# Population Based Incremental Learning with Guided MutationVersus Genetic Algorithms: Iterated Prisoners Dilemma

**Timothy Gosling**

Dept. of Computer Science
University Of Essex
C04 3SQ, England
Email: tdbgos@essex.ac.uk

**Nanlin Jin**

Dept. of Computer Science
University Of Essex
C04 3SQ, England
Email: njin@essex.ac.uk

**Edward Tsang**

Dept. of Computer Science
University Of Essex
C04 3SQ, England
Email: edward@essex.ac.uk

**Abstract-** Axelrod's original experiments for evolving IPD player strategies involved the use of a basic GA. In this paper we examine how well a simple GA performs against the more recent Population Based Incremental Learning system under similar conditions. We find that GA performs slightly better than standard PBIL under most conditions. This differnce in performance can be mitigated and reversed through the use of a 'guided' mutation operator.

## I. INTRODUCTION

Experiments in to evolving strategies to play Iterated Prisoners Dilemma (IPD) were initially carried by Axelrod in 1987 [1]. These experiments found that from an evolutionary stand point Tit-For-Tat was a dominant strategy. In the years since these early results others have attempted to evolve IPD strategies and made claims about the dominance of various other approaches to playing the game [2][3][4]. This paper is not directly concerned with the arguments for or against various strategies but is rather concerned with how the more recent statistical approaches to evolutionary computation compare with traditional GA approaches in evolving those strategies.

To this end a series of experiments have been run comparing the effectiveness of IPD strategies evolved independently by straightforward GA and PBIL implementations. Where PBIL has proven less effective a guided mutation operator is shown to improve performance.

The paper introduces the strategy representation scheme used, the idea behind PBIL and the guided mutation operator, the way in which PBIL and GA were compared and the results of that comparison.

## II. REPRESENTATION

IPD strategies may be represented in a number of different ways partially dependent upon the particular variation of IPD being investigated [5][6]. For the GA/PBIL comparison being undertaken the original Axelrod representation for a three-round memory based game was used [1].

Under this system a 0 represents co-operation and 1 represents defect. A player's memory contains information on the last three rounds of play each of these rounds being represented by a bit pair. The bit pair consists of a record of both the player's move and its opponent's move in one of the previous rounds. Organising the pairs in time order provides a six-bit string that can be interpreted as a number between 0 and 63. Assuming a simple cooperate or defect reaction to any memory there are therefore a maximum of 64 possibly responses to a 3-round memory and the IPD strategy can thus be represented as a 64 bit string, each position in the string providing the response co-operate or defect to a specific memory. The below diagram illustrates how the player memory and strategy work together to produce a players move in the current round:
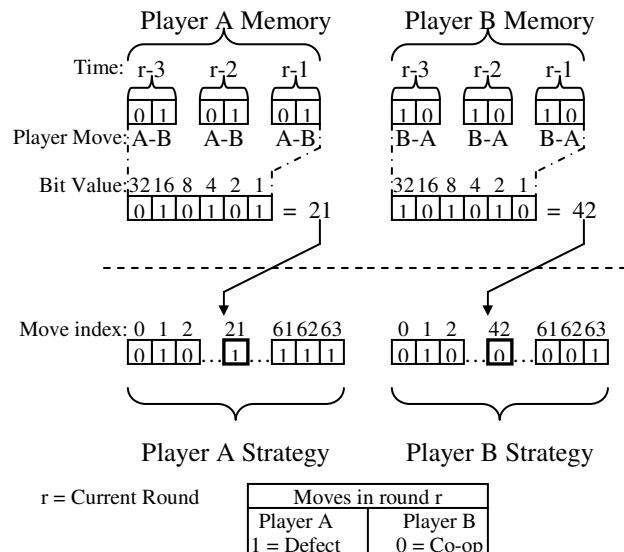


Figure 1. Player Memory and Strategy Representation

Since at the beginning of a game players start with no memory of previous rounds additional information is required as part of the representation. To this end an additional 6 bits is used to provide the player with an initial

starting memory and so its first index into the strategy. As play progresses the starting memory is eventually forgotten and play continues purely based upon true memories of the current interaction. The starting information could therefore be considered as a predisposition of a player towards its opponent. The full IPD representation is shown below:

Initial Memory | Strategy

5 4 3 2 1 0 | 0 1 2 3 4 5 ... 58 59 60 61 62 63
0 1 0 1 0 1 | 0 1 0 1 0 1 ... 0 1 0 1 0 1

Figure 2. Complete IPD Representation

## III. INTRODUCTION TO PBIL

Population based incremental learning (PBIL [7]) is a statistical approach to evolutionary computation that combines elements of GAs and Reinforcement Learning.

