

Behavioral Diversity, Choices and Noise in the Iterated Prisoner's Dilemma

Siang Y. Chong, *Member, IEEE*, and Xin Yao, *Fellow, IEEE*

Abstract—Real-world dilemmas rarely involve just two choices and perfect interactions without mistakes. In the iterated prisoner's dilemma (IPD) game, intermediate choices or mistakes (noise) have been introduced to extend its realism. This paper studies the IPD game with both noise and multiple levels of cooperation (intermediate choices) in a coevolutionary environment, where players can learn and adapt their strategies through an evolutionary algorithm. The impact of noise on the evolution of cooperation is first examined. It is shown that the coevolutionary models presented in this paper are robust against low noise (when mistakes occur with low probability). That is, low levels of noise have little impact on the evolution of cooperation. On the other hand, high noise (when mistakes occur with high probability) creates misunderstandings and discourages cooperation. However, the evolution of cooperation in the IPD with more choices in a coevolutionary learning setting also depends on behavioral diversity. This paper further investigates the issue of behavioral diversity in the coevolution of strategies for the IPD with more choices and noise. The evolution of cooperation is more difficult to achieve if a coevolutionary model with low behavioral diversity is used for IPD games with higher levels of noise. The coevolutionary model with high behavioral diversity in the population is more resistant to noise. It is shown that strategy representations can have a significant impact on the evolutionary outcomes because of different behavioral diversities that they generate. The results further show the importance of behavioral diversity in coevolutionary learning.

Index Terms—Behavioral diversity, coevolution, coevolutionary learning, evolutionary computation, iterated prisoner's dilemma (IPD), representation.

I. INTRODUCTION

THE ABSTRACT mathematical game of the iterated prisoner's dilemma (IPD) has been studied widely in many areas [1]. In its classical form, it is a two-player, nonzero sum, and noncooperative game. It has long been used to model social, economic, biological, and other interactions, and to study the conditions that allow for mutual cooperation to occur within a group of selfish individuals. In the IPD game, two isolated players engage in repeated interactions and can choose between two choices: cooperate and defect. It is "justifiably famous as an elegant embodiment of the tension between individual rationality (reflected in the incentive of both sides to be selfish) and group rationality (reflected in the higher payoff to both sides for mutual cooperation)" [2]. In the IPD, defection is not always the best choice. Axelrod [2], [3] showed through tournament games that a particular form of cooperative play can be a viable

strategy. Many others have also shown that cooperative play can be learned from a set of random strategies using an evolutionary algorithm [4]–[8] even though the evolution is unstable [9].

Early studies of the IPD using the coevolutionary approach [4]–[8] each assumed only two choices, i.e., cooperation and defection, for each player. Further studies have extended the IPD to bridge the gap between the model and real-world dilemmas. The IPD with more choices [10]–[13], i.e., multiple levels of cooperation, was studied. It allows modeling subtle behavioral interactions that are not possible with just two extreme choices. It was shown that with more choices, the evolution of cooperation is unstable [10] and more difficult to achieve [11]–[13].

The other extension to the classical IPD is the consideration of noise, i.e., misinterpretations or mistakes, which is used to model uncertainty in behavioral interactions [14]–[19]. Julstrom [20] studied noise in the classic *two-choice* IPD. Wahl and Nowak [21], [22] considered a continuous prisoner's dilemma using linear reactive strategies, with and without noise. (Their noise model is different from what we are using here.) However, none of the previous studies has investigated the impact of noise on the coevolutionary learning of strategies in the IPD with more choices. In classical evolutionary games [14]–[18], [21], [22], strategies do not adapt as they do in coevolutionary learning from an initial set of random strategies. It is also worth noting that continuous prisoner's dilemma is different from the discrete IPD with multiple levels of cooperation.

Darwen and Yao [11]–[13], [23] were among the first who studied the IPD with multiple levels of cooperation in a coevolutionary learning environment. They showed that cooperative strategies were more difficult to evolve as the number of cooperation levels increased. However, a very simple reputation could help to encourage cooperation [23]. Noise was not considered in their studies.

This paper presents a detailed empirical study of the impact of behavioral diversity in the IPD with multiple levels of cooperation and noise in a coevolutionary learning environment. Two major research issues are examined. First, we study how the level of noise (from low to high probability levels of making mistakes) influences the evolution of cooperation in the IPD with multiple levels of cooperation. Results show that different levels of noise lead to different evolutionary outcomes. More importantly, different coevolutionary models exhibit different evolutionary dynamics and lead to different evolutionary outcomes in the noisy IPD games with more choices. It turns out that behavioral diversity, rather than genetic diversity, plays a crucial role here. Second, to support this explanation, we study the role of behavioral diversity explicitly. A comparative study

Manuscript received July 31, 2004; revised June 8, 2005.

The authors are with the School of Computer Science, The University of Birmingham, Edgbaston, Birmingham B15 2TT, U.K. (e-mail: S.Y.Chong@cs.bham.ac.uk; X.Yao@cs.bham.ac.uk).

Digital Object Identifier 10.1109/TEVC.2005.856200

	Cooperate	Defect
Cooperate	R	T
Defect	S	P

Fig. 1. Payoff matrix for the two-player IPD. The payoff listed in the lower left-hand corner is assigned to the player choosing the move.

	Cooperate	Defect
Cooperate	4	5
Defect	0	1

Fig. 2. Payoff matrix for the two-player IPD used in this paper. $T = 5$, $R = 4$, $P = 1$, and $S = 0$.

is conducted between two different strategy representation schemes, which can produce different behavioral diversities. We show that through direct evolution of strategy's behaviors (using the lookup table representation), behavioral diversity of the coevolving population increases and, as a result, a smooth evolution toward cooperation can be achieved easily. We also distinguish between the effects of direct behavioral evolution and that of noise on the coevolutionary process.

The rest of this paper is organized as follows. Section II introduces the classical and extended IPD game models used in our experiments. Section III describes the two strategy-representation schemes, neural networks and lookup tables, and the coevolutionary model. Section IV presents our experimental studies. Section V discusses the experimental results. Finally, Section VI concludes this paper with some remarks on future studies.

II. IPD OVERVIEW

A. Classic IPD

In the classic IPD, each player has two choices: cooperation and defection. The payoff a player receives depends on a payoff matrix (Fig. 1) that must satisfy the following three conditions:

- 1) $T > R$ and $P > S$ (defection always pays more);
- 2) $R > P$ (mutual cooperation beats mutual defection);
- 3) $R > (S + T)/2$ (alternating does not pay).

