 Motivations for designing strategies for the IPD

The primary shortcomings of strategies like TFT in the noisy IPD is that their memory of past interactions is too short—TFT quickly reacts to any defection whether it is a genuine defection or one caused by a degree of noise. It seems plausible that a longer memory would permit strategies to score more points by not reacting to an opponents defection, but rather to the opponent's behaviour over a period time.

Secondly, in the IPD, strategies like TFT have a tendency to lock into a spiral of mutual defections which can prove to be detrimental to the overall fitness of the strategies.

Another less important factor is TFT's potential to be exploited by non-nice naive strategies.

By running some simulations and investigating the properties of the successful strategies we were able to add further weight to the hypotheses forwarded in the preceding paragraphs.

We argue that in both clean and noisy environments, the willingness to break spirals of mutual defection by incorporating a degree of forgiveness is a useful trait. Furthermore, we argue that as environments become noisy a longer memory is beneficial; i.e. one should base the reaction not just on the previous few moves but on a longer memory of the game.

In subsequent sections, we discuss some results in both clean and noisy environments.

6 IPD in a clean environment

6.1 Introduction

Our strategy, *forgiving*[2] attempts to take the following factors into account:

- Don't be exploited by periodic strategies
- Try to re-establish cooperation by forgiving

Note that the above modifications do not violate the first three recommendations (generally accepted) forwarded by Axelrod—never defect first, be retaliatory, be forgiving. There are cases when the exploitation of periodic strategies can damage performance: where a pattern is recognised as a periodic strategy and we adopt an `ALL-D` approach to avoid exploitation. This can quickly result in a spiral of mutual defections if the opponent is not really periodic (but appears to be).

The degree of forgiveness in our strategy is of length 2, i.e we play two consecutive Cs. The length of the spiral of mutual defections is set to 5 (i.e. once 5 pairs of defections are encountered an effort is made to re-establish cooperation). These figures were initially chosen rather arbitrarily; we merely wished the spiral to be of some reasonable length and the cooperative gesture to be sufficient to re-establish mutual cooperation.

In summary, our strategy is like tit-for-tat, with the following amendments—exploit periodic strategies and forgive when interactions are spiraling into an ongoing defection.

6.2 Results

The initial experiments carried out on the strategy were in a round-robin tournament with 37 other well-known strategies. The strategies chosen were those included as default in the simulation package created by Beaufils and Delahaye[3]. This very useful package allows experiments in both a round-robin and an evolutionary setting. A set of well-known strategies are included in the package and the addition of new strategies is facilitated. We added our strategies to the pool of strategies provided.

In the evolutionary simulation, each successive generation contains strategies with frequency proportional to their score in the current generation. The performance of the strategies over a number of generations was plotted (Figure 1).

---

[2] It is important to note, to avoid confusion, that Axelrod's notion of contrition is fundamentally different to the notion of *forgiveness* [27] discussed in later sections. Axelrod's notion of contrition effectively amounts to modifying TFT such that it "avoids responding to the other player's defection after its own unintended defection"

[3] Available at http://www.lifl.fr/IPD/ipd.frame.html

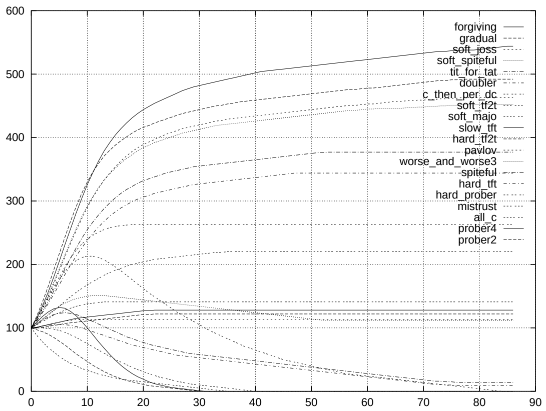

Figure 1: Evolutionary Setting

Following these initial results we wished to investigate which of the two aspects of the *forgiving* strategy accounted for its good performance (its exploitation of periodic strategies or its ability to re-establish cooperation). The performance of the strategy in the environmental setting indicates that its exploitation of periodic strategies, while useful, is not necessary for its success as the strategies (periodic) upon which it preys die off at a relatively early stage (e.g. *per-cd*, *per-ddc*).

To provide empirical evidence, we also include two variations—*forgiving-1* which does not exploit periodic strategies but attempts to re-establish cooperation and *forgiving-2* which attempts to exploit periodic strategies only. The graph in Figure 2 shows their performance. As can be seen, the two strategies that attempt to forgive and re-establish cooperation do well in the evolutionary setting.