Under a simple PBIL scenario the basic representation of a solution can be the same as in a GA but instead of storing each possibility explicitly the population is replaced by a probability distribution. To elaborate further. If we consider a single member of a GA population it consists of a chromosome with a number of alleles. Each allele often represents some single variable in the solution and may take on a number of possible values, for the given population member though the value of each alleles is fixed. Within the population of solutions values for particular alleles exist with differing frequencies, if each population member was identical then only one value for each alleles would exist in the population as a whole, at the other extreme a wide variety of values would exist for each alleles within the population with little variation in frequency among them. PBIL essentially represents these frequencies directly and dispenses with the population itself. Thus under PBIL, each value of each alleles has a frequency (or probability) of existing within a hypothetical population associated with itself - the probability of each value within an allele must add up to 1. To generate a real solution string it is possible to select allele values probabilistically from the PBILs probability distribution. A diagram illustrating the difference between GA and PBIL representation can be found towards the end of this section (Figure 3).

To update a GAs population, population members are first evaluated and then recombine in some way to generate a new population. Members with a higher fitness have a greater probability of either finding their way in to the new population or helping generate new population members. Mutation is usually used to help increase diversity and reintroduce information that may have been lost at an earlier stage.

PBIL updates in a rather different manner. What needs to be updated is the probability distribution rather than a fixed population. The simplest way to perform the update is to find a good candidate solution and then increase the probability of each of the values of its alleles in the distribution (positive learning). The reverse can be done with a bad candidate solution with probabilities of values being reduced (negative learning). The rules for updating the probability of values can be quite simple and are usually tied to a learning rate (LR). The learning rate determines by how much the probability of a value under a given allele should increase and thus by how much the remaining value probabilities should be reduced. Fixed or variable learning rates can be used; if the LR is variable it may be tied to the relative fitness of the candidate solution being used to update the distribution.

Since a PBIL system will often have no real population of solutions to draw candidates from a temporary pool of solutions maybe generated from the distribution. The solutions in this pool can then be evaluated and the best and worst used to update the distribution.
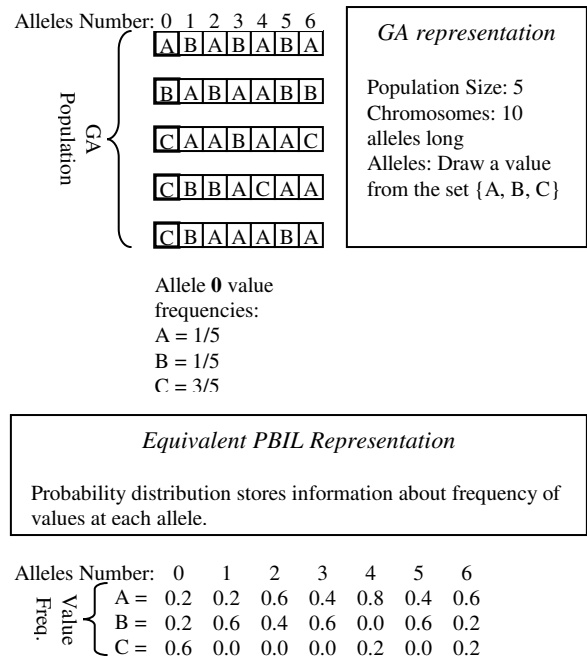
Alleles Number: 0 1 2 3 4 5 6

A B A B A B A
B A B A A B B
C A A B A A C
C B B A C A A
C B A A A B A

GA Population

*GA representation*

Population Size: 5
Chromosomes: 10 alleles long
Alleles: Draw a value from the set {A, B, C}