There are many possible values for T , R , P , and S that satisfy the above three conditions. We use $T = 5$, $R = 4$, $P = 1$, and $S = 0$ (Fig. 2).

B. IPD With Multiple Levels of Cooperation and Noise

The IPD with multiple levels of cooperation and noise follows previous work [11]–[13], [23], [24]. The model is based on a simple interpolation of the classic *two-choice* IPD. Specifically,

		PLAYER B			
		+1	$+\frac{1}{3}$	$-\frac{1}{3}$	-1
PLAYER A	+1	4	$2\frac{2}{3}$	$1\frac{1}{3}$	0
	$+\frac{1}{3}$	$4\frac{1}{3}$	3	$1\frac{2}{3}$	$\frac{1}{3}$
	$-\frac{1}{3}$	$4\frac{2}{3}$	$3\frac{1}{3}$	2	$\frac{2}{3}$
	-1	5	$3\frac{2}{3}$	$2\frac{1}{3}$	1

Fig. 3. Payoff matrix for the two-player *four-choice* IPD used in this paper. Each element of the matrix gives the payoff for Player A.

the n -choice IPD is linearly interpolated from the *two-choice* IPD using the following equation [13]:

$$p_A = 2.5 - 0.5c_A + 2c_B, \quad -1 \leq c_A, c_B \leq 1$$

where p_A is the payoff to player A, given that c_A and c_B are the cooperation levels of the choices that players A and B make, respectively.

In generating the payoff matrix for the IPD with n -choices, the following conditions must be satisfied:

- 1) for $c_A < c'_A$ and constant c_B :
 $p_A(c_A, c_B) > p_A(c'_A, c_B)$;
- 2) for $c_A \leq c'_A$ and $c_B < c'_B$:
 $p_A(c_A, c_B) < p_A(c'_A, c'_B)$;
- 3) for $c_A < c'_A$ and $c_B < c'_B$:
 $p_A(c'_A, c'_B) > (1/2)(p_A(c_A, c'_B) + p_A(c'_A, c_B))$.

The above conditions are analogous to those for the *two-choice* IPD's. The first condition ensures that defection always pays more. The second condition ensures that mutual cooperation has a higher payoff than mutual defection. The third condition ensures that alternating between cooperation and defection does not pay in comparison to just playing cooperation.

Given the payoff equation and the three conditions above, an n -choice IPD can be formulated [11]. Here, a *four-choice* IPD (Fig. 3) is used, which is sufficient to investigate the two major research issues mentioned in Section I. Having more choices than four increases the computational cost of the experiments and would have made the analysis of results unnecessarily complicated.

Fig. 3 illustrates two important points. First, the payoffs in the four corners of an n -choice IPD payoff matrix are the same as those in the *two-choice* IPD. Second, any 2×2 submatrix of the $n \times n$ matrix of the n -choice IPD is itself a *two-choice* IPD.

Noise is modeled as the mistake that a player makes. It captures the situation where a player intends to play a choice but ends up with a different choice instead. It has a certain probability of occurrence p_n and is fixed throughout a game. Consider an example of a *four-choice* IPD game with $p_n = 0.05$. Suppose a player chooses to play at $+1/3$ cooperation level. There will be a 0.05 probability that this cooperation level will not be made. When this occurs, one of the other three possible cooperation levels, i.e., $+1$, $-1/3$, and -1 , will be chosen uniformly at random. This simplifies the implementation and ensures that there is no bias of play toward a specific cooperation level. The noise model here is different from some other implementations [22] where noise is modeled as errors that occur in the interpretation of opponent's move. As in [22], we note that the noise

model here is an extreme case, where players do not know that mistakes have occurred.

III. STRATEGY REPRESENTATION AND COEVOLUTIONARY MODEL

Neural networks have often been used to represent strategies in a coevolutionary learning environment [6], [10]–[13], [23]. However, it has been pointed out that the many-to-one mapping from genotypes (neural networks) to phenotypes (strategy behaviors) may have contributed to the difficulty in evolving cooperative strategies in the IPD with more choices. It is unknown whether this would still be the case for the IPD with both more choices and noise. To study the issue in more depth, we compare two strategy representation schemes in our experiments. One is based on neural networks. The other is based on lookup tables.

We restrict the memory length of every strategy to be one in order to facilitate the analysis of the results and reduce the computational cost.

A. Neural Network Representation of Strategies

In the IPD, strategies can be represented by binary strings [4], [8], finite-state machines (FSMs) [5], [7], [25], and neural networks [6], [10]–[13], [23]. Neural network representation is attractive because it can scale up easily as the number of cooperation levels increases.

1) *Neural Network Architecture:* We use a fixed-architecture feed-forward multilayer perceptron (MLP) to represent IPD strategies. Specifically, the neural network consists of an input layer, a single hidden layer, and an output node. The input layer consists of the following four input nodes.

- 1) The neural network's previous choice, i.e., level of cooperation, in $[-1, +1]$.
- 2) The opponent's previous level of cooperation.
- 3) An input of +1 if the opponent played a lower cooperation level compared to the neural network, and 0, otherwise.
- 4) An input of +1 if the neural network played a lower cooperation level compared to the opponent, and 0, otherwise.

Since the last two inputs are derived from the first two inputs, the input layer is a function of just two variables: the neural network's previous choice and the opponent's previous choice. The reason for the additional two input nodes is to facilitate learning the recognition of being exploited and exploiting [11], [23].

The number of hidden nodes used in the experiments is ten and has not been optimized. The network is fully connected and strictly layered. That is, there is no shortcut connection from the input layer to the output node. The transfer (activation) function used for all nodes is the hyperbolic tangent function $\tanh(x)$.

Given the inputs, the neural network's output determines the choice for its next move. The neural network output is a real value between +1 and -1. It is discretized to either +1, +1/3, -1/3, or -1, depending on which discrete value the neural network output is closest to.

With regards to the choice of using a single hidden layer of ten nodes, previous studies on IPD with multiple choices have shown that the coevolution of neural networks with ten hidden

nodes is capable of producing mutual cooperative behaviors [13].

2) *Evolution of Neural Network Weights:* As summarized by Yao [26], real-valued weights of a neural network can be evolved using self-adaptive evolutionary algorithms, such as evolutionary programming [27]. Here, a neural network is associated with a self-adaptive parameter vector $[\sigma_i(j)]$ that controls the mutation step size of the respective weights and biases of the neural network $[w_i(j)]$.