6.2.1 Evolutionary Search for *forgiving* strategies

The previous sections provide some results to justify the incorporation of forgiveness into strategies. To explore more fully the effect of degrees of forgiveness, an evolutionary computation method was adopted.

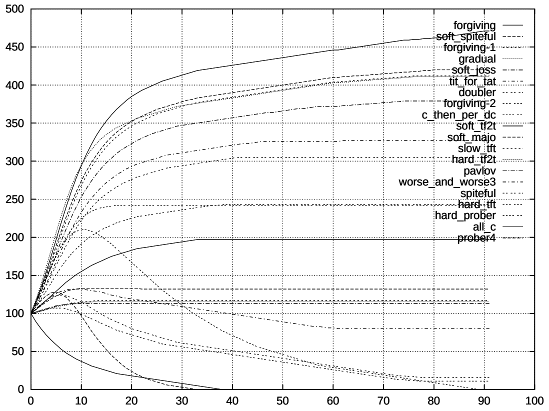

Figure 2: Evolution of different versions of forgiving

We use a genetic algorithm to breed successful strategies to play against a well-known collection of strategies (the same used in the previous strategies).

Evolutionary computation approaches have been previously used by other researchers to explore the range of strategies. These include Nowak and Sigmund[25], Linster[19], Beaufils[5], Harrald and Fogel[15], Cohen et al.[9].

In our initial experiments, we encode, in each chromosome, several aspects of a strategy's behaviour. These include:

- behaviour on the first move.
- behaviour following a defection by opponent.
- behaviour following a cooperation by opponent.
- number of mutual defections allowed before changing to *forgiving* behaviour.
- number of successive cooperative gestures to make in order to *forgive.*
- whether a strategy will use calculations based on a longer memory to choose the next move.
- a threshold (number of previous cooperations less the number of previous defections) under which strategies will not cooperate.

On all trials the first three genes converge extremely quickly and confirm findings by many others—cooperate on first move, cooperate following a cooperation and defect following a defection.

The fifth gene does not converge to zero in any of the experiments, indicating that *forgiving* behaviour is a useful feature for any strategy to maintain. The degree of *forgiveness* and the length of the spiral have varied in the trials but have consistently been much larger than those in our designed strategy.

The above experiments allowed us to explore a range of attributes of forgiving strategies.

6.3 Co-evolution of strategies

The experiments involved evolving a set of strategies where the initial set of strategies created randomly and then evolved in the environment of creatures playing the iterated prisoner's dilemma. This was to remove any bias that may exist in our experiments by the choice of strategies with whom our evolved strategies were playing.

In these experiments, the noise level was set to zero, so we were dealing with a noise free environment. In both of these experiments, the same settings were used apart from the mutation rate which was changed to a higher value in the second experiment.

In both experiments we see interesting behaviour. In the first (Figure 3), we see that rather surprisingly, at first glance, the strategies converge such that strategies initially defect on their first move. This seems unusual given that the selection of cooperation on the first move in evolutionary settings is much reported in experiments similar to our own. We also see that the evolved strategies retaliate and do not immediately forgive following a defection by their opponent.

However, we see that *forgiveness* quickly becomes the norm; strategies choose to forgive strategies with whom they are locked in a mutual spiral of defection. We believe the presence of this behaviour in the population reduces the evolutionary advantage conferred on those that initially cooperate; the danger of mutual defection is not as costly as the strategies are likely to forgive. The notion of maintaining a long memory has not been selected for in the evolutionary setting.