Allele **0** value frequencies:
A = 1/5
B = 1/5
C = 3/5

*Equivalent PBIL Representation*

Probability distribution stores information about frequency of values at each allele.

| Alleles Number: | 0 | 1 | 2 | 3 | 4 | 5 | 6 |
|---|---|---|---|---|---|---|---|
| A = | 0.2 | 0.2 | 0.6 | 0.4 | 0.8 | 0.4 | 0.6 |
| B = | 0.2 | 0.6 | 0.4 | 0.6 | 0.0 | 0.6 | 0.2 |
| C = | 0.6 | 0.0 | 0.0 | 0.0 | 0.2 | 0.0 | 0.2 |

Value Freq.

Figure 3.Differences between GA and PBIL representation

Mutation is often used with PBIL to help increase the search space much as with GA. Various schemes to implement mutation exist however two common approaches are either to vary the value frequencies by some amount with low probability or, alternatively, apply mutation with a low probability to generated population members before they are evaluated. Recently the guided mutation operator has been

proposed ([10]) and a variation of this is used here. The guided mutation operator works in a similar manner to the second mutation mechanism discussed. A generated population member may have some of its values selected at random for adjustment. In the case of normal mutation the adjustment itself is also random. Guided mutation uses a different approach, drawing upon the last best solution. In this approach the last best solution from the previous generation is remembered (this would often be the case anyhow) and any time a mutation occurs the mutated value is adjusted to reflect the value of the last best solution. In this way the operator acts to constrain variation to solutions that have been known to be at least reasonably effective in the past.

While the above explanation of PBILs operation is sufficient to explain how alleles with a discrete set of values or symbols may be represented it does not explain how continuous ranges maybe dealt with. Since continuous ranges are not required in the formation of the IPD strategies used here, no explanation of how this is accomplished will be provided.

## IV. PBIL SPECIFICS

The PBIL implementation used for running the IPD experiments represents each allele as a cell with a specific numeric range and number of symbols. For the purposes of IPD each cell is ranged 0 to 1 with two symbols, i.e. 0 and 1. At the beginning of an experiment each cell was set such that 0 and 1 had an equal probability.

The positive reinforcement rule used in the experiments was simple (taken from [7]). The mechanism used is described below:

$$prob_c = prob_c \times (1.0 - LR) + (cand_c - LR)$$

$prob_c$ - Probability of a 1 in cell c of the distribution

$LR$ - The learning rate (between 0 and 1)

$cand_c$ - Bit value at position c in the candidate solution

No negative learning was used in the experiments. Learning rates were varied between experiments but were constant within a give experiment. The range of learning rates used was between 10% and 0.5%.

Two different mutation operators were used within the experiments. In the first instance it was applied by changing generated population members with a probability of 0.7% per allele. When changed a value would simply flip from 1 to 0 or vice versa. In the second instance guided mutation was used, the last best solution being recalled from the previous (non test pool) population. A guided mutation rate of 1% was used. The last best solution was taken as the last population member used to positively reinforce the probability distribution. Many experiments were also run with no mutation operator in effect.

A generation in the sense of PBIL consists of the creation of a population, evaluation of that population and an update of the distribution by the fittest population member. In all the experiments described here PBIL was used with a greater number of generations than the GA for reasons described later.

To perform an update of the probability distribution two sets of IPD strategies were generated. The first 'update' or 'population' set was relatively small and used for updating the distribution; the second 'test' set was used purely for evaluating the first set. An example would be a population set of 10 coupled with a test set of 99. In this case each of the 10 population members would be tested against each of the 99 testing members to find its fitness. The population member with the highest mean score against each of its 99 opponents would be used to update the distribution.