Offspring neural networks ($[w'_i(j)]$ and $[\sigma'_i(j)]$) are generated from parent neural networks ($[w_i(j)]$ and $[\sigma_i(j)]$) through mutation. Two different mutations, Gaussian and Cauchy, are used in order to investigate the effects of variation operators on evolving IPD strategies.

For the self-adaptive Gaussian mutation, offspring neural networks are generated according to the following equations:

$$\begin{aligned}\sigma'_i(j) &= \sigma_i(j) * \exp(\tau * N_j(0, 1)); \\ i &= 1, \dots, 15, \quad j = 1, \dots, N_w \\ w'_i(j) &= w_i(j) + \sigma'_i(j) * N_j(0, 1); \\ i &= 1, \dots, 15, \quad j = 1, \dots, N_w\end{aligned}$$

where $N_w = 63$, $\tau = (2(N_w)^{0.5})^{-0.5} = 0.251$, and $N_j(0, 1)$ is a Gaussian random variable (zero mean and standard deviation of one) resampled for every j . N_w is the total number of weights, biases, and the pre-game inputs required for an IPD strategy based on memory length of one.

For the self-adaptive Cauchy mutation [27], the following equations are used:

$$\begin{aligned}\sigma'_i(j) &= \sigma_i(j) * \exp(\tau * N_j(0, 1)); \\ i &= 1, \dots, 15; \quad j = 1, \dots, N_w \\ w'_i(j) &= w_i(j) + \sigma'_i(j) * C_j(0, 1); \\ i &= 1, \dots, 15; \quad j = 1, \dots, N_w\end{aligned}$$

where $C_j(0, 1)$ is a Cauchy random variable (centered at zero and with a scale parameter of 1) resampled for every j . All other variables remain the same as those in the self-adaptive Gaussian mutation.

B. Lookup Table Representation of Strategies

The lookup table representation used here is different from that used in [4], [8], where the genotypic representation used is a binary string. Our lookup table is used to represent directly the strategy map. It is treated as a two-dimensional representation. The variation operators used to generate offspring are applied to the strategy maps directly and explicitly. As such, the evolutionary process is applied directly at the phenotypic (behavioral) level of the IPD strategies. This avoids the potential problem of having many-to-one mappings between genotypes and phenotypes [28], [29] associated with other representations such as neural networks.

Fig. 4 illustrates the lookup table representation for the IPD strategies with four choices and memory length of one. m_{ij} specifies the choice to be made, given the inputs i (player's own previous choice) and j (opponent's previous choice). Rather

		Opponent's Previous Move			
		+1	$+\frac{1}{3}$	$-\frac{1}{3}$	-1
Player's Previous Move	+1	m_{11}	m_{12}	m_{13}	m_{14}
	$+\frac{1}{3}$	m_{21}	m_{22}	m_{23}	m_{24}
	$-\frac{1}{3}$	m_{31}	m_{32}	m_{33}	m_{34}
	-1	m_{41}	m_{42}	m_{43}	m_{44}

Fig. 4. The lookup table representation for the two-player IPD with four choices and memory length one.

than using pre-game inputs (two for memory length one strategies), the first move is specified independently. Instead of using quaternary values (0, 1, 2, 3), each of the table elements can take any of the possible four choices (+1, +1/3, -1/3, -1).

A simple mutation operator is used to generate offspring. Mutation replaces the original element m_{ij} by one of the other three possible choices with an equal probability. For example, if mutation occurs at $m_{13} = +1/3$, then the mutated element m'_{13} can take either +1, -1/3 or -1 with an equal probability. Each table element has a fixed probability p_m of being replaced by one of the remaining three choices. The value p_m is not optimized. Crossover is not used in any of the experiments.

Although the lookup table used here and the one used by Axelrod in [4] similarly specify a strategy's responses that are indexed by previous moves, they are different if we consider the evolutionary algorithm as a whole. An evolutionary algorithm depends on the interactions between variation operators and the representation. There is an important difference between our approach (e.g., varying the lookup table elements directly) and Axelrod's approach (e.g., varying lookup table encoded as binary strings).

C. Coevolutionary Model

The coevolutionary procedure involves a population of 30 strategies. A small population size in coevolution will reduce the level of behavioral diversity in the population, allowing strategies to overspecialize and, thus inducing cycles between cooperation and defection. This facilitates the investigation into how behavioral diversity is influenced by different parts of the model and how the evolutionary outcome is influenced by different behavioral diversity.

Coevolution starts with random initialization of 15 parents, which then generate 15 offspring. In initializing the strategies, additional steps are taken to ensure that the initial population has sufficient behavioral diversity in addition to genotypic diversity [11]. The following setting is adopted after some preliminary experimentations.

For experiments using neural networks, the weights and biases of the parent neural networks are initialized randomly by sampling from a uniform distribution over $[-2.45, +2.45]$. Each component of the self-adaptive parameter vector of the parent neural networks is initialized to 0.5 for consistency with the initialization range of neural networks' weights and biases. After that, outputs of each neural network are evaluated and




Fig. 5. Evolution of cooperation for C-CEP and C-FEP in the four-choice IPD.

discretized. For an IPD game with four choices and memory length of one, we required that the frequency at which each of the four choices (+1, +1/3, -1/3, -1) is selected for each of the parent neural networks' strategies (input-output mappings) is roughly the same. The parents are resampled as many times as necessary to satisfy this condition. A similar check is made for experiments that used lookup tables.

After generating the offspring, each strategy in the population competes with every other strategy, including itself, in a round-robin fashion (e.g., each strategy competes in a total of 30 games). The fitness of a strategy is obtained by taking the average payoff that it receives from all the games it played. The best half of the population is selected as parents for the next generation. The cycle of generating offspring, competition, and selection is repeated for 600 generations. Each experiment is repeated for 30 independent runs.

D. Shadow of the Future

In the IPD, the shadow of the future refers to the situation whereby the number of iterations of a game is known in advance. In this situation, there is no incentive to cooperate in the last iteration because there is no risk of retaliation from the opponent. However, if every player defects on the last iteration, then there is no incentive to cooperate in the iteration prior to the last one. If every player defects in the last two iterations, then there is no incentive to cooperate in the iteration before that, and so forth. So we would end up with mutual defection in all iterations.

One popular way to address this issue and to allow for cooperation to emerge is to have a fixed probability in ending the game on every iteration, thereby keeping the game length uncertain. In our experiments, neither of our representations includes the information on the number of iterations. It is impossible for our strategies to count how many iterations have been played and how many remain. We simply use a fixed game length of 150 iterations (move start from 0) in all our game plays. The game length is chosen to be comparable to the experiments conducted in [4] and [30] so that comparisons with previous work can be made.