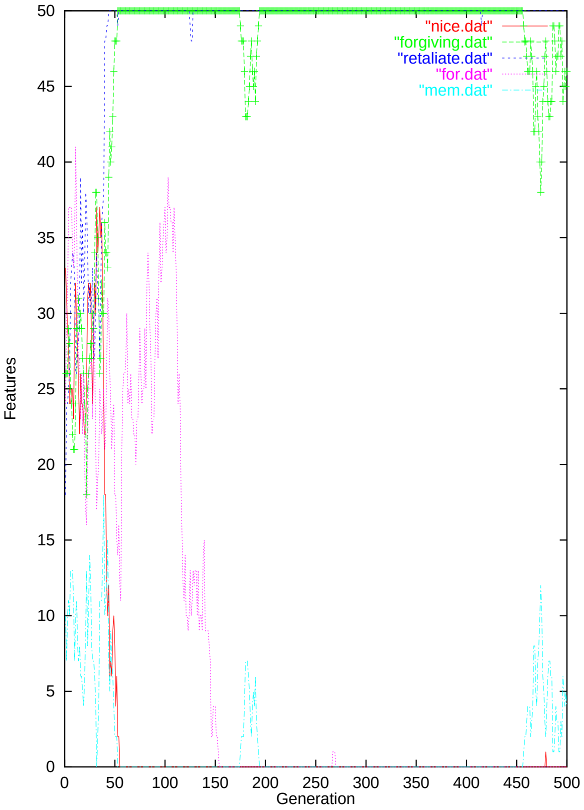

Figure 3: Population Size = 50. No of generations = 500. No noise. Mutation: 0.001%, Crossover 0.80%

In another run, with a higher mutation rate in the genetic algorithm, we see another scenario evolving. In this case (Figure 5), cooperation in the first move becomes a dominant feature in the population. Immediate retaliation and forgiveness are also selected. We also see a favouring of long-term forgiveness as an evolved trait. This seems

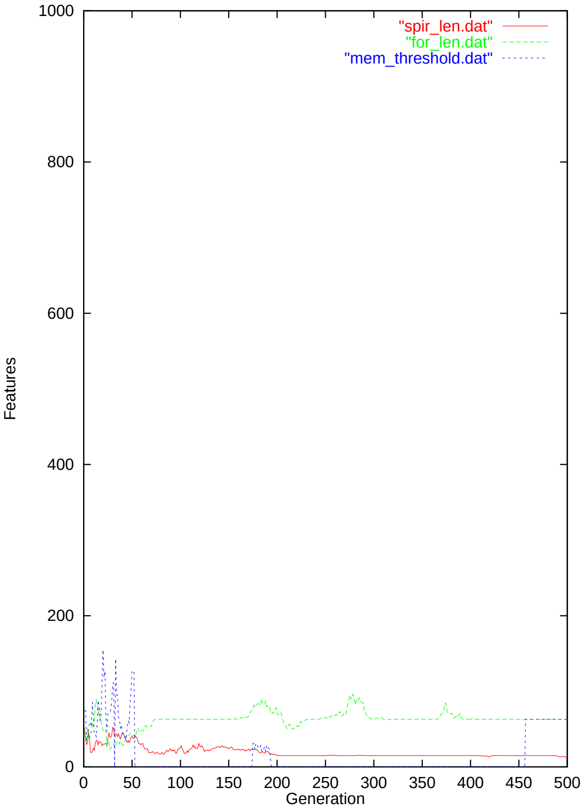

Figure 4: Forgiving, Spirals, and memory; no noise; mutation 0.01%

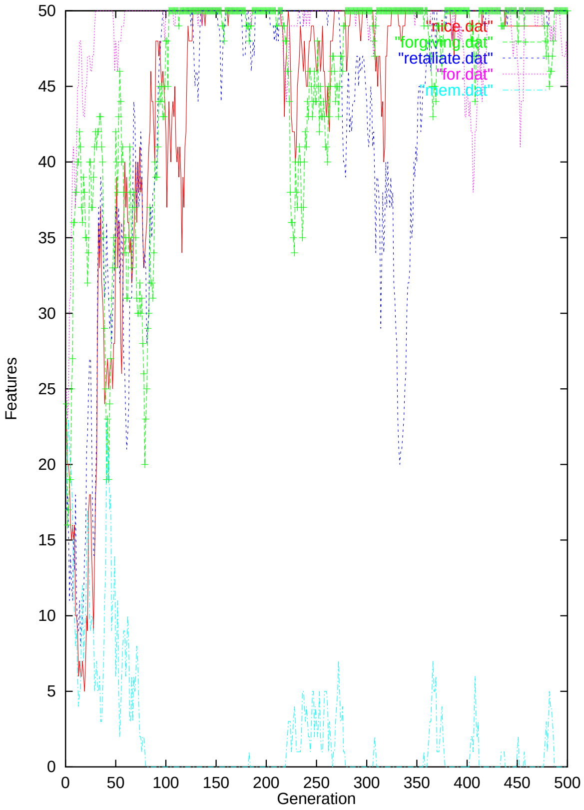

Figure 5: Population 50; Length : 1000; Noise: 0%; Mutation: 0.01%, Crossover 0.80%

to indicate that irrespective of whether or not the population is initially nice or not, the willingness to forgive upon encountering mutual spirals of defection is a desirable trait.