This mechanism is necessary to even out the disparity in information use between the GA and PBIL when updating. While the GA (in a sense) makes use of its entire population to create a new population the PBIL system only uses one population member to update the distribution. Different ways of resolving this disparity might be used but the one above was selected for its ease of understanding. By selecting a smaller population size played against a larger testing pool the quality of individual evaluations may be maintained. By repeating the process for a larger number of generations than the GA a fairer use of information by both is maintained. In each case the total number of games played and solution evaluations must be maintained or bias will be introduced. To help prevent bias the following must hold:

$$GA_{evals} = GApop \times (GApop - 1) \times GAgens$$
$$PBIL_{evals} = PBILpop \times PBILtestpop \times PBILgens$$
$$GA_{evals} = PBIL_{evals}$$

GAevals/PBILevals - Total number of solution evaluations

GApop - The GA population size

GAgens - The number of generations the GA runs for

PBILpop - The population size of the PBIL system

PBILtestpop - The testing pool size of the PBIL system

PBILgens - The number of generation the PBIL system runs

With the above in mind a population size 100 GA running for a 100 generations could be played against a PBIL system using a population size of 10 and a testing pool of 99 running for 1000 generations. If however the GAs population was reduced a corresponding reduction in the

population size, test pool size or number of generations would be required by the PBIL system. Some combinations evidently will be more effective than others and in the course of these experiments different variations were considered.

For the purposes of comparison with the GA ten IPD strategies are generated at various points through out the experimental run (the GA records corresponding data after every generation). The ten strategies are generated in exactly the same way as for evaluation or testing. The interval between points at which comparison strategies are recorded can be determined by:

$$ComparisonOutputInterval = PBILgens / GAgens$$

The experimental random number generator seeds were select from the system clock, most of the PBIL experiments were conducted using the random number generator found in [8], the remaining experiments used the standard Java JDK random number generator.

## V. GA SPECIFICS

Individual strings (strategies) in the initial population are generated randomly with 50% possibility of choosing "Defection" and 50% possibility of choosing "Cooperation" at every bit of every 70-bits string.

Performance (fitness) of a string is evaluated by the average score that it earns from playing Iterated Prisoners' Dilemma with every other string in the same population.

Like natural selection, individuals having higher fitness are selected with higher probability. First of all, the fittest string is ensured to be selected as a parent, which is called "Elitism". Each of rest parents is chosen using the "Roulette-Wheel Algorithm". A random number $r \in [0, f_1 + f_2 + \ldots + f_{popSize})$ is created, then the string i whose fitness notated $f_i$ is selected, where $f_1 + f_2 + \ldots + f_{i-1} \leq r < f_1 + f_2 + \ldots + f_i$ .[9]

Strings are selected pair-wise and undergo one-point crossover, exchanging portions of strings of each other. Newly created intermediate strings mutate with very low rate (0.7%) by randomly alternating one bit of "cooperation" to "defection", or vice verse.

The offspring of the parent strings go on to form a completely new population for the next generation.

Strings used for comparison with PBIL are the first 10 strings chosen of every generation by Roulette-Wheel selection without Elitism.

## VI. GAMES

Part of the evaluation process of an IPD strategy involves playing against other IPD strategies. To this end each game used for evaluation consisted of 150 moves being played and the standard score grid below being used:



Figure 4. Standard Pay-off table

The same game parameters were used for the comparison discussed below.

## VII. GA AND PBIL COMPARISON

Providing a comparison between GA and PBIL systems in a way that provides neither with an advantage is difficult, however as much bias as possible has been removed. The following describes how the comparison was eventually realized.

To provide a comparison between the GA and PBIL systems for evolving IPD strategies, both systems were run independently and their resulting strategies tested against one another.

In running the comparison between GA and PBIL it was important to not provide a significant advantage to either, it was also critical that the comparison mechanism itself not be unfair or subject to too much uncertainty.

To counter the first problem the PBIL and GA systems were run for a differing number of generations with differing population sizes, the total number of evaluations and individual evaluation quality was maintained however. See 'PBIL Specifics' above for more details. To sensibly compare the GA and PBIL strategies at comparable intervals in the runs, ten strategies were generated by the PBIL and recorded for comparison and ten strategies were selected probabilistically and recorded for comparison by the GA. The recording process began at initialization and was performed at regular intervals up until the end of the run in each system. In the experiments reported here GA runs

lasted 100 or 300 generations, PBIL runs lasted between 1000 and 6000 generations.