Fig. 6. Frequencies of choices played for C-CEP in the *four-choice* IPD.

IV. EXPERIMENTAL STUDIES

A. Behavioral Diversity Helps the Evolution of Cooperation for the IPD With More Choices

Behavioral diversity indicates how different the played choices are in the IPD games. We examine two major sources of behavioral diversity in our experiments. The first is induced by genetic diversity, e.g., different weights and biases for neural networks. However, genetic diversity among neural networks may not lead to behavioral diversity among strategies. The second source to induce diversity is through direct and explicit generation of diversity by mutating a lookup table of strategy maps. We will compare the two different representations using the *four-choice* IPD.

Let the term “coevolutionary classical evolutionary programming” (C-CEP) refer to the coevolutionary model that uses the neural network representation with self-adaptive Gaussian mutation. We refer to the model that uses self-adaptive Cauchy mutation as “coevolutionary fast evolutionary programming” (C-FEP). We refer to the model that uses lookup table representation as C-PM05, C-PM10, C-PM15, C-PM20, and C-PM25, respectively, to indicate models that use p_m of 0.05, 0.10, 0.15, 0.20, and 0.25, respectively.

1) *Results With the Neural Network Representation:* Fig. 5 plots the average payoff over 30 runs for C-CEP and C-FEP. It shows that the evolutionary behaviors for C-CEP and C-FEP are similar to that of the classical *two-choice* IPD [4]. For both C-CEP and C-FEP, Fig. 5 shows a brief initial drop in average payoff before rising to a value greater than 3 (i.e., mutual cooperation).

Figs. 6 and 7 show the average frequencies for each of the four choices that were played at each generation for C-CEP and C-FEP, respectively. The frequency for each choice started around 25%, indicating good behavioral diversity in the initial

population. However, after a short period (around 10–15 generations), the frequencies for the four choices diverged. The frequency for full cooperation (+1) was significantly higher than other choices, indicating that there were more full cooperative plays by the evolving strategies.

Figs. 8 and 9 show the plots for five individual runs of C-CEP and C-FEP, respectively. The figures indicate that not all runs had converged to fully cooperative strategies within 600 generations. Some runs stayed at the intermediate levels of cooperation, while a few converged to defection. This is quite different from the case for classical *two-choice* IPD games [4], [8], where each run converged to mutual cooperation quite consistently and quickly. It illustrates that more choices have made cooperation more difficult to evolve.

2) *Results With the Lookup Table Representation:* Similar experiments were conducted for C-PM05, C-PM10, C-PM15, C-PM20, and C-PM25, respectively. The average results over 30 runs showed general evolutionary trends similar to those observed from the classic *two-choice* IPD [4] and to those in Fig. 5, i.e., the average cooperation level decreased initially before restoring and increasing to full cooperation. The figures are omitted from this paper for brevity.

Fig. 10 shows the average frequencies of four different choices played at each generation for C-PM05. It is evident that full cooperation (+1) was played far more often than others. A similar observation can be made for C-PM10, C-PM15, C-PM20, and C-PM25 as well. Corresponding figures are again omitted for brevity.

Fig. 11 plots five individual runs for C-PM05, which show that the majority of the runs evolved to a very high cooperation level (≥ 3.5). Similar results were obtained for other C-PMs. No runs converged to defection.

Table I summarizes the experiment results for the *four-choice* IPD. Three observations can be made.




Fig. 7. Frequencies of choices played for C-FEP in the *four-choice* IPD.




Fig. 8. Five sample runs of C-CEP in the *four-choice* IPD.




Fig. 9. Five sample runs of C-FEP in the *four-choice* IPD.

First, the most significant result in the table is that no run converged to defection when the lookup table representation was used. This demonstrates the effectiveness of direct and explicit manipulation of behavioral diversity in coevolving strategies.

Second, comparing the results from C-CEP and C-FEP, it appears that Cauchy mutation is less effective in increasing behavioral diversity in the population in comparison with the case where a lookup table is used. There was no statistically significant difference between C-CEP and C-FEP in spite of Cauchy mutation, which increases the genetic diversity by providing larger variations to the neural network's weights [27]. As observed previously [11], an increase in genetic diversity does not mean an increase in behavioral diversity. It is behavioral diversity that helps to coevolve cooperation [13]. An indirect method of varying a strategy's behavior (e.g., neural networks) is less

effective than a direct method (e.g., lookup tables) in increasing behavioral diversity.

Third, a small mutation rate for the lookup table representation, i.e., C-PM05 and C-PM10, can increase and maintain behavioral diversity significantly better than others, while not degenerating the coevolutionary algorithm into random search. High mutation rates appear to disrupt evolution and make convergence to the high level of cooperation more difficult. However, no defection was observed in any of our runs.

B. How Does Noise Affect Coevolutionary Learning of Strategies?

We now compare the evolutionary behaviors of different coevolutionary models using the *four-choice* IPD with noise. We study the effects of various noise probability settings. For low




Fig. 10. Frequencies of choices played for C-PM05 in the *four-choice* IPD.




Fig. 11. Five sample runs of C-PM05 in the *four-choice* IPD.

p_n settings, we use 0.05%¹ ($p_n = 0.0005$), 0.10%, 0.15%, 0.20%, 0.25%, 0.50%, and 1.50%. For high p_n settings, we use 5%, 10%, 15%, and 20%.

1) *Coevolutionary Models Are Robust Against Low Noise:* Tables II and III summarize the comparisons of different coevolutionary models when p_n was set to 0.05% and 1.50%, respectively. Results for other low p_n values are omitted.

Comparing Tables II and III against Table I, there was no statistically significant difference between the low noise and noiseless cases at the end of the evolutionary process, which indicates that the coevolutionary models used here were quite robust against low noise.

On the evolutionary process itself, there were more runs that exhibited large fluctuations between full cooperation and full

defection when noise was introduced to C-CEP and C-FEP. When $p_n = 0.05\%$, the fluctuations were less obvious (Figs. 12 and 13). When p_n setting was increased, however, the fluctuations were more obvious. For example, cycles between full cooperation and full defection can be seen clearly from Fig. 14 (C-CEP) and Fig. 15 (C-FEP) when $p_n = 1.50\%$.