Figures 4 and 6 show graphically, plotted against time, the average spiral length, the average length of forgiveness and the threshold used by agents adopting a longer memory.

The lengths are quite small given the maximum possible but yet are longer than in the original designed strategies. In the second of these graphs, the levels for spirals and forgiveness are quite small, which is reasonable given that there will be less spirals given the large proportion of strategies partaking in cooperative behaviour.

6.4 IPD in a noisy environment

Given the results and evidence obtained from experiments in the noise-free environment, we wished to investigate the effects that introducing noise would have on strategies and their features.

By ranking the strategies according to their proportion following 100 generations of evolution in a noisy environment, we see different strategies (e.g. doubler, soft-joss) attaining higher degrees of fitness,

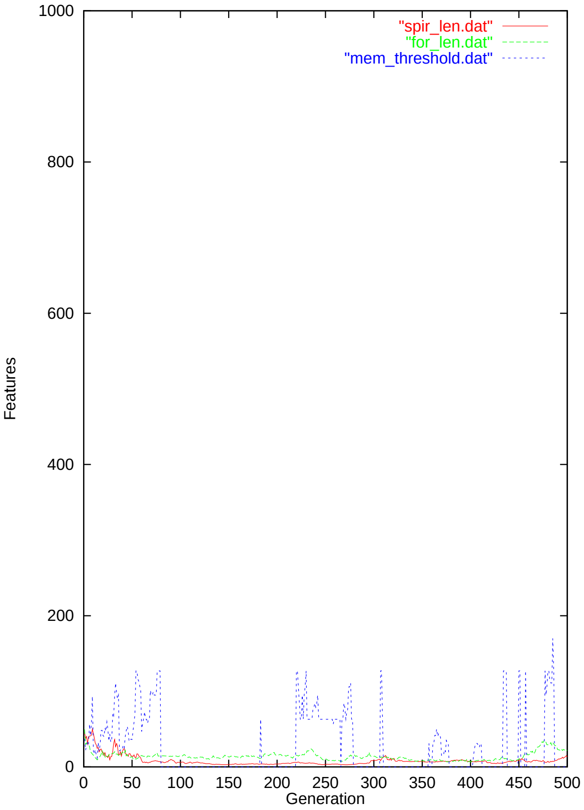

Figure 6: Forgiving, Spirals, and memory

with strategies such as tft and *forgiving* attaining lower degrees of fitness than witnessed in earlier experiments in a clean, noise-free environment.

1 percent noise

```
------  Generation score in 101 ------
1          doubler =   1447
2        soft_joss =   1129
3    soft_spiteful =    596
4        soft_tf2t =    504
5      hard_prober =    158
6      forgiving-1 =     96
7        forgiving =     44
8      tit_for_tat =     15
9        hard_tf2t =      6
10         slow_tft =      0 stopped in 76
```

10 percent noise

```
------  Generation score in 101 ------
1          doubler =   1189
2        soft_joss =    865
3    soft_spiteful =    776
4        soft_tf2t =    407
5        soft_majo =    311
6      forgiving-1 =    234
7    prob_c_4_on_5 =    111
8        forgiving =     58
9      tit_for_tat =     21
10        hard_tf2t =     18
```

On first analysis of the results we see that the strategies that do well are those that base their responses on calculations over the entire history of the games, i.e., a longer memory seems advantageous.

We repeated the co-evolutionary experiments again but introduced a small degree of noise to the simulations. A genetic algorithm is again used with the same chromosome layout as in the earlier experiments.

In the first experiment with a noise at a rate of 1% we see the population fluctuates a lot with respect to the different features. We plot only willingness to forgive, memory and niceness in the graph to avoid clutter in the graph. As can be seen from the graph, the society lacks stability and fails to converge on any fixed value for any of the features.