To provide a comparison between a single GA and PBIL experiment for a given time, each of the ten strategies from each was played against all the strategies from the opposition. The results of these games, the mean scores and standard deviations were recorded. The mean strategy score for the GA and PBIL at the specified time can then be found easily. A single comparison such as this, results in 100 games being played. When comparing GA and PBIL runs in total this process is repeated for all compatible comparison points in both the PBIL and GA systems (either 100 or 300 points).

To improve the validity of the results each PBIL and GA experiment was repeated ten times (unless otherwise stated). Each of the PBIL and GA experiments could then be compared to one another and the results averaged. This results in 100 comparisons being done for a single time instance and so 10000 games being played between 200 strategies.

## VIII. EXPERIMENTAL RESULTS

While a large number of PBIL configurations and somewhat smaller number of GA configurations were tried, the most effective GA and PBIL configurations used are shown below:

| Type | GA | PBIL | PBIL+GM |
|------|-----|------|---------|
| Population Size | 100 | 5 | 5 |
| Learning Pool | NA | 99 | 99 |
| Mutation Rate | 0.007 | 0 | 0.9 |
| Learning Rate | NA | 0.025 | 0.025 |
| Generations | 300 | 6000 | 6000 |
| Data Points | 300 | 300 | 300 |

Table 1: Comparison of most effective GA and PBIL configurations

The first two of these configurations form the basis of the comparison at the beginning of this section. The third configuration, PBIL with Guided Mutation, is shown later to be more effective than the basic PBIL configuration.

The graphs, shown below, use 'relative mean fitness' to demonstrate the relative effectiveness of algorithms being compared. This is determined by the mean fitness of all of player A's strategies divided by the mean fitness of all of player B's strategies. Player A and B are determined by what the graph is attempting to demonstrate. When comparing various PBIL configurations against GA, player A will always be the best GA configuration while player B is each of the tested PBIL configurations. As a result any score above 1.0 shows greater effectiveness for player A

strategies and anything below 1.0 shows greater effectiveness for player B strategies.

The following diagram shows the relative performance of different PBIL configurations against GA. In this case the GA parameters are identical to those above and the PBIL parameters only vary in terms of the learning rate (Figure 5), PBIL uses no mutation in this instance:



Figure 5: Best GA v Population 5, no mutation PBIL with varying learning rates. PBIL initial fares well or poorly depending on LR but ultimately does worse than GA.

If standard mutation is applied to PBIL its performance is generally reduced. At low learning rates this effect is minimal but at higher learning rates the negative effect becomes more evident The diagrams below show the reduction in effectiveness caused by standard mutation at different learning rates (Figure 6, 7):
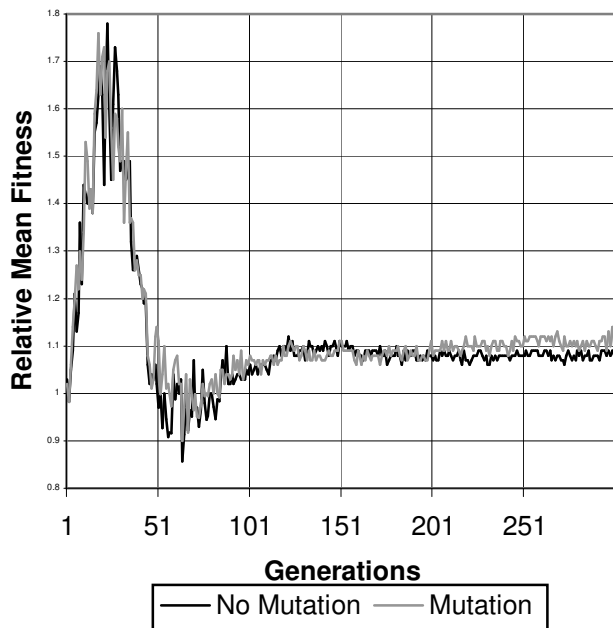
Figure 6: PBIL performance against GA with and without standard mutation a at low learning rate (0.005). Generation 1- 51, GA does very well compared to PBIL. Generation 51-101, PBIL recovers somewhat. Generation 101 onwards, PBIL performance stabilizes but is less effective than GA.