Comparing various C-PMs with C-CEP and C-FEP shows that there were more runs in the C-PMs that evolved to higher cooperation levels (≥ 3.5) for all low p_n values. For C-PMs, no run exhibited large fluctuations between full cooperation and full defection that were associated with the runs in C-CEP and C-FEP. For example, Fig. 16 (C-PM05) and Fig. 17 (C-PM25) have shown quite different characteristics from those in Fig. 12 (C-CEP) and Fig. 13 (C-FEP). They have also shown that consistently high levels of cooperation could be evolved with low level of noise. In other words, the lookup table representation is more robust in the face of low noise.

2) *High Noise Causes Misunderstanding and Destroys Cooperation:* Axelrod [15] noted that with increasing noise in the IPD, cooperation becomes less likely as generosity invites more exploitation. Given high p_n values, we investigate our coevolutionary models' resistance to defection. Tables IV–VI summarize the results for C-CEP, C-FEP, and C-PM05, respectively. Results for other C-PMs are omitted for space.

Tables IV and V show that for models that used the neural network representation, there was an increasing tendency for runs to evolve toward mutual defection as p_n increased. Observation on individual runs (not detailed in this paper) showed a consistent trend toward defection. On the other hand, Table VI shows that for similar p_n settings, no run evolved to defection for C-PM05. Observation on individual runs showed cycles between full cooperation and full defection for C-PM05.

The results indicate that C-PMs are more robust than C-CEP and C-FEP. Even for high noise, C-PMs could still converge to

¹%'s are not probabilities.

TABLE I
COMPARISON BETWEEN COEVOLUTIONARY MODELS THAT USE DIFFERENT STRATEGY REPRESENTATIONS FOR THE *FOUR-CHOICE* IPD. RESULTS WERE TAKEN FROM THE FINAL GENERATION. "MEAN" INDICATES THE AVERAGE OVER 30 RUNS. "STD DEV" INDICATES THE STANDARD DEVIATION. "MAX" INDICATES THE HIGHEST POPULATION PAYOFF. "MIN" INDICATES THE LOWEST POPULATION PAYOFF. " $No \leq 1.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH DEFECTION. " $1.5 < No < 3.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH INTERMEDIATE PAYOFFS. " $No \geq 3.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH FULL COOPERATION. THE *t*-TESTS COMPARE C-CEP WITH OTHER MODELS

Experiment	Mean \pm Std Dev	Max	Min	$No \leq 1.5$	$1.5 < No < 3.5$	$No \geq 3.5$	<i>t</i> -test
C-CEP	3.24 ± 1.11	4.00	1.03	4	6	20	
C-FEP	3.08 ± 0.75	4.00	1.10	1	19	10	0.62
C-PM05	3.72 ± 0.34	4.00	2.87	0	5	25	-2.41 [†]
C-PM10	3.70 ± 0.37	4.00	1.92	0	3	27	-2.45 [†]
C-PM15	3.57 ± 0.25	3.92	2.79	0	6	24	-1.58
C-PM20	3.48 ± 0.15	3.75	3.16	0	16	14	-1.24
C-PM25	3.39 ± 0.19	3.85	3.09	0	20	10	-0.77

[†] The *t* value with 29 degrees of freedom is statistically significant at a 0.05 level of significance using a two-tailed test.

TABLE II
COMPARISON BETWEEN COEVOLUTIONARY MODELS THAT USE DIFFERENT STRATEGY REPRESENTATIONS FOR THE *FOUR-CHOICE* IPD WITH p_n SET AT 0.05%. RESULTS WERE TAKEN FROM THE FINAL GENERATION. "MEAN" INDICATES THE AVERAGE OVER 30 RUNS. "STD DEV" INDICATES THE STANDARD DEVIATION. "MAX" INDICATES THE HIGHEST POPULATION PAYOFF. "MIN" INDICATES THE LOWEST POPULATION PAYOFF. " $No \leq 1.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH DEFECTION. " $1.5 < No < 3.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH INTERMEDIATE PAYOFFS. " $No \geq 3.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH FULL COOPERATION. THE *t*-TESTS COMPARE C-CEP WITH OTHER MODELS

Experiment	Mean \pm Std Dev	Max	Min	$No \leq 1.5$	$1.5 < No < 3.5$	$No \geq 3.5$	<i>t</i> -test
C-CEP	3.49 ± 0.94	4.00	1.00	3	4	23	
C-FEP	3.21 ± 0.77	3.99	1.00	2	15	13	1.16
C-PM05	3.86 ± 0.21	3.99	2.90	0	2	28	-2.07 [†]
C-PM10	3.80 ± 0.12	3.98	3.55	0	0	30	-1.80
C-PM15	3.64 ± 0.22	3.93	2.78	0	3	27	-0.87
C-PM20	3.50 ± 0.13	3.70	3.28	0	16	14	-0.06
C-PM25	3.39 ± 0.16	3.77	3.11	0	25	5	0.54

[†] The *t* value with 29 degrees of freedom is statistically significant at a 0.05 level of significance using a two-tailed test.

a certain degree of cooperation and avoid full defection. This was not the case for C-CEP and C-FEP. However, cycles between full cooperation and full defection could be observed for C-PMs under very high noise during the evolutionary process.

V. DISCUSSION

A. Behavioral Diversity Is Crucial for Evolution to Higher Cooperation Level in the IPD With More Choices and Noise

McNamara *et al.* [31] showed that stable evolution to cooperation can be achieved in the IPD if sufficient variations in the behavior of strategies are maintained. In the continuous IPD with noise [22], it was shown that evolution to cooperation was unstable. Cycles between cooperation and defection could occur. Our studies show that stable evolution to cooperation can be achieved in the IPD with more choices and low noise if behavioral diversity is maintained, e.g., the case of C-PMs. The coevolutionary model using the lookup table representation is quite robust against noise.

The cycles between full cooperation and full defection observed in some of our experiments can be explained in terms of overspecialization of strategies in the same population, which

makes strategies vulnerable to mutants and noise. The generalization issue in coevolutionary learning was first discussed for the classical noise-free two-player IPD by Darwen and Yao [8]. According to their study [8], the evolutionary process could lead to a population of naive cooperators, through drift and selection, that were vulnerable to defectors. The defectors would then be replaced by retaliatory cooperators such as tit-for-tat. After tit-for-tat type strategies took over the population, drifts toward naive cooperators would happen again.

Although overspecialization in C-CEP and C-FEP could be addressed by using larger population sizes (e.g., a population size of 100), a smaller population can avoid cyclic evolutionary behaviors if behavioral diversity is maintained explicitly, as is the case for C-PMs. Analysis of individual strategies evolved by C-PMs has showed that their populations included a good mixture of different strategies. The population did not overspecialize. Naive cooperators were not observed to dominate the population at any time of evolution. The behavioral diversity can be very high for the lookup table representation because every mutation will result in a behavioral change. The degree of change is controlled directly and explicitly by p_m . For the neural network representation, not every weight change will lead to a behavioral change.