7 Conclusion

The iterated prisoner's dilemma is an oft-studied game in many domains. This paper examined some features of strategies playing the game. Initially we explored the concept of forgiveness (the ability of a strategy to cooperate following a spiral of defections). We found that in a fixed environment a designed forgiving strategy flourished, in an evolutionary setting forgiveness was also selected for and finally in a co-evolutionary setting we showed that a selective

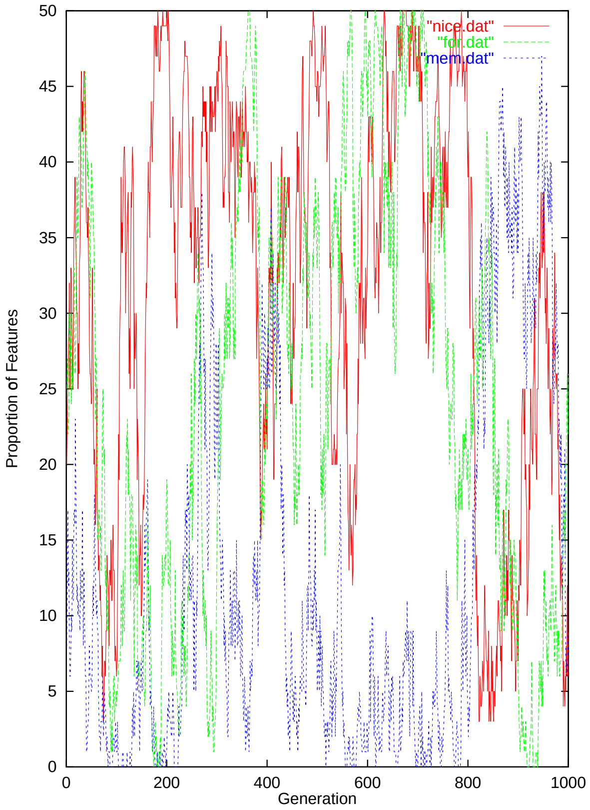

Figure 7: Population 50; Length : 1000; Noise: 1%; Mutation: 0.01%, Crossover 0.80%

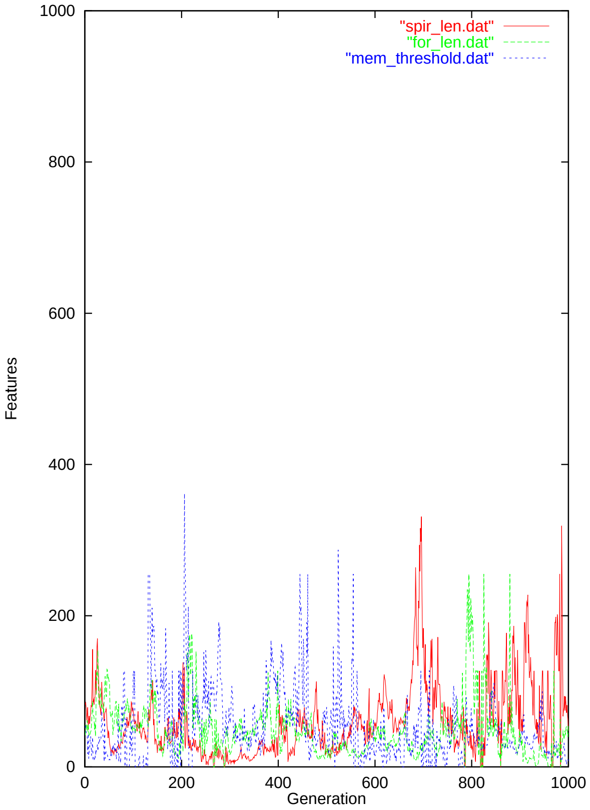

Figure 8: Forgiving, Spirals, and memory

advantage was conferred upon forgiving strategies.

We also wished to investigate the importance of ‘memory’, i.e. the ability of a strategy to maintain statistics of the entire game and not just the immediate past. We saw that in our evolutionary and co-evolutionary settings, ‘forgiveness’ was selected over ‘memory’.

We also investigated the effect of introducing noise into the system. The introduction of a low level of noise into the system has an immediate effect on well-known strategies and we saw that their willingness to react to defections causes a decrease in their fitness. In the evolutionary experiments we see that forgiveness, although useful, does not result in a stable society. The use of a longer memory tends to be more useful in these experiments but similarly does not produce a stable society.

References

[1] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

[2] R. Axelrod and L. D'Ambrosio. An annotated bibliography on the evolution of cooperation. *http://www.ipps.lsa.umich.edu/ipps/papers/coop/Evol_of_Coop_Bibliography.txt*, Oct. 1994.

[3] A. Bazzan and R. Bordini. Evolving agents with moral sentiments in an iterated prisoner's dilemma

exercise. *2nd Workshop on Game Theoretic and Decision Theoretic Agents*, 2000.