Figure 7: PBIL performance against GA with and without standard mutation at high learning rate (0.05). Generations 1-51, PBIL initial performs well but rapidly looses out to GA. Generation 51 onwards, PBIL stabalizes doing worse than PBIL. Mutated configuration PBIL does worse than non mutation equivelent

Increasing the number of tournaments played while reducing the number of generations of the PBIL yields results that are generally not quite as favorable although

early performance does tend to be a little better (Figure 8). In the results show below the PBIL population size has been increased to 10 while the number of generations has been reduced to three thousand, various learning rates are shown:



Figure 8: Best GA v Population size 10, no mutation PBIL with varying learning rate. Generation 1-51, PBIL performs well or poorly depending on LR. Generation 51-101, PBIL performance begins to stabalize if it hasn't already. Generation 101 onwards, PBIL performance stabailizes but does worse than GA.

The population size 10 PBIL experiments run with standard mutation enabled, again perform somewhat worse (Figure 9):



Figure 9: GA v Population size 10 PBIL with standard mutation and varying learning rates. Generation 1-51, PBIL performs according to learning rate well or poorly. Generation 51-101, PBIL performance

stabalizes if it hasn't already. Generation 101 onwards, PBIL performance stabalizes and is worse than GA, often increasingly so.

The above results show that PBIL is beaten by GA in all of its basic configurations. The use of standard mutation generally reduces performance still further. We now introduce the use of guided mutation and gauge how it affects performance.

In the graph below the best GA configuration is pitted against PBIL with guided mutation enabled at various, fairly low, levels (Figure 10):



Figure 10: PBIL with guided mutation v standard configuration GA. PBIL performs well against GA to begin with losing ground around generation 50. Low (0.005) and high (0.4) levels of guided mutation do not provide a performance improvement. A level of 0.01 allows PBIL to compare favourably with GA by around generation 200. The larger sample is 50 PBIL and GA experiments being used.

The result above shows that guided mutation allows PBIL to recover and compete favourably with the GA when used at low levels. Far higher levels of guided mutation have also been investigated and are shown below (Figure 11):



Figure 11: PBIL with guided mutation v best configuration GA. Highest and lowest levels of guided mutation (0.99 and 0.8) perform well in the

first 25 generations with performance dropping of thereafter. Guide mutation at 0.9 starts well and continues to perform better than GA continually. The larger sample is 20 PBIL experiments v 50 GA experiments.

The above results show that a high level of guided mutation can boost PBIL performance beyond that of the standard best GA configuration. This seems somewhat surprising and suggests that a smaller variation in the PBIL population from generation to generation can make a significant contribution to performance.

As a final example we show that PBIL is effective against GA in the case of only a small population size is available. Approximating the original Axelrod experiments [1], the GA population is reduced to size 20 with a mutation rate of 0.7%. The tes is run for 300 generations, somewhat longer than the orginal experiments. This GA is run against an opponent PBIL using a population size of 5, training pool of 19 and running for 1200 generations without any mutation.
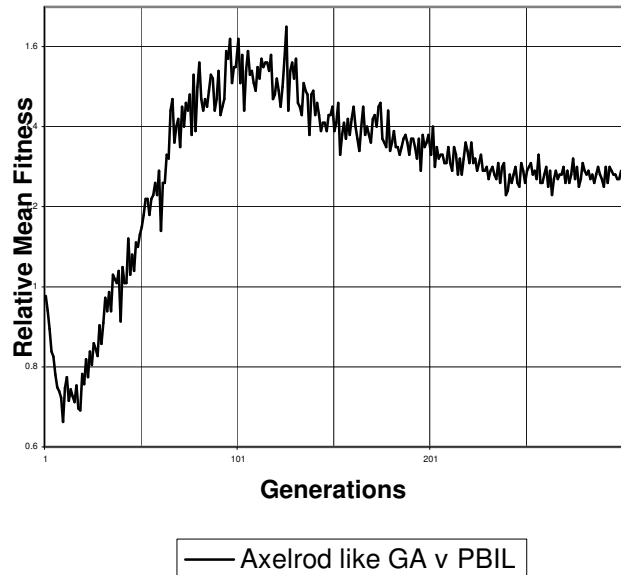


Figure 12: Axelrod like GA versus PBIL. Generation 1-51, PBIL initially does badly against the GA due to its low learning rate but is rapidly gaining ground after about 25 generations. Generation 51-101, PBIL continues to gain grownd and reaches a peak in realtive performance. Generation 101-201, PBILs relative performance is degraded but is still better than GA. Generation 201 onwards, PBILs performance stabilizes and does better than GA.