TABLE III
COMPARISON BETWEEN COEVOLUTIONARY MODELS THAT USE DIFFERENT STRATEGY REPRESENTATIONS FOR THE *Four-Choice* IPD WITH p_n SET AT 1.5%. RESULTS WERE TAKEN FROM THE FINAL GENERATION. “MEAN” INDICATES THE AVERAGE OVER 30 RUNS. “STD DEV” INDICATES THE STANDARD DEVIATION. “MAX” INDICATES THE HIGHEST POPULATION PAYOFF. “MIN” INDICATES THE LOWEST POPULATION PAYOFF. “ $N_o \leq 1.5$ ” INDICATES THE NUMBER OF RUNS THAT ENDED WITH DEFECTION. “ $1.5 < N_o < 3.5$ ” INDICATES THE NUMBER OF RUNS THAT ENDED WITH INTERMEDIATE PAYOFFS. “ $N_o \geq 3.5$ ” INDICATES THE NUMBER OF RUNS THAT ENDED WITH FULL COOPERATION. THE t -TESTS COMPARE C-CEP WITH OTHER MODELS

Experiment	Mean \pm Std Dev	Max	Min	$N_o \leq 1.5$	$1.5 < N_o < 3.5$	$N_o \geq 3.5$	t -test
C-CEP	3.25 ± 0.87	3.97	1.03	2	12	16	
C-FEP	3.09 ± 0.77	3.94	1.08	2	19	9	0.92
C-PM05	3.73 ± 0.14	3.88	3.22	0	1	29	-2.88^\dagger
C-PM10	3.68 ± 0.12	3.85	3.35	0	3	27	-2.61^\dagger
C-PM15	3.47 ± 0.15	3.74	3.19	0	17	13	-1.41
C-PM20	3.38 ± 0.13	3.64	3.10	0	24	6	-0.82
C-PM25	3.21 ± 0.15	3.41	2.86	0	30	0	0.25

† The t value with 29 degrees of freedom is statistically significant at a 0.05 level of significance using a two-tailed test.




Fig. 12. Five sample runs of C-CEP using the *four-choice* IPD with $p_n = 0.05\%$.




Fig. 14. Five sample runs of C-CEP using the *four-choice* IPD with $p_n = 1.50\%$.




Fig. 13. Five sample runs of C-FEP using the *four-choice* IPD with $p_n = 0.05\%$.




Fig. 15. Five sample runs of C-FEP using the *four-choice* IPD with $p_n = 1.50\%$.

B. Noise Does Not Necessarily Promote Behavioral Diversity

It was mentioned recently that sources of variations in models such as mutation and noise might encourage cooperation [31].

However, mutation and noise produce behavioral diversity differently. Mutation introduces strategies with different behaviors

TABLE IV

COMPARISON OF C-CEP BETWEEN VARIOUS HIGH p_n SETTINGS FOR THE *FOUR*-CHOICE IPD. RESULTS WERE TAKEN FROM THE FINAL GENERATION. "MEAN" INDICATES THE AVERAGE OVER 30 RUNS. "STD DEV" INDICATES THE STANDARD DEVIATION. "MAX" INDICATES THE HIGHEST POPULATION PAYOFF. "MIN" INDICATES THE LOWEST POPULATION PAYOFF. " $No \leq 1.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH DEFECTION. " $1.5 < No < 3.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH INTERMEDIATE PAYOFFS. " $No \geq 3.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH COOPERATION. THE t -TESTS COMPARE THE NOISELESS CASE WITH THE NOISY ONES

p_n (%)	Mean \pm Std Dev	Max	Min	$No \leq 1.5$	$1.5 < No < 3.5$	$No \geq 3.5$	t -test
0	3.24 \pm 1.11	4.00	1.03	4	6	20	
5	3.12 \pm 0.90	3.90	1.10	4	11	15	0.46
10	2.77 \pm 0.97	3.80	1.20	7	14	9	1.75
15	2.52 \pm 0.93	3.70	1.30	8	15	7	2.68 [†]
20	2.04 \pm 0.83	3.51	1.39	18	11	1	5.14 [†]

[†] The t value with 29 degrees of freedom is statistically significant at a 0.05 level of significance using a two-tailed test.

TABLE V

COMPARISON OF C-FEP BETWEEN VARIOUS HIGH p_n SETTINGS FOR THE *FOUR*-CHOICE IPD. RESULTS WERE TAKEN FROM THE FINAL GENERATION. "MEAN" INDICATES THE AVERAGE OVER 30 RUNS. "STD DEV" INDICATES THE STANDARD DEVIATION. "MAX" INDICATES THE HIGHEST POPULATION PAYOFF. "MIN" INDICATES THE LOWEST POPULATION PAYOFF. " $No \leq 1.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH DEFECTION. " $1.5 < No < 3.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH INTERMEDIATE PAYOFFS. " $No \geq 3.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH COOPERATION. THE t -TESTS COMPARE THE NOISELESS CASE WITH THE NOISY ONES

p_n (%)	Mean \pm Std Dev	Max	Min	$No \leq 1.5$	$1.5 < No < 3.5$	$No \geq 3.5$	t -test
0	3.08 \pm 0.75	4.00	1.10	1	19	10	
5	2.40 \pm 0.96	3.90	1.10	9	18	3	3.00 [†]
10	2.10 \pm 0.77	3.23	1.20	11	19	0	4.48 [†]
15	1.87 \pm 0.72	3.48	1.29	17	13	0	7.22 [†]
20	1.58 \pm 0.43	2.88	1.39	26	4	0	10.95 [†]

[†] The t value with 29 degrees of freedom is statistically significant at a 0.05 level of significance using a two-tailed test.