[4] B. Beaufils, J. P. Delahaye, and P. Mathieu. Our meeting with gradual, A good strategy for the iterated prisoner's dilemma. In C. Langton, editor, *Artificial Life, Proc. 5th International Workshop on the Synthesis and Simulation of Living Systems*, volume 5, 1996. Artificial Life 5, Nara, Japan, May 16-18.

[5] B. Beaufils, P. Mathieu, and J.-P. Delahaye. Complete classes of strategies for the classical iterated prisoner's dilemma. In *Proceedings of Evolutionnary Programming VII, LNCS vol 1447*, pages 33–41. Springer-Verlag, 1998.

[6] J. Bendor, R. Kramer, and S. Stout. When in doubt cooperation in a noisy prisoner's dilemma. *Journal of Conflict Resolution*, 35(4):691–719, 1991.

[7] J. Bendor and D. Mookherjee. Institutional structure and the logic of ongoing collective action. *American Political Science Review*, (81):129–154, 1987.

[8] R. Boyd. Mistakes allow evolutionary stability in the repeated prisoner's dilemma. *Journal of Theoretical Biology*, 136:47–56, 1989.

[9] M. D. Cohen, R. L. Riolo, and R. Axelrod. The emergence of social organization in the prisoner's dilemma: How context-preservation and other factors promote cooperation. Technical report, SantaFe Institute, 1999.

[10] L. Dugatkin. Do guppies play tit for tat during predator inspection visits? *Behavioral Ecology and Sociobiology*, 2(3):395–399, 1988.

[11] L. Dugatkin. N-person games and the evolution of co-operation: A model based on predator inspection in fish. *Journal of Theoretical Biology*, (142):123–135, 1990.

[12] M. Fisher and M. Wooldrige. On the formal specification and verification of multi-agent systems. *Intelligent Agents Workshop*, 1995.

[13] D. Fundunberg and E. Maskin. Evolution and cooperation in noisy repeated games. *The American Economic Review*, 80(2):274–179, 1990.

[14] P. Gmytrasiewicz and C. Lisetti. Using decision theory to formalize emotions in multi-agent systems, 2000.

[15] P. Harrald and D. Fogel. Evolving continuous behaviours in the iterated prisoner's dilemma. *Biosystems*, 37:135–145, 1996.

[16] R. Hoffman. Twenty years on: The evolution of cooperation revisited. *Journal of Artificial Societies and Simulation*, 3(2), 2000.

[17] L. Kahn and J. Murnighan. Conjecture, uncertainty, and cooperation in prisoner's dilemma games: Some experimental evidence. *Journal of Economic Behaviour & Organisation*, 22:91–117, 1993.

[18] D. Kraines and V. Kraines. Learning to cooperate with pavlov an adaptive strategy for the iterated prisoner's dilemma with noise. *Theory and Decision*, 35:107–150, 1993.

[19] B. Linster. Evolutionary stability in the infinitely repeated prisoner's dilemma played by two-state moore machines. *Southern Economic Journal*, (58):880–903, 1992.

[20] M. Lombardo. Tree swallows and tit for tat : Response to koenig. *Ethology and Sociobiology*, (11):521–528, 1990.

[21] B. Lomborg. Nucleus and shield: The evolution of social structure in the interated prisoner's dilemma. *American Sociological Review*, (61):278–307, 1996.

[22] L. Marinoff. Maximizing expected utilities in the prisoner's dilemma. *Journal of Conflict Resolution*, 36(1):183–216, 1992.

[23] J. Miller. The coevolution of automata in the repeated prisoner's dilemma. *Journal of Economic Behaviour & Organisation*, 29:87–112, 1996.

[24] U. Mueller. Optimal retaliation for optimal cooperation. *Journal of Conflict Resolution*, 31(4):692–724, 1988.

[25] M. Nowak and K. Sigmund. The evolution of stochastic strategies in the prisoner's dilemma. *Acta Applicandae Mathematicae*, (20):247–265, 1990.

[26] H. S. Nwana. Software agents : An overview. *Knowledge Engineering Review*, 11(3):1–40, 1996.

[27] C. O' Riordan. A *forgiving* strategy for the iterated prisoner's dilemma. *Journal of Artificial Societies and Social Simulation*, 3(4), 2000.

[28] S. Sen. Reciprocity: A foundational principle for promoting cooperative b ehavior among self-interested agents. In V. Lesser, editor, *Proceedings of the First International Conference on Multi-Age nt Systems*. MIT Press, 1995.

[29] E. Sober. Stable cooperation in iterated prisoner's dilemmas. *Economics and Philosophy*, 8:127–139, 1992.