As can be seen from the above, when the GAs population is quite small the use of a probability distribution by PBIL comes in to its own – it might not use any more evaluations but it can certainly generate more effective strategies even if it takes it a while.

## IX. CONCLUSIONS

The results above show that the use of guided mutation in conjunction with PBIL is able to beat a standard GA in evolving IPD strategies.

Standard PBIL without guided mutation is unable to beat GA at evolving IPD strategies and consistently does slightly worse once the systems have stabilized.

The initial relative performance of PBIL may be superior to GA if the learning rate is high enough as it is able to learn fairly effective strategies quickly; however this tends to be accompanied by later worse performance due to overly rapid convergence. Setting the learning rate of the PBIL system lower tends to allow the GA a better run in the beginning and reduces the negative effects of rapid convergence.

Increasing the PBIL population size seems to cause far greater variation in performance relative to the learning rate used and may help reduce a performance gap when guided mutation is not used.

Introducing standard mutation into the PBIL system doesn't seem to remedy the effects of overly rapid convergence and indeed seems to generally exacerbate the problem. At higher learning rates and larger population sizes mutation would seem to help smooth performance out a little over time if not actually improve it.

When the GA population size is reduced PBIL is able to perform significantly better than GA. While the GA struggles under a lack of evaluations the PBIL is able to use what evaluations it has in the context of a hypothetically infinite population and respond accordingly. When the GA population is increased PBIL's advantage tends to be lost.

Any set of comparisons is difficult. The set of comparisons shown here only encompasses one particular problem using a limited number of configurations of both systems, as such its results shouldn't be considered definitive in arguing whether PBIL or GA is more appropriate under much wider conditions. It does however lend additional support to the use of guided mutation.

Further work on the effectiveness of comparable GA and PBIL systems should be undertaken to better understand when each is more applicable; some form of co evolutionary approach may yield interesting results for instance. It is important to establish when a technique is effective and when it is not.

## REFERENCES

[1]  Robert Axelrod, *The Evolution of Strategies in the Iterated Prisoner's Dilemma,* in Davis, L. (ed.), Genetic algorithms and simulated annealing, Research notes in AI, Pitman/Morgan Kaufmann, 1987, 32-41

[2]  Ken Binmore, *Game Theory and the Social Contract I, playing fair*, MIT Press, 1994

[3]  Ken Binmore, *Game Theory and the Social Contract II, just playing*, MIT Press, 1998

[4]  Linster, B., *Evolutionary Stability in the Repeated Prisoners' Dilemma Played by Two-State Moore Machines*, Southern Economic Journal, 1992, pages 880-903

[5]  Fogel, D.B. (1993). Evolving behaviors in the iterated prisoner's dilemma. Evolutionary Computation, 1(1), 77-97.

[6]  Crowley, P.H. (1996). Evolving cooperation: strategies as hierarchies of rules.    BioSystems, 37:67-80.

[7]  Shumeet Baluja, Population Based Incremental Learning – A Method for Integrating Genetic Search Based Function Optimisation and Competitive Learning, *(Tech. Rep. No. CMU-CS-94-163). Pittsburgh, PA: Carnegie Mellon University (1994)*

[8]  Numerical Recipes in C++: The Art of Scientific Computing – p? - William H. Press, Saul A. Teukolsky, William T. Vetterling, Brian P. Flannery

[9]  Course cc385, "Genetic Programming and its Applications", 2003 – University Of Essex, Prof Riccardo Poli

[10] Qingfu Zhang, Jianyong Sun and Edward Tsang, Evolutionary Algorithm with Guided Mutation for the Maximum Clique Problem, TEC 2005