TABLE VI

COMPARISON OF C-PM05 BETWEEN VARIOUS HIGH p_n SETTINGS FOR THE *FOUR*-CHOICE IPD. RESULTS WERE TAKEN FROM THE FINAL GENERATION. "MEAN" INDICATES THE AVERAGE OVER 30 RUNS. "STD DEV" INDICATES THE STANDARD DEVIATION. "MAX" INDICATES THE HIGHEST POPULATION PAYOFF. "MIN" INDICATES THE LOWEST POPULATION PAYOFF. " $No \leq 1.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH DEFECTION. " $1.5 < No < 3.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH INTERMEDIATE PAYOFFS. " $No \geq 3.5$ " INDICATES THE NUMBER OF RUNS THAT ENDED WITH COOPERATION. THE t -TESTS COMPARE THE NOISELESS CASE WITH THE NOISY ONES

p_n (%)	Mean \pm Std Dev	Max	Min	$No \leq 1.5$	$1.5 < No < 3.5$	$No \geq 3.5$	t -test
0	3.72 \pm 0.34	4.00	2.87	0	5	25	
5	3.44 \pm 0.24	3.69	2.61	0	14	16	3.54 [†]
10	3.14 \pm 0.24	3.39	2.46	0	30	0	7.01 [†]
15	2.93 \pm 0.16	3.25	2.33	0	30	0	10.98 [†]
20	2.55 \pm 0.24	2.87	1.99	0	30	0	15.85 [†]

[†] The t value with 29 degrees of freedom is statistically significant at a 0.05 level of significance using a two-tailed test.

into the population. Noise allows other parts of a strategy's behavior that are not played otherwise in a noiseless IPD game to be accessed.

Our results show that noise does not necessarily promote behavioral diversity in the population that lead to a stable evolution to cooperation. This is especially true in the cases of C-CEP and C-FEP. Increasingly higher levels of noise did not improve the evolution of cooperation at all. In fact, it increased the number of runs that converged to full defection.

With higher levels of noise, closer inspection of evolved strategies revealed that overspecialization had occurred in the population, which led to cyclic behaviors in the evolutionary process. This was true even in experiments that used

lookup tables. High mutation rates and high noise were supposed to generate high diversity according to previous thinking in the literature. Our results show that this is not always the case.

More importantly, noise and mutation have different impacts on the evolutionary process. For example, increasingly higher levels of noise lead to mutual defection outcomes. Given a very noisy environment, strategies overspecialized to play defection only. This was not observed in the noiseless case of the IPD with increasingly more mutations. For example, increasingly higher mutation rates in C-PMs did not lead to mutual defection outcomes. Strategies were not observed to overspecialized to play defection, or any specific play.




Fig. 16. Five sample runs of C-PM05 using the *four-choice* IPD with $p_n = 1.50\%$.




Fig. 17. Five sample runs of C-PM25 using the *four-choice* IPD with $p_n = 1.50\%$.

VI. CONCLUSION

Behavioral diversity is shown to be important in coevolutionary learning of strategies for the IPD with multiple levels of cooperation and noise. A detailed empirical study has been carried out to find out how behavioral diversity affects the learning of strategies in the IPD game. Two main issues have been addressed here: how different levels of noise affect the evolutionary outcome and how behavioral diversity plays a key role in the evolution to cooperation.

First, our study shows that low levels of noise can help in the evolution of cooperation by introducing generosity in the strategies. However, increasingly higher levels of noise discourage cooperation because of increasing misunderstandings in the interactions. Instead of encouraging different behaviors, high levels of noise cause the population of strategies to overspecialize. This makes the strategies vulnerable to invading mutants, and can be observed by the cyclic dynamics between cooperation and defection in the evolutionary process. With very high levels of noise introduced in the game, strategies evolved to mutual defection.

Second, our study shows that behavioral diversity helps with the evolution to cooperation in noisy IPD games. In particular, it has been shown that strategy representations have a significant impact on the evolutionary outcomes because of the different levels of behavioral diversity that they generate. A coevolutionary model that uses lookup tables that allow direct evolution of strategy behaviors can induce higher levels of behavioral diversity and is more robust to the effects of noise in the IPD game. However, very high levels of noise still lead to mutual defection.

For future research directions, a more rigorous analysis of evolved strategies from generalization's point of view should be carried out. Comparisons between our noise model and Wahl and Nowak's [22] can be carried out. The impact of noise on more realistic and complex IPD games can also be investigated, e.g., IPD with neighborhoods and IPD with more than two players.

ACKNOWLEDGMENT

The authors are grateful to The Centre of Excellence for Research in Computational Intelligence and Applications (CERCIA), School of Computer Science, The University of Birmingham, for providing computing support in running the experiments. The authors are grateful to Dr. D. B. Fogel and four anonymous reviewers for their constructive comments that have helped greatly in improving this paper.

REFERENCES

- [1] R. Axelrod, *The Evolution of Cooperation*. New York, 1984. Basic Books.
- [2] —, "More effective choice in the prisoner's dilemma," *J. Conflict Resolution*, vol. 24, no. 3, pp. 379–403, Sep. 1980.
- [3] —, "Effective choice in the prisoner's dilemma," *J. Conflict Resolution*, vol. 24, no. 1, pp. 3–25, Mar. 1980.
- [4] —, "The evolution of strategies in the iterated prisoner's dilemma," in *Genetic Algorithms and Simulated Annealing*, L. D. Davis, Ed. San Mateo, CA: Morgan Kaufmann, 1987, ch. 3, pp. 32–41.
- [5] D. B. Fogel, "The evolution of intelligent decision making in gaming," *Int. J. Cybern. Syst.*, vol. 22, pp. 223–236, 1991.
- [6] —, "Evolving behaviors in the iterated prisoner's dilemma," *Evol. Comput.*, vol. 1, no. 1, pp. 77–97, 1993.
- [7] —, "On the relationship between the duration of an encounter and the evolution of cooperation in the iterated prisoner's dilemma," *Evol. Comput.*, vol. 3, no. 3, pp. 349–363, 1996.
- [8] P. Darwen and X. Yao, "On evolving robust strategies for iterated prisoner's dilemma," in *Progress in Evolutionary Computation*, ser. Lecture Notes in Artificial Intelligence, 1995, vol. 956, pp. 276–292.
- [9] J. Batali and P. Kitcher, "Evolution of altruism in optional and compulsory games," *J. Theor. Bio.*, vol. 175, pp. 161–171, 1995.
- [10] P. G. Harrald and D. B. Fogel, "Evolving continuous behaviors in the iterated prisoner's dilemma," *BioSystems (Special Issue on the Prisoner's Dilemma)*, vol. 37, pp. 135–145, 1996.
- [11] P. Darwen and X. Yao, "Does extra genetic diversity maintain escalation in a co-evolutionary arms race," *Int. J. Knowl. Based Intell. Eng. Syst.*, vol. 4, no. 3, pp. 191–200, Jul. 2000.
- [12] —, "Why more choices cause less cooperation in iterated prisoner's dilemma," in *Proc. 2001 Congr. Evol. Comput.*, 2001, pp. 987–994.
- [13] —, "Co-evolution in iterated prisoner's dilemma with intermediate levels of cooperation: Application to missile defense," *Int. J. Comput. Intell. Appl.*, vol. 2, no. 1, pp. 83–107, 2002.
- [14] P. Molander, "The optimal level of generosity in a selfish, uncertain environment," *J. Conflict Resolution*, vol. 29, no. 4, pp. 611–618, 1985.
- [15] R. Axelrod and D. Dion, "The further evolution of cooperation," *Science*, vol. 242, no. 4884, pp. 1385–1390, 1988.
- [16] J. Wu and R. Axelrod, "How to cope with noise in the iterated prisoner's dilemma," *J. Conflict Resolution*, vol. 39, no. 1, pp. 183–189, 1995.

- [17] J. Bendor, R. M. Kramer, and S. Stout, "When in doubt. . . cooperation in a noisy prisoner's dilemma," *J. Conflict Resolution*, vol. 35, no. 4, pp. 691–719, 1991.
- [18] D. B. Neill, "Optimality under noise: Higher memory strategies for the alternating prisoner's dilemma," *J. Theor. Bio.*, vol. 211, no. 2, pp. 159–180, 2001.
- [19] B. Sainty, "Achieving greater cooperation in a noisy prisoner's dilemma: An experimental investigation," *J. Econ. Behav. Org.*, vol. 39, no. 4, pp. 421–435, 1999.
- [20] B. A. Julstrom, "Effects of contest length and noise on reciprocal altruism, cooperation, and payoffs in the iterated prisoner's dilemma," in *Proc. 7th Int. Conf. Genetic Algorithms*, 1997, pp. 386–392.
- [21] L. M. Wahl and M. A. Nowak, "The continuous prisoner's dilemma: I. Linear reactive strategies," *J. Theor. Bio.*, vol. 200, no. 3, pp. 307–321, 1999.
- [22] —, "The continuous prisoner's dilemma: II. Linear reactive strategies with noise," *J. Theor. Bio.*, vol. 200, no. 3, pp. 323–338, 1999.
- [23] X. Yao and P. Darwen, "How important is your reputation in a multi-agent environment," in *Proc. 1999 Conf. Systems, Man, Cybern.*, 1999, pp. 575–580.
- [24] S. Y. Chong and X. Yao, "The impact of noise on iterated prisoner's dilemma with multiple levels of cooperation," in *Proc. Congr. Evol. Comput.*, 2004, pp. 348–355.
- [25] E. A. Stanley, D. Ashlock, and M. D. Smucker, "Prisoner's dilemma with choice and refusal of partners: Evolutionary results," in *Proc. 3d Eur. Conf. Adv. Artif. Life*, 1995, pp. 490–502.
- [26] X. Yao, "Evolving artificial neural networks," in *Proc. IEEE*, vol. 87, Sep. 1999, pp. 1423–1447.
- [27] X. Yao, Y. Liu, and G. Lin, "Evolutionary programming made faster," *IEEE Trans. Evol. Comput.*, vol. 3, no. 2, pp. 82–102, Jul. 1999.
- [28] D. B. Fogel, "An introduction to simulated evolutionary optimization," *IEEE Trans. Neural Netw.*, vol. 5, no. 1, pp. 3–14, Jan. 1994.
- [29] W. Atmar, "Notes on the simulation of evolution," *IEEE Trans. Neural Netw.*, vol. 5, no. 1, pp. 130–147, Jan. 1994.
- [30] X. Yao, "Introduction," *Informatica (Special Issue on Evolutionary Computation)*, vol. 18, pp. 375–376, 1994.
- [31] J. M. McNamara, Z. Barta, and A. I. Houston, "Variation in behavior promotes cooperation in the prisoner's dilemma," *Nature*, vol. 428, pp. 745–748, 2004.



Siang Y. Chong (M'99) received the B.Eng. (Honors) and M.Eng.Sc. degrees in electronics engineering from Multimedia University (MMU), Melaka, Malaysia, in 2002 and 2004, respectively. He is currently working towards the Ph.D. degree in computer science at the University of Birmingham, Birmingham, U.K.

He is a former Tutor for the Faculty of Engineering and Technology, MMU, and is currently a Member of the Natural Computation Research Group, School of Computer Science, University of Birmingham.

His research interests include evolutionary computation, neural networks, and game theory.

Mr. Chong is the recipient of the Student Travel Grant for the 2003 Congress on Evolutionary Computation (CEC03). He is a Member of the Institution of Engineers, Malaysia (IEM).



Xin Yao (M'91–SM'96–F'03) received the B.Sc. degree from the University of Science and Technology of China (USTC), Hefei, in 1982, the M.Sc. degree from the North China Institute of Computing Technology, Beijing, in 1985, and the Ph.D. degree from USTC in 1990.

He was an Associate Lecturer and Lecturer from 1985 to 1990 at USTC, while working towards his Ph.D. He took up a Postdoctoral Fellowship in the Computer Sciences Laboratory, Australian National University (ANU), Canberra, in 1990, and continued

his work on simulated annealing and evolutionary algorithms. He joined the Knowledge-Based Systems Group, CSIRO Division of Building, Construction and Engineering, Melbourne, in 1991, working primarily on an industrial project on automatic inspection of sewage pipes. He returned to Canberra in 1992 to take up a lectureship in the School of Computer Science, University College, University of New South Wales (UNSW), Australian Defence Force Academy (ADFA), where he was later promoted to a Senior Lecturer and Associate Professor. Attracted by the English weather, he moved to the University of Birmingham, Birmingham, U.K., as a Professor of Computer Science in 1999. Currently, he is the Director of the Centre of Excellence for Research in Computational Intelligence and Applications (CERCIA), a Distinguished Visiting Professor of the University of Science and Technology of China, Hefei, and a visiting professor of three other universities. He has more than 200 publications. He is an associate editor or editorial board member of several journals. He is the Editor of the World Scientific Book Series on *Advances in Natural Computation*. He has given more than 35 invited keynote and plenary speeches at conferences and workshops worldwide. His major research interests include evolutionary artificial neural networks, automatic modularization of machine learning systems, evolutionary optimization, constraint handling techniques, computational time complexity of evolutionary algorithms, coevolution, iterated prisoner's dilemma, data mining, and real-world applications.

Dr. Yao was awarded the President's Award for Outstanding Thesis by the Chinese Academy of Sciences for his Ph.D. work on simulated annealing and evolutionary algorithms. He won the 2001 IEEE Donald G. Fink Prize Paper Award for his work on evolutionary artificial neural networks. He is the Editor-in-Chief of the IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